



CERGY PARIS  
UNIVERSITÉ



THEMA

théorie économique,  
modélisation et applications

PhD Thesis

# Development of a Dynamic Transport Simulator for Policy Evaluation: Applications to Ride-Sharing and Low Emission Zone in Paris

presented and publicly defended by

**Lucas Javaudin**

on December 9, 2024

prepared under the direction of

**André de Palma**

delivered by

**CY Cergy Paris Université**

## Jury

**Nicolas Coulombel**

Associate Professor (HDR), École des Ponts et Chaussées

Reviewer

**Robin Lindsey**

Professor Emeritus, University of British Columbia

Reviewer

**Moez Kilani**

Professor, Université du Littoral Côte d'Opale

Examiner

**Katheline Schubert**

Professor, Paris School of Economics, University Paris 1 Panthéon-Sorbonne

Jury president

**André de Palma**

Professor Emeritus, CY Cergy Paris Université

Supervisor



# Remerciements / Acknowledgments

Pendant quatre années et quelques mois, mon esprit a suivi une route tantôt express, tantôt sinueuse, avec parfois des détours ou même des demi-tours, et avec aussi bien sûr des arrêts pour se reposer. Même si je n'ai pas pris le chemin le plus court pour réaliser cette thèse, ce fut un voyage formidable !

Tout en remerciant les personnes qui ont parcouru une partie de ce chemin avec moi, je souhaite ici faire une rétrospective de ces années marquées par un enchaînement de passions, d'émotions fortes, de joies, de fatigues, et de rires, avec parfois, il faut l'avouer, des périodes de congestion mentale sévère.

*While the thesis is written exclusively in English*, la plupart des interactions sur mon chemin ont été effectuées en Français, c'est pourquoi je m'autorise ici à parler *sometimes in English* et parfois en Français.

Mon parcours a commencé lorsqu'André de Palma m'a accordé l'opportunité de travailler à l'amélioration du simulateur METROPOLIS, puis m'a fait confiance pour créer une version améliorée, METROPOLIS2, en partant de zéro. Il m'a prévenu des difficultés, qu'il avait lui-même déjà rencontrées, et du risque de passer des mois sans résultats concrets. Néanmoins, il m'a fait confiance et m'a guidé tout au long du chemin. Je suis profondément reconnaissant de sa confiance ainsi que de l'aide qu'il m'a apporté pour rester sur la bonne voie. Je le remercie également de m'avoir impliqué dans des opportunités enrichissantes (et chronophages), notamment avec la Société des Grands Projets, et d'avoir accepté de financer mes dernières années de thèse.

*Along the path, I was fortunate to meet many brilliant researchers with whom I had the pleasure to collaborate and from whom I received valuable advice. Many thanks to Robin Lindsey for his meticulous comments, to Andrea Araldo for his dedication to our shared article, to Nathalie Picard for her insightful guidance, to Nicolas Coulombel for teaching opportunities and for his support; and also to Patrick Stokkink, Dominique Bureau, Moez Kilani, Nadège Blond, Paolo Delle Site, Moshe Ben-Akiva, Guillaume Monchambert. Their guidance ensured that my thesis journey did not go off-road.*

*I thank the four jury members, brilliant researchers who were each a source of inspiration, for accepting to review my thesis at the end of the journey.*

Tout en revenant en arrière sur ce parcours, je tiens à remercier le laboratoire THEMA, qui m'a accueilli durant ce périple. Un grand merci à Éric Danan, directeur du laboratoire et "allocateur de bureaux", ainsi qu'aux administratrices du secrétariat, toujours présentes pour régler les soucis et les voyages.

Si le paysage n'est peut-être pas aussi pittoresque que la Bretagne, Cergy (ou Ygrec pour les intimes) est néanmoins un lieu chaleureux : la base de loisir, les courts de tennis, les balades à vélo dans le Vexin, les restaurants, le cinéma. . .

Au sein du laboratoire, j'ai eu la chance d'être entouré de merveilleux doctorants. Il y a ceux qui m'ont montré la voie : Éloïse, Gabrielle et Margaux ; ceux qui arrivent dans le même wagon : Mélanie, Anderson et Romuald ; et ceux qui sont dans le rétroviseur pleins d'encouragements : Ben, Aurélie et Ivan. Ces "collègues" ont été un carburant inestimable pour continuer à avancer.

J'ai également eu le plaisir de travailler et d'échanger avec des doctorants et étudiants utilisateurs ou contributeurs de METROPOLIS2: Filippo, Pouingnè, Ismail, Ahmed, Théotime, Joseph et Youssef. *With a special thanks to Samarth, with whom I shared an office, offering us the best view on Cergy's worst road intersection, for our passionate discussions on congestion modeling, that helped me navigate through my own congestion of ideas.*

*Teaching how to use the METROPOLIS2 simulator was an enriching and exhausting*

*adventure that made my work feel concrete. Talking alone is not recommended, so thanks to the students of all ages who listened patiently as I talked passionately for seven classes.* J'ai également apprécié former des étudiants motivés de l'École des Ponts et Chaussées à l'utilisation de METROPOLIS2.

*I would also like to thank two people who played an important role in this thesis, even though I never met them: Fabrice Marchal and Yurii Nesterov who, together with André de Palma, developed the METROPOLIS simulator, first of the name, which served as a major inspiration for my work.*

Pour mieux travailler, les moments de repos (intellectuel) sont indispensables ! Le tennis a toujours été au rendez-vous, même si le travail n'était jamais loin, que ce soit en jouant avec d'autres doctorants ou en affrontant un certain chercheur renommé en économie des transports.

Je n'oublie pas, bien sûr, les amis qui étaient déjà là au départ du voyage, m'ont encouragé tout au long du chemin et me félicitent à l'arrivée : Eoghan, Robin, Mickaël, Margot.

Le voyage nous amène parfois à traverser des épreuves difficiles. Je dédie cette thèse à mon papa, avec qui j'aurais aimé partager tellement plus d'expériences. Je remercie infiniment ma maman et ma sœur ; l'amour que nous partageons est plus forte que les épreuves que nous affrontons.

Je remercie également toute ma famille, qui m'a accompagné tout au long de mon parcours, avec une pensée particulière pour Robin, Béatrice et Titouan, toujours soucieux de prendre des nouvelles, et pour ma mamie, qui pourra enfin arrêter de me demander si j'ai réussi mes examens...

Enfin, même si Rust n'est pas mal non plus, la plus belle découverte de cette thèse reste Gabrielle. Merci de m'avoir soutenu et aidé sans interruption. *I hope to share many more scenic routes with you. . .*

# Résumé

Cette thèse explore l'utilisation des simulations dynamiques de transport pour évaluer les politiques de mobilité urbaine, en se concentrant sur le covoiturage et la Zone à Faibles Émissions dans la région métropolitaine de Paris. Elle s'appuie sur deux simulateurs de transport : METROPOLIS1, déjà existant, et METROPOLIS2, un simulateur développé dans le cadre de cette thèse. Les deux simulateurs utilisent la théorie du choix discret pour estimer les réactions des usagers aux changements de politique et évaluer les impacts sur leur utilité. En utilisant un cadre mésoscopique, ces simulateurs modélisent efficacement la congestion routière à l'échelle urbaine et régionale. METROPOLIS2 améliore son prédécesseur en permettant d'évaluer une gamme plus large de politiques, telles que les restrictions basées sur les véhicules, avec une efficacité et une précision accrues.

Le premier chapitre examine un schéma de covoiturage avec des conducteurs inflexibles qui conservent des heures de départ et des itinéraires fixes, qu'ils aient des passagers ou non. METROPOLIS1 est utilisé pour modéliser les choix d'itinéraires et d'heures de départ de chaque usager, ainsi que les niveaux de congestion en résultant. L'appariement optimal entre conducteurs et passagers, ainsi que les points de prise en charge et de dépose, est obtenu en résolvant un problème de programmation linéaire en nombres entiers. Les simulations des trajets du matin dans Paris montrent que, même avec une faible participation, le covoiturage peut réduire la congestion, la consommation de carburant et les émissions de CO<sub>2</sub>. Des avantages supplémentaires peuvent être obtenus en augmentant la capacité des véhicules ou en offrant des incitations financières, sans compromettre l'inflexibilité des conducteurs.

Le deuxième chapitre présente METROPOLIS2, un simulateur mésoscopique à base d'agent, capable de modéliser les décisions de voyage (mode, heure de départ et itinéraire) en se basant sur la théorie du choix discret dans un cadre dynamique et en temps continu. Ce simulateur améliore METROPOLIS1 en intégrant des chaînes de trajets, plusieurs types de véhicules, une plus grande flexibilité dans la spécification des utilités, etc. Son efficacité et sa précision sont validées par deux études de cas : la première réplique des résultats analytiques du modèle dit *bottleneck* à une route ; la seconde démontre une meilleure vitesse et convergence par rapport à son prédécesseur dans un scénario à grande échelle de la région parisienne.

Le troisième chapitre applique METROPOLIS2 pour évaluer la Zone à Faibles Émissions (ZFE) de Paris. Des données ouvertes sont utilisées pour générer une population synthétique et un réseau routier, tandis que des techniques d'apprentissage automatique (comme la régression Lasso et l'optimisation bayésienne) sont utilisées pour calibrer la simulation et reproduire les temps de trajet et les comportements observés. L'analyse évalue les effets de la ZFE sur la qualité de l'air, la congestion et les inégalités, mettant en évidence les bénéfices pour les résidents du centre-ville, mais révélant des impacts néfastes potentiels pour les populations suburbaines dépendantes de véhicules plus anciens.

# Summary

This thesis investigates the use of dynamic transport simulations to evaluate urban mobility policies, focusing on ride-sharing and Low Emission Zone in the Paris metropolitan area. It relies on two transport simulators: the existing METROPOLIS1 and METROPOLIS2, a novel simulator developed as part of this PhD. Both simulators employ discrete-choice theory to estimate responses of commuters to policy changes and assess the impact on utility. Using a mesoscopic framework, the simulators efficiently model road congestion at the city and regional scales. METROPOLIS2 builds upon its predecessor, enabling the evaluation of a broader range of policies, such as vehicle-based restrictions, with enhanced efficiency and accuracy.

The first chapter explores a ride-sharing scheme with inflexible drivers who maintain fixed departure times and routes regardless of whether they carry passengers. METROPOLIS1 is used to model the departure time and route chosen by each commuter, as well as the resulting congestion levels. The optimal matching of drivers with passengers, as well as the pick-up and drop-off points, are obtained by solving an integer linear programming problem. Simulations of Paris's morning commute demonstrate that, even with low participation, ride-sharing can reduce congestion, fuel consumption, and CO<sub>2</sub> emissions. Additional benefits can be achieved by increasing vehicle capacity or offering monetary incentives, without compromising driver inflexibility.

The second chapter introduces METROPOLIS2, a mesoscopic agent-based transport simulator capable of modeling travel decisions (mode, departure time, and route) based



on discrete-choice theory within a dynamic, continuous-time framework. The simulator improves upon its predecessor METROPOLIS1 by incorporating trip chaining, multiple vehicle types, greater flexibility in utility specification, etc. Its efficiency and accuracy are validated through two case studies: replicating analytical results from the standard single-road bottleneck model and demonstrating superior speed and convergence compared to its predecessor on a large-scale scenario of the Paris region.

The third chapter applies METROPOLIS2 to evaluate the Low Emission Zone (LEZ) in Paris. Open data are used to generate a synthetic population and road network, while machine-learning techniques (such as Lasso regression and Bayesian optimization) calibrated the simulation to replicate observed travel times and behaviors. The analysis assesses the LEZ's effects on air quality, congestion and inequalities, highlighting benefits for city-center residents but revealing potential disadvantages for suburban populations dependent on older vehicles.

# Contents

Remerciements / Acknowledgments	1
Résumé	4
Summary	6
List of Figures	11
List of Tables	13
General Introduction	15
<b>1 Ride-Sharing</b>	<b>33</b>
1.1 Introduction . . . . .	35
1.2 Literature Review . . . . .	38
1.3 Methodology . . . . .	40
1.3.1 Ride-sharing Scheme . . . . .	40
1.3.2 Four-Step Procedure . . . . .	42
1.3.3 Traffic Simulator . . . . .	43
1.3.4 Ride-sharing Cost . . . . .	45
1.3.5 Matching . . . . .	46
1.4 Case Study . . . . .	48
1.4.1 Network Modelling . . . . .	48
1.4.2 Travel Demand . . . . .	49
1.4.3 Results . . . . .	52
1.4.4 Detailed Results for the 30 % Scenario . . . . .	54
1.4.5 Allowing Multiple Riders in the Same Car . . . . .	58
1.4.6 Proposing Incentives to the Riders . . . . .	59
1.5 Conclusion . . . . .	60
<b>2 METROPOLIS2</b>	<b>64</b>
2.1 Introduction . . . . .	66
2.2 Literature Review . . . . .	68
2.3 From Models to Simulations . . . . .	75
2.3.1 The Bottleneck Model . . . . .	75
2.3.2 Challenges of Simulation . . . . .	77

2.3.3	Simulating the Bottleneck Model . . . . .	78
2.4	Simulator Overview . . . . .	86
2.5	Implementation . . . . .	93
2.6	Comparison with METROPOLIS1 . . . . .	94
2.7	Conclusion . . . . .	102
2.A	Numerical Application . . . . .	104
2.A.1	Definitions and Notations . . . . .	104
2.A.2	Comparison to the Analytical Results . . . . .	107
2.A.3	Sensitivity Analysis . . . . .	112
2.A.4	Considering Heterogeneity . . . . .	117
2.B	Input Data . . . . .	119
2.B.1	Road-Network Input . . . . .	119
2.B.2	Agents . . . . .	120
2.B.3	Parameters . . . . .	121
2.C	Output Data . . . . .	122
2.C.1	Agent Output . . . . .	122
2.C.2	Road-Network Output . . . . .	122
2.C.3	Aggregated Results . . . . .	122
2.D	Network Conditions . . . . .	123
2.E	Description of the Models . . . . .	123
2.E.1	Demand Model . . . . .	123
2.E.2	Supply Model . . . . .	126
<b>3</b>	<b>Low Emission Zone</b> . . . . .	<b>127</b>
3.1	Introduction . . . . .	129
3.2	Literature Review . . . . .	132
3.2.1	Calibration of agent-based transport simulators . . . . .	132
3.2.2	Evaluation of Low Emission Zones . . . . .	133
3.3	Theoretical Foundations . . . . .	136
3.3.1	Demand Side . . . . .	136
3.3.2	Supply Side . . . . .	139
3.4	Data Input . . . . .	140
3.4.1	Supply Side . . . . .	140
3.4.2	Demand Side . . . . .	142
3.4.3	Mode-specific Processing . . . . .	145
3.5	Calibrating the Model . . . . .	147
3.5.1	Free-Flow Travel Times Calibration . . . . .	148
3.5.2	Congested Travel Times Calibration . . . . .	150
3.5.3	Departure-Time Distribution Calibration . . . . .	154
3.5.4	Mode Shares Calibration . . . . .	157
3.6	Low Emission Zone Policy . . . . .	164
3.6.1	Crit'Air Classification System . . . . .	165
3.6.2	Simulating the LEZ . . . . .	167
3.6.3	Air Quality Impact . . . . .	168
3.6.4	Expected Results . . . . .	169

3.6.5	Limits of the Methodology . . . . .	170
3.7	Results . . . . .	171
3.7.1	Convergence . . . . .	171
3.7.2	Global results . . . . .	172
3.7.3	Pollution Results . . . . .	177
3.7.4	Spatial and Economic Inequalities . . . . .	183
3.8	Conclusion . . . . .	189
<b>General Conclusion</b>		<b>193</b>
<b>Bibliography</b>		<b>197</b>

# List of Figures

1.1	Ride-sharing trip example . . . . .	41
1.2	Île-de-France road network . . . . .	49
1.3	Desired arrival time distribution of the four traveller groups . . . . .	50
1.4	Mode shifts from the reference scenario (left) to the 30 % participation rate scenario (right) . . . . .	54
1.5	Origins and destinations of matched riders in Île-de-France . . . . .	55
1.6	Distribution of walking distance for the matched riders . . . . .	56
1.7	Distribution of schedule delay for the riders . . . . .	57
1.8	Distribution of the decrease of the riders' generalised cost, compared to the reference scenario . . . . .	57
2.1	Illustration of the road network in the bottleneck model . . . . .	75
2.2	Analytical solution of the bottleneck model . . . . .	76
2.3	Illustration of the use of inverse sampling to draw a departure time from a cumulative distribution function . . . . .	80
2.4	Agents' travel times and interpolated travel-time function (with $N = 200$ and $\delta = 4$ minutes) . . . . .	82
2.5	Overview of METROPOLIS2 process . . . . .	88
2.6	Bottleneck speed-density function on an example road . . . . .	95
2.7	Comparison of the convergence of departure times over iterations, between METROPOLIS1 and METROPOLIS2 . . . . .	99
2.8	Comparison of the convergence of travel-time expectation error over iterations, between METROPOLIS1 and METROPOLIS2 . . . . .	100
2.9	Comparison of the departure-time distribution in the Paris simulation with METROPOLIS1 and METROPOLIS2 . . . . .	101
2.10	Comparison of the travel-time distribution in the Paris simulation with METROPOLIS1 and METROPOLIS2 . . . . .	101
2.11	Comparison of the link-level flows for a Paris simulation with METROPOLIS1 and METROPOLIS2 . . . . .	102
2.12	Analytical result for the rate of departures from origin . . . . .	106
2.13	Analytical result for the travel-time function . . . . .	106
2.14	Comparison of the rate of departures from origin in the simulation and in the analytical results . . . . .	108
2.15	Comparison of the travel-time function from origin to destination in the simulation and in the analytical results . . . . .	109

2.16	Convergence of $\text{RMSE}_{\kappa}^{\text{dep}}$ and $\text{RMSE}_{\kappa}^T$ . . . . .	109
2.17	Comparison of the rate of departures when averaging over 10 simulations . . .	111
2.18	Comparison of the rate of departures when using systematic sampling . . . .	111
2.19	Impact of the smoothing factor on $\text{RMSE}_{\kappa}^{\text{dep}}$ convergence . . . . .	113
2.20	Impact of the number of agents on $\text{RMSE}_{\kappa}^{\text{dep}}$ convergence . . . . .	114
2.21	Impact of the breakpoint interval on $\text{RMSE}_{\kappa}^{\text{dep}}$ convergence . . . . .	115
2.22	Travel-time function for different values of the breakpoint interval . . . . .	115
2.23	Probability distribution of the departure-time choice for different values of the scale of the utility's random component (assuming a constant travel time) . .	116
2.24	Impact of the scale of the utility's random component on $\text{RMSE}_{\kappa}^{\text{dep}}$ convergence	117
2.25	Distribution of the $t_n^*$ values for the simulations with observed heterogeneity	118
2.26	Comparison of the rate of departures between homogeneous and heterogeneous simulations . . . . .	119
2.27	Structure of an agent . . . . .	120
2.28	Structure of a trip chain . . . . .	120
2.29	Timing dynamics of a trip chain with two trips . . . . .	121
3.1	Example activity pattern of an agent with two tours . . . . .	137
3.2	Road segments imported from OpenStreetMap, near Cergy . . . . .	141
3.3	Process to assign an origin / destination node from activity coordinates . . .	145
3.4	Density of OD-level free-flow travel times (TomTom, calibrated, uncalibrated)	151
3.5	TomTom vs calibrated congested travel times at the OD-level . . . . .	154
3.6	Cumulative distribution of departure time, for the 18 clusters . . . . .	155
3.7	Comparison between the overall departure-time distribution in the travel sur- vey and in METROPOLIS2 after calibration . . . . .	157
3.8	Comparison between the observed and simulated departure-time distribution for 4 clusters . . . . .	158
3.9	Comparison between observed and simulated mode shares over all clusters . .	163
3.10	Comparison between global observed and simulated mode shares . . . . .	164
3.11	Area of the Low Emission Zone within Île-de-France . . . . .	165
3.12	Evolution of Crit'Air category shares in the Île-de-France vehicle fleet . . . .	166
3.13	Share of cars with Crit'Air 3 or worse at the municipality level for the 2025 vehicle fleet . . . . .	167
3.14	Variation in daily average travel times at the road-segment level . . . . .	176
3.15	Variation of daily flows on public-transit segments . . . . .	177
3.16	Vehicle-kilometers by Crit'Air category . . . . .	179
3.17	Daily $\text{PM}_{2.5}$ emissions from car trips on a 500 m grid . . . . .	180
3.18	$\text{PM}_{2.5}$ concentrations on a 500 m grid . . . . .	181
3.19	Day-average population density on a 500 m grid . . . . .	182
3.20	Daily exposure cost to $\text{NO}_2$ and $\text{PM}_{2.5}$ pollutants on a 500 m grid . . . . .	183
3.21	Cumulative distribution function of the travel and health surplus (agent-level)	184
3.22	Share of LEZ winners by municipality . . . . .	188
3.23	Share of LEZ losers by municipality . . . . .	188
3.24	Share of LEZ winners as a function of the municipality mean income . . . .	189
3.25	Share of LEZ losers as a function of the municipality mean income . . . . .	190

# List of Tables

1.1	Preference parameters for the four groups of travellers, in €/h . . . . .	51
1.2	Comparison of results for the reference scenario and five scenarios with a different participation rate . . . . .	53
1.3	Average subsidy and cost of the policy as a function of the inconvenience cost for the drivers . . . . .	58
1.4	Comparison of results when drivers can have at most 1, 2 or 3 passengers in their car, for 30 % of people willing to participate in the ride-sharing scheme	59
1.5	Comparison of results with different incentive amount, for 30 % of people willing to participate in the ride-sharing scheme (at most 1 passenger per car)	61
2.1	Comparison of 3 transport simulators: General and software . . . . .	69
2.2	Comparison of 3 transport simulators: Demand definition . . . . .	70
2.3	Comparison of 3 transport simulators: Choice models . . . . .	71
2.4	Comparison of 3 transport simulators: Supply . . . . .	73
2.5	Main notations used in METROPOLIS2 . . . . .	87
2.6	Comparison of the aggregate results for a Paris simulation with METROPOLIS1 and METROPOLIS2 . . . . .	98
2.7	Default values for the parameters of the model . . . . .	105
2.8	Results of running METROPOLIS2 to simulate the bottleneck model . . . .	108
2.9	Average, minimum and maximum values for the results with 10 different random seeds . . . . .	110
2.10	Sensitivity of the results to the smoothing factor . . . . .	112
2.11	Sensitivity of the results to the number of agents . . . . .	113
2.12	Sensitivity of the results to the breakpoint interval $\delta$ . . . . .	114
2.13	Sensitivity of the results to the scale of the utility's random component . . .	116
2.14	Comparison between the homogeneous and heterogeneous bottleneck models	118
3.1	Summary of the calibration process . . . . .	148
3.2	Calibrated capacities by road type (in PCE per hour per lane) . . . . .	153
3.3	Calibrated parameter values for departure-time choice . . . . .	157
3.4	Calibrated mode utility constants . . . . .	161
3.5	Calibrated mode-specific values of time . . . . .	162
3.6	Calibrated values for the fuel factor and implicit car value of time . . . . .	163
3.7	Crit'Air categories . . . . .	166
3.8	Statistics on the convergence of the simulations . . . . .	172

3.9	Tour characteristics by car ownership . . . . .	172
3.10	Measures of effectiveness for the simulated average day . . . . .	174
3.11	Transition matrix of modes from baseline to LEZ . . . . .	175
3.12	Pollution-related results . . . . .	178
3.13	Basic benefit-cost analysis of the LEZ policy . . . . .	184
3.14	Travel and health surplus by population group . . . . .	185
3.15	Characteristics of the agents and their trips by population group . . . . .	186
3.16	Mode shares in the baseline and LEZ scenario for the LEZ winners and losers	187



# General Introduction

## Context and Transportation Policies

Transportation plays an essential role in modern societies, enabling the movement of people to workplaces, facilitating access to essential services, and connecting both goods and individuals to commercial areas.

However, the transportation sector is currently facing several challenges, particularly in the context of sustainability. One major issue is its environmental impact – transport accounts for nearly a quarter of global CO<sub>2</sub> emissions, making it a key contributor to climate change.<sup>1</sup> In addition to its ecological footprint, transportation also poses serious public health risks. Road transport is responsible for 37% of nitrogen oxides emissions, pollutants that contribute to chronic respiratory illnesses and cause approximately 40 thousand premature deaths annually across Europe.<sup>2</sup> Furthermore, noise pollution from road traffic is a growing concern, with the World Health Organization estimating that at least 1.6 million healthy years of life are lost due to exposure to road traffic noise in Europe.<sup>3</sup>

Congestion is another important issue, particularly in large urban centers like Paris, where drivers lost an average of 97 hours in traffic in 2023.<sup>4</sup> This congestion has major

---

<sup>1</sup>Source: IEA (2020), Global energy-related CO<sub>2</sub> emissions by sector, IEA, Paris <https://www.iea.org/data-and-statistics/charts/global-energy-related-co2-emissions-by-sector>, Licence: CC BY 4.0.

<sup>2</sup>Sources: <https://www.eea.europa.eu/publications/air-quality-in-europe-2021/health-impacts-of-air-pollution> [accessed 2024/09/11]; <https://www.eea.europa.eu/publications/air-quality-in-europe-2022/sources-and-emissions-of-air> [accessed 2024/09/11].

<sup>3</sup>Source: World Health Organization. (2018). Environmental noise guidelines for the European region.

<sup>4</sup>Source: INRIX 2023 Traffic Scorecard Report <https://inrix.com/scorecard/> [accessed 2024/10/21].

economic and social impacts, as lost time leads to reduced productivity and increased stress for commuters (Fattah et al., 2022). Additionally, road accidents in the Paris metropolitan area resulted in 268 fatalities and 19 457 injuries in 2023, highlighting the larger public-health and safety problems related to transportation.<sup>5</sup>

In response to these challenges, cities and governments around the world are implementing various policies aimed at mitigating the adverse effects of transportation. These policies generally fall into one of four categories: investments, restrictions, incentives, and price-based policies. Investment policies focus on developing transportation infrastructure by constructing new roads or public-transit lines, or by expanding the capacity of existing infrastructure. Restrictions policies are regulations intended to influence individual behavior. Examples include the European Union’s scheduled ban on the sale of internal combustion engine vehicles by 2035, or high-occupancy vehicle (HOV) lanes, which restrict certain traffic lanes to vehicles carrying at least one passenger (in addition to the driver). Incentive policies offer subsidies or nudges to encourage individuals to adopt more sustainable behaviors, such as subsidies for purchasing electric vehicles or incentives to commute by car-pooling. Finally, price-based policies use price signals to influence behaviors. Road tolls and public-transit fares are example of policies that adjust costs to influence travel choices.

In Paris and its metropolitan area, which serves as the main area of study for this thesis, a Low Emission Zone was introduced in 2015 to restrict the entry of the most polluting vehicles into certain parts of the region. This policy aims primarily to improve air quality while also reducing congestion, and limiting CO<sub>2</sub> emissions and noise pollution.

In addition, Paris has taken further steps to address transportation-related issues. The speed limit on the *Boulevard Périphérique*, a major highway surrounding the city, was reduced from 70 km/h to 50 km/h in an effort to reduce vehicle emissions and noise. The Île-de-France region has also introduced subsidies to encourage car-pooling, aiming to decrease the number of vehicles on the road by promoting shared journeys, which can help reduce

---

<sup>5</sup>Source: DRIEAT Île-de-France <https://www.drieat.ile-de-france.developpement-durable.gouv.fr/barometres-mensuels-de-l-accidentalite-routiere-d-a11231.html> [accessed 2024/10/28].

congestion and lower overall pollution levels. Finally, a major project currently reshaping the Paris metropolitan area is the *Grand Paris Express*, an ambitious initiative involving the construction of four new métro lines, adding 200 kilometers of track and 68 new stations. By creating an extensive and modernized public transport network, the Grand Paris Express is expected to induce a significant modal shift, encouraging commuters to switch from private cars to public transit, with in turn could substantially reduce congestion and pollution.

## Individual Transport-Related Decisions

Before implementing policies such as those described, it is crucial to evaluate their potential impacts on travel times and air quality, for example, not only on a global scale but also on specific population segments (Donais et al., 2019). This task is particularly challenging, as it requires a deep understanding of how individuals make transport-related decisions. These decisions vary significantly in scope, ranging from short-term to long-term considerations, and they can be influenced by numerous factors.

In the short term, individuals must decide how to manage their daily commutes. For example, when driving, they must determine the optimal route to take, considering current and anticipated congestion levels (Ben-Akiva and Bierlaire, 1999). Similarly, when using public transit, passengers must decide which itinerary to follow based on real-time disruptions and passenger flows. Timing is also a key factor: commuters choose when to depart in order to reach their destination on time, given the unpredictability of traffic or transit conditions (Small, 1982; Lemp et al., 2010). Additionally, they face the decision of which mode of transport to use – whether to drive, take public transit, cycle, or walk – depending on travel times, costs, and personal preferences (Train and McFadden, 1978; Miller et al., 2005).

In the medium term, individuals choose which activities to perform throughout the days and weeks – working, shopping, or engaging in leisure activities – and where to perform them, taking into account time constraints and travel times between locations with different opportunities (Rasouli and Timmermans, 2014).

In the long term, individuals make decisions about which vehicle to purchase – bicycle, electric car, or traditional car – and whether to subscribe to services such as public transit or bike-sharing (Jong et al., 2004). These choices depend on anticipated travel needs over the coming months or years, as well as the available infrastructure (e.g., roads, cycleways, transit systems) and policies (e.g., vehicle restrictions in certain zones).

Over an even longer horizon, individuals also face major life decisions, such as whether and where to work, and where to live (Waddell et al., 2007). These decisions often involve balancing factors such as proximity to job opportunities, accessibility to public transit, and the availability of amenities.

Public policies, including transportation-related measures, can influence many of these decisions. For instance, the introduction of a Low Emission Zone may prompt individuals to reconsider their vehicle choice – particularly if their current vehicle is banned within the zone. It can also affect mode choice, encouraging a switch to public transit if driving is no longer viable, and alter route choice by forcing drivers to reroute around restricted areas. Thus, understanding how people make these complex, interrelated decisions is essential for predicting how they will react to new policies.

However, predicting individual responses to policy interventions is complicated by the fact that these decisions are often not made in isolation. Individuals' choices are interdependent, both at the household level and on a broader regional scale.

At the household level, decision-making is often shared among members. For example, vehicles are typically used by multiple individuals, meaning that vehicle availability must be coordinated (Picard et al., 2018). Similarly, certain activities, such as grocery shopping or leisure activities, are frequently shared or coordinated within households (Ho and Mulley, 2013). One person might shop for the entire household, or families may plan leisure activities together. These internal agreements mean that individual transportation decisions are rarely independent.

On a larger scale, transportation infrastructure is shared across the entire population,

which introduces further complexity. Roads, public transit systems, and other networks are used by multiple individuals, leading to phenomena such as congestion. The infrastructures can even be shared with freight transport or public services. When too many vehicles use the same road, they slow each other down, and when too many passengers board the same bus or train, overcrowding can occur, forcing some passengers to stand or even wait for the next vehicle. These shared infrastructures and the collective behavior of individuals create feedback loops that influence transport decisions, complicating the task of predicting responses to new policies.

This thesis positions itself within the field of transportation economics, a branch of economics that seeks to understand how individuals make transport-related decisions and the interdependencies between these decisions. Closely linked to urban, labor, and public economics, this field demonstrates that long-term decisions like where to live, where to work, or which vehicle to buy are shaped by anticipated short-term needs, such as commute routes, while short-term choices are often limited by previous long-term commitments. Transport-related decisions are thus deeply interrelated and cannot typically be studied in isolation, highlighting the need for robust tools to capture these interdependencies.

## **Transport Simulators**

When evaluating transportation policies at the city or regional scale, a common approach in transportation economics is to use transport simulators. These models are capable of predicting the decisions made by large populations of individuals interacting within road and public-transit networks. Generally, they focus on short-term decisions, such as route, departure time, and mode. However, they can be integrated with activity-based models (Rasouli and Timmermans, 2014), car ownership models (Jong et al., 2004), or land-use models (Waddell, 2002) to account for longer-term decisions and their broader impacts.

The four-step model (McNally, 2007) is a classical approach to estimate decisions made by individuals given factors such as the road network, public-transit systems, built environment,

and public policies. This model consists of four stages: trip generation, where the number of trips originating and arriving in each zone is estimated; trip distribution, which calculates the flow of trips between origin and destination zones; modal choice, which predicts the proportion of trips using each mode of transport; and finally, route assignment, where the specific routes for the trips are determined.

While this model has been extensively applied, it does have notable limitations. One key drawback is that it does not account for all relevant decisions, such as departure-time choice, which is not modeled explicitly. Instead, the model aggregates trips over a given time period, with congestion levels treated as constant throughout the period. Another limitation is its reliance on aggregated flows, meaning the model struggles to accurately assess the impact of a policy on specific population segments. Although the four-step model can sometimes divide flows by population group, it remains limited in how much individual heterogeneity it can capture.

With advances in computing power, more sophisticated simulation-based transport models have emerged. These transport simulators can be classified into various categories, depending on how precisely they model the demand side (the population making decisions) and the supply side (the transport infrastructure). In contrast to four-step models, which describe demand through aggregated flows, newer models – referred to as agent-based models – represent each individual, or “agent”, separately (Huang et al., 2022). This approach allows for greater flexibility in accounting for individual differences. However, in practice, due to computational constraints or limited data, individuals are still often grouped into homogeneous classes.

On the supply side, four-step models are classified as macroscopic models, meaning they represent flows of vehicles between origins and destinations, using aggregate congestion laws and in a static way. At the other extreme, microscopic models simulate each vehicle independently, incorporating detailed driving behaviors, such as lane changes and overtaking, as well as a fine-grained representation of network infrastructure, including traffic lights and

lane restrictions. An example of a microscopic model is SUMO (Lopez et al., 2018). Due to their complexity, microscopic models are generally limited to smaller urban areas or city quarters, and they also often lack detailed demand-side modeling, assuming fixed modes and departure times. Consequently, these models are more commonly used in civil engineering than in transportation economics.

Between these two extremes are mesoscopic models, which represent individual vehicles but use aggregate congestion laws at the link (or road) level, which simplifies the computation while still capturing vehicle-level interactions. This balance allows mesoscopic models to simulate city- or region-size networks within a reasonable time. Their ability to model congestion at the vehicle level makes them particularly well-suited for integration with agent-based demand models, making them effective tools for evaluating public policies at the city or regional scale.

Unlike static macroscopic models, which assume constant conditions, microscopic and mesoscopic models are typically dynamic, meaning they are able to simulate traffic flows over time, often using steps as small as one second. This allows them to account for changes in congestion as vehicles interact with one another and with the transportation network, providing a more detailed and accurate representation of traffic conditions.

An example of mesoscopic, agent-based transport simulator, used in the first chapter of this thesis, is METROPOLIS (de Palma et al., 1997). On the demand side, METROPOLIS simulates a population of agents, grouped into segments of homogeneous preferences, making decisions regarding mode, departure time, and route. All agents make a single trip, from an origin zone to a destination zone. Mode choice is modeled using a binary Logit model, with a choice between private car and public transit, where the latter is assumed to have constant travel times. Departure-time choice is represented by a Continuous Logit model (Ben-Akiva and Watanatada, 1981), using the alpha-beta-gamma preferences (Vickrey, 1969; Arnott et al., 1990). These preferences reflect agents' trade-offs when choosing their departure time over a continuous time period, balancing the desire to reduce travel time with the need to

arrive at their destination on time. Lastly, route choice is modeled dynamically, with agents making decisions at each intersection to minimize their expected arrival time by choosing the optimal next road.

On the supply side, METROPOLIS uses an event-based model to simulate the chronological sequence of all events occurring on the road network, such as vehicles departing from an origin or reaching an intersection. Congestion is represented by speed-density functions, which reduce vehicle speed as road density increases. The primary speed-density function used in METROPOLIS is designed to simulate bottleneck congestion (Vickrey, 1969). In this case, the flow of vehicles is constrained by road capacity, with waiting times proportional to the number of vehicles queued and inversely proportional to the road's capacity. Additionally, the model accounts for the limited storage capacity of roads – cars are prevented from entering roads that are fully occupied.

## **Motivations and Outline of the Thesis**

This thesis is motivated by the need for tools that can evaluate transportation policies at a detailed level, particularly in metropolitan areas like Paris where challenges such as congestion, air pollution, and sustainability are increasingly prominent. By developing and applying a dynamic transport simulator, this research aims to analyze the effects of specific policies, with a focus on ride-sharing and Low Emission Zones (LEZs), to better understand their potential impacts at both a global and individual scale.

The thesis is organized as follows. Chapter 1 explores a ride-sharing scheme in Paris metropolitan area using the METROPOLIS simulator. Chapter 2 details the development of the METROPOLIS2 simulator, including its architecture and improvements over other simulators. Chapter 3 applies the simulator to analyze the impact of the Low Emission Zone in Paris, providing insights into both the global and individual impacts of the policy. Finally, the General Conclusion summarizes the thesis contributions and suggests areas for future research to expand transport simulation models.



## Chapter 1: Ride-sharing with Inflexible Drivers in the Paris Metropolitan Area<sup>6</sup>

### Motivations

The first chapter of this thesis uses the METROPOLIS simulator to evaluate the potential of a specific ride-sharing system, with a case study focusing on the Paris metropolitan area. In this context, ride-sharing refers to the practice where two or more individuals share a car for their journey. At the individual level, ride-sharing offers the benefit of splitting fuel costs, while on a broader scale, it has the potential to reduce congestion and pollution by lowering the number of vehicle-kilometers traveled. However, for ride-sharing to be successful, the individuals sharing the ride must have similar destinations (spatial match) and compatible schedules (temporal match), as important detours or delays can undermine the benefits (Liu et al., 2020).

Typically, ride-sharing occurs within households, companies, or through dedicated platform where drivers offer their services, and passengers select a suitable option on a first-come, first-served basis. However, several challenges limit the broader adoption of ride-sharing. First, drivers are often unwilling to deviate from their original routes or alter departure times to accommodate passengers' needs. Second, passengers usually pay for their ride and may struggle to find an optimal itinerary. Lastly, the first-come, first-served matching system is suboptimal, often failing to provide the best possible match for passengers.

To address these limitations, this chapter proposes a hypothetical ride-sharing system that could be implemented by governments with three key characteristics. First, drivers are not required to change their route or departure time and may receive a small subsidy, making them no worse off than if they were traveling alone. Second, passengers do not pay for the ride, which makes the option more attractive. Finally, all drivers' and passengers' requests are collected in advanced, allowing for an optimal match to be determined. This

---

<sup>6</sup>In collaboration with André de Palma, Patrick Stokkink (TU Delft) and Léandre Tarpin-Pitre.

ensures that passengers are offered better options (in terms of travel time, walking time, and schedule) than what would be available under a first-come, first-served system.

## **Methodology**

To assess the potential of this ride-sharing system, a specific methodology is proposed. After collecting the routes and departure times of all potential drivers, as well as passengers' preferences, the methodology computes the utility (or generalized cost) for passengers for each potential driver match. This process involves determining the optimal pick-up and drop-off points along the driver's route to maximize the passenger's utility, taking into account travel time in the vehicle, walking time, and schedule-delay costs (the cost of arriving early or late compared to a desired arrival time).

An Integer Linear Programming problem is then solved to identify the matches between drivers and passengers that maximize total individual utility. This optimization ensures that passengers are only matched if the ride-sharing option increases their utility and that drivers are not matched with more passengers than a predefined limit.

## **Results**

The methodology is applied to a case study in the Paris metropolitan area using simulations from METROPOLIS. In the main scenario, where 30% of the population is assumed to be willing to participate to the ride-sharing system, the results show a reduction in the share of car trips from 74.5% to 72.4%, with 3.3% of the population matched as ride-sharing passengers. This leads to a 1.9% decrease in vehicle-kilometers traveled and CO<sub>2</sub> emissions. The quality of the matches is high, with passengers experiencing an average walking time of less than five minutes.

Further analysis suggests that the system's benefits could be enhanced by offering incentives to passengers to accept lower-quality options or by allowing more than one passenger per driver.

## Contributions

This chapter presents a methodology for assessing ride-sharing systems with inflexible drivers, considering the spatial and temporal matching between drivers and passengers. Through an application to the Paris metropolitan area, it explores the potential popularity of this system and evaluates its benefits in reducing traffic congestion and CO<sub>2</sub> emissions.

Building on the work of de Palma et al. (2022b), this chapter implements the matching optimization they propose, extending it to a large-scale dynamic application with endogenous congestion.

While Coulombel et al. (2019) also use a transport simulator to evaluate ride-sharing potential in the Paris metropolitan area, this chapter extends their work by incorporating an actual matching process between drivers and passengers. In contrast, their approach does not use a matching algorithm and assumes instead that ride-sharing merely increases vehicle occupancy.

## Limits

One major limitation of this study is its exclusive focus on morning home-to-work commutes. This narrow scope introduces two specific challenges: (i) trip-chaining constraints, such as dropping off children at school before work, are not modeled; and (ii) passengers are not guaranteed a high-quality match for the evening return trip, as the driver's evening schedule may not align with theirs. These simplifications facilitate the matching process compared to real-world conditions, where passengers would need viable ride-sharing options for both morning and evening commutes.

Furthermore, since METROPOLIS simulates trips between zones, accurately estimating walking times for passengers matched with drivers within the same zone is challenging. Finally, when simulating the impact of ride-sharing on congestion, drivers' departure times and routes may diverge from those assumed during matching, due to METROPOLIS's inability to impose specific departure times and routes.

## Chapter 2: METROPOLIS2: Bridging Theory and Simulation in Agent-Based Transport Modeling<sup>7</sup>

### Motivations

As pointed out in the first chapter, the scope of the METROPOLIS simulator limits the policies that can be evaluated. On the demand side, the simulator is restricted to only two modes of transport (car and public transit) and agents can perform only a single trip, preventing the modeling of trip chaining or detours. Moreover, trips are modeled from zone to zone, which limits the precision with which the first and last segments of trips are represented. On the supply side, congestion is modeled using speed-density functions, which, while effective in some cases, do not always satisfy the first-in-first-out (FIFO) property. This can lead to an inaccurate representation of real-world congestion dynamics. Additionally, METROPOLIS only supports a single vehicle type, meaning that trucks and other non-car vehicles cannot be simulated.

To address these limitations, and given that no existing simulator fully meets our requirements, the second chapter of this thesis focuses on the development of a new transport simulator, built from the ground up. This new simulator, named METROPOLIS2, builds on the same foundations as METROPOLIS, by simulating agents' mode, departure time, and route choices using discrete-choice models within a dynamic mesoscopic framework. However, it introduces many improvements on both the demand and supply sides of the model.

### Methodology

On the demand side, METROPOLIS2 allows agents to perform activities at multiple locations. Unlike METROPOLIS, which is limited to a single trip from an origin zone to a destination zone, agents in METROPOLIS2 can thus make more than one trip, with the

---

<sup>7</sup>In collaboration with André de Palma.

origins and destinations specified as exact coordinates. Agents can also choose from an arbitrary number of modes for these trips, with choices modeled through either a Multinomial Logit model or a deterministic choice model. The available modes include road-based modes (where trips are simulated from one intersection to another on the road network) and modes with constant travel times. For departure-time choice, METROPOLIS2 continues to support the Continuous Logit model, but it also offers the option to use a Multinomial Logit or to impose specific departure times. Unlike METROPOLIS, route choice in METROPOLIS2 is determined entirely before departure, using state-of-the-art routing algorithms. Additionally, METROPOLIS2 allows for greater flexibility in the utility specification, with preference parameters that can be defined at the agent level, and by supporting more general utility functions.

On the supply side, METROPOLIS2 simulates congestion in continuous time, an improvement over the 1-second time step used in METROPOLIS. This allows for explicit modeling of bottleneck congestion, with implications for the formation of queues and the calculation of waiting times, ensuring that the FIFO property remains satisfied. The simulator also supports multiple vehicle types, each with distinct characteristics such as vehicle length, passenger car equivalent (PCE) values, speed limits, and road access permissions, making it suitable for simulations involving mixed vehicle types, such as freight transport.

The development of the simulator was completed entirely by myself in Rust, a high-performance programming language. By making efficient use of parallelism whenever possible, the enhancements over METROPOLIS have been achieved without compromising the computational efficiency of the simulator.

## Results

Chapter 2 highlights METROPOLIS2's ability to replicate analytical results, demonstrated through its application to the single-road bottleneck model with homogeneous agents (de Palma et al., 1983; Arnott et al., 1990). Achieving this replication is a great accomplishment,

as simulating homogeneous agents is particularly challenging due to the risk of all agents attempting to depart at the same time, causing extreme congestion. This replicability is made possible by METROPOLIS2's continuous departure-time model, which ensures that agents distribute their departure times across a specific period.

The chapter concludes with a large-scale application to Île-de-France, comparing the performance of METROPOLIS2 with that of METROPOLIS. This comparison shows that METROPOLIS2 produces results consistent with those from METROPOLIS, while running faster and achieving better convergence.

## **Contributions**

This chapter contributes to the literature by introducing METROPOLIS2, a new transport simulator that improves key aspects of existing models like METROPOLIS (de Palma et al., 1997) and MATSim (Horni et al., 2016). Compared to METROPOLIS, METROPOLIS2 enhances the modeling of congestion and the decision-making process, while expanding the range of policies that can be effectively modeled. In relation to MATSim, which employs a co-evolutionary algorithm, METROPOLIS2 combines discrete-choice models with a convergence algorithm, offering improved potential for achieving a Nash equilibrium. This advantage is demonstrated through the successful replication of analytical results from the bottleneck model.

## **Limits**

This chapter does not address certain aspects essential for applying the simulator to real-world scenarios, such as the generation of a synthetic population and the calibration of the model. Additionally, a public-transit model could be incorporated to the simulator, enabling the consideration of in-vehicle congestion in route choices for public-transit trips.

## Chapter 3: Impact of Low Emission Zones on Spatial and Economic Inequalities using a Dynamic Transport Simulator

### Motivation

The third chapter is dedicated to evaluating the Low Emission Zone (LEZ) policy in the Paris metropolitan area. METROPOLIS2 is particularly well-suited for this task, as its features allow for the simulation of policies that affect agents differently based on their vehicles. The LEZ restricts the most polluting vehicles from entering a zone that includes Paris and 76 neighboring municipalities. Under the version of the policy scheduled for January 2025, 20.8% of the regional vehicle fleet would be banned from entering the zone.

### Methodology

In order to evaluate the LEZ's impact, this chapter calibrates a baseline simulation of METROPOLIS2 for the Île-de-France region. The baseline simulation represents an average working day in the region, encompassing all activity purposes and five modes: car driver, car passenger, public transit, bicycle, and walking. In total, approximately 9.37 million trips are simulated, including 600 thousand truck trips, on a road network with around 340 thousand road segments, covering about 43 700 km of roads.

The calibration process follows a four-step methodology, leveraging machine-learning techniques to match the simulated values in METROPOLIS2 with observed data while minimizing the number of simulations required. The four steps sequentially calibrate free-flow travel times, congested travel times, departure-time distributions, and mode shares. This methodology replicates data collected from the TomTom API and the regional travel survey, using techniques such as Lasso regression, Bayesian Optimization, and Random Forest regression.

## Results

Using both the calibrated baseline simulation and a counterfactual simulation with the LEZ, this chapter evaluates the policy’s overall and individual-level impacts. Results indicate that the LEZ would reduce car mode share from 36.6 % to 34.7 %, leading to a 3.9 % reduction in vehicle-kilometers traveled. METROPOLIS2 is also coupled with a model to compute global and local pollutant emissions and health impacts (Le Frioux et al., 2024), showing expected reduction of emissions by 4.5 % for CO<sub>2</sub>, 9.2 % for NO<sub>x</sub>, and 7.6 % for PM<sub>2.5</sub>. Moreover, premature deaths due to NO<sub>2</sub> and PM<sub>2.5</sub> exposure are projected to decrease by 9.9 % and 13.0 %, respectively.

At the individual level, health benefits are relatively evenly distributed, but variations in travel surplus show significant disparities. The majority of the costs are borne by a small minority of the population, owners of banned vehicles living within the LEZ, who must switch mode, resulting in longer travel times.

## Contributions

This chapter presents a methodology for calibrating simulations and evaluating LEZs or other transportation policies using agent-based simulators. The calibration procedure extends and improves several aspects of previous approaches, such as that of Ziemke et al. (2019), by incorporating additional data sources and new techniques. This approach enables calibration of a broader range of parameters than was previously possible and reduces the need for manual intervention.

Our evaluation methodology complements empirical studies of LEZs, such as that of Holman et al. (2015), by allowing detailed, scenario-based analysis. It also extends prior research on policy evaluation through agent-based transport simulations, such as the work of Yin et al. (2024), by offering a flexible framework capable of modeling differentiated policies and their heterogeneous impacts.



## Limits

The analysis in this chapter focuses on short-term decisions, such as mode, departure time, and route choice, while variables like car ownership, activity patterns, and residential location are held constant. In the long term, individuals may adapt to the policy by purchasing authorized vehicles or relocating closer to public transport services.

Moreover, the model does not account for correlations between vehicle ownership and socio-demographic characteristics within municipalities. Integrating a car-ownership model to METROPOLIS2 could address this limitation, providing a better understanding of how the LEZ affects different population segments.

## Other Contributions of the Thesis

In addition to evaluating Low Emission Zones, the advanced features of METROPOLIS2 make it well-suited for assessing a wider range of policies compared to METROPOLIS. For instance, METROPOLIS2's trip-chaining capabilities enable the simulation of ride-sharing scenarios where drivers make detours to pick up or drop off passengers. During this thesis, in collaboration with Samarth Ghoslya (Sapienza University of Rome) and under the supervision of André de Palma and Paolo Delle Site (Niccolò Cusano University), we initiated a project using METROPOLIS2 to compare the ride-sharing system explored in Chapter 1 with a different system. Instead of assuming that drivers maintain their routes while passengers walk to specific pick-up or drop-off points, this work explores a scenario where drivers may make detours to pick up passengers at their origin, or passengers walk to the driver's origin, with similar options at the destination. The optimization problem is also extended to minimize fuel costs and CO<sub>2</sub> emissions resulting from the road trips.

METROPOLIS2 is also well-suited for evaluating new public-transit infrastructure, particularly because trip origins and destinations can be specified as exact coordinates. This allows for more accurate estimates of the access (how to reach the first main station) and the egress (how to reach the final destination) portions of trips. As part of this thesis, I had the

opportunity to collaborate with the *Société des Grands Projets*, the institution overseeing the construction of the *Grand Paris Express*, to assess the impact of an expansion of a new métro line. Different scenarios regarding which stations would be served were explored, with METROPOLIS2 estimating passenger flows between any two stops on the expanded lines. Additionally, METROPOLIS2 provided insights into how the expanded line would affect passenger flows on the other lines of the public-transit network.

In a different area of study, this thesis explores the potential of personalized incentives policies as a tool for encouraging individuals to adopt more socially beneficial alternatives. These policies involve offering incentives to individuals to induce them to switch to behaviors that result in better social outcomes, such as using more sustainable modes of transportation. The incentives are personalized in the sense that the amount offered depends on each individual's willingness to switch, and the policy targets those who are most likely to be influenced by the incentive.

In collaboration with Andrea Araldo (Télécom SudParis, Institut Polytechnique de Paris) and André de Palma, we formalize the problem of determining the optimal allocation of personalized incentives to maximize social welfare, while remaining within the budget constraint of the regulator. Our work demonstrates that no other subsidy policy can further improve social welfare under this constraint, and we propose an algorithm to approximate the optimal allocation of incentives.

An example application of this model in Lyon is proposed, where individuals are offered incentives to switch to more ecologically-friendly transportation modes. This research has lead to the following publications:

Javaudin, L., Araldo, A., & de Palma, A. (2022). Large-scale allocation of personalized incentives. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)* (pp. 4151-4156). IEEE.

Javaudin, L., Araldo, A., & de Palma, A. (2023). Personalised incentives with constrained regulator's budget. *Transportmetrica A: Transport Science*, 1-43.

# Chapter 1

## Ride-sharing with Inflexible Drivers in the Paris Metropolitan Area

This chapter has been published in an academic journal as:

de Palma, A., Javaudin, L., Stokkink, P., & Tarpin-Pitre, L. (2024). Ride-sharing with inflexible drivers in the Paris metropolitan area. *Transportation*, 51(3), 963-986.

### Abstract

In ride-sharing, commuters with similar itineraries share a vehicle for their trip. Despite its clear benefits in terms of reduced congestion, ride-sharing is not yet widely accepted. We propose a specific ride-sharing variant, where drivers are completely inflexible. This variant can form a competitive alternative against private transportation, due to the limited efforts that need to be made by drivers. However, due to this inflexibility, matching of drivers and riders can be substantially more complicated, compared to the situation where drivers can deviate.

In this work, we propose a four-step procedure to identify the effect of such a ride-sharing scheme. We use a dynamic mesoscopic traffic simulator which computes departure-time choices and route choices for each commuter. The optimal matching of potential drivers and riders is obtained outside the simulation framework through an exact formulation of the problem. We evaluate the potential of this ride-sharing scheme on a real network of the Paris metropolitan area for the morning commute. We show that even with inflexible drivers and when only a small share of the population is willing to participate in the ride-sharing scheme, ride-sharing can alleviate congestion. Further improvements can be obtained by increasing the capacity of the vehicles or by providing small monetary incentives, but without jeopardizing the inflexibility of the drivers. Thereby, we show that ride-sharing can lead to fuel savings, CO<sub>2</sub> emission reductions and travel time savings on a network level, even with a low participation rate.

**Keywords:** ride-sharing; carpooling; matching; dynamic congestion

**JEL Codes:** R41; R48

## 1.1 Introduction

Ride-sharing, also known as carpooling, is a non-profit shared ride service where a car owner shares his / her vehicle with another person heading in the same direction to share expenses. It aims to solve one key problem of urban congestion: low vehicle occupancy, especially for commuting trips. In the Paris region, there are 1.05 persons per vehicle on average for commuting trips (Enquête Globale Transport, 2010). This rate has been decreasing since 1976 (Cornut, 2017). In urban areas, congestion also has severe implications with regards to air pollution. Ride-sharing offers the opportunity to raise average vehicle occupancy and to address public health and climate change issues. For travellers, it is also an opportunity to save on fuel cost.

In this work, we propose a ride-sharing scheme quite similar to conventional hitch-hiking. Drivers do not deviate from their predetermined itinerary, meaning they determine their optimal departure time and exact route without considering a potential rider. The rider then adapts to the itinerary of the matched driver. This implies that the rider may need to walk to reach the driver and to reach his / her final destination after being dropped off by the driver. However, similar to hitch-hiking, the trip is assumed to be completely free of charge for the rider.

Our key hypothesis is that the segment of the rider's trip spent in a personal vehicle, will be offered by a driver who has already planned to travel that segment on his / her own trip. This is the key feature of our system: drivers are inflexible. This hypothesis is explained by the observation that one of the setbacks in the development of the ride-sharing process is that drivers are reluctant to change their routes and schedules. Another disadvantage of ride-sharing, not considered in this paper, is the inconvenience of having another person in the car; later on, one can think of some certification systems to reduce the uncertainty to have somebody else (unknown) in their own car. Of course, such certification (of the car, the insurance status or the driving license) should preserve anonymity and should not be incompatible with privacy rules.

The emergence of ride-sharing services such as UberPool, Lyft Line, and Blablacar Daily (not to be confused with the ride-sourcing services provided by Uber and Lyft) has been a major competitor to the practice of ride-sharing (Shaheen and Cohen, 2019). Ride-sharing consists of people with similar travel needs travelling together, whereas ride-sourcing consists of car owners offering paid lifts to gain money. The social benefits of ride-sharing are manifold: less traffic congestion (Cici et al., 2014; Xu et al., 2015), less CO<sub>2</sub> and NO<sub>x</sub> emissions leading to better air quality (Bruck et al., 2017), and better transit accessibility in suburban areas (Teubner and Flath, 2015; Li et al., 2016; Kong et al., 2020). Moreover, ride-sharing brings about travel cost sharing for riders and drivers (Malichová et al., 2020). However, the popularity of ride-sharing remains low for commuting trips.

Many forms of ride-sharing have been studied over the years to increase the mode's convenience and maximise the societal gains it provides in terms of traffic congestion as well as of emissions reduction. Nonetheless, each of them presents certain drawbacks. For instance, multi-hop ride-sharing explores the possibility for a rider to use multiple cars to complete his / her trip at the cost of a transfer penalty and waiting time. As for detours created by door-to-door ride-sharing services, they increase the driver's travel distance and time, all the more in the case of multiple passengers. Some ride-sharing companies have stopped door-to-door service and now ask riders to walk in order to reduce the extent of the detours (Lo and Morseman, 2018; Schaller, 2021).

This research proposes a four-step procedure to evaluate the effect of a ride-sharing scheme where the driver makes no detour at all and no concession on his / her schedule. The procedure can assess the impact of the ride-sharing scheme on congestion and CO<sub>2</sub> emissions reduction. The first step consists in running the mesoscopic dynamic traffic simulator METROPOLIS to identify the departure time and route chosen by the drivers. In the second step, the ride-sharing costs are computed. The ride-sharing scheme is such that the ride is free of charge for the riders and the inconvenience of the driver is completely compensated by state subsidies. In the third step, the optimal matching is obtained by solving an Integer

Linear Programming (ILP) problem. The matching is such that each rider is matched with a driver who has similar a origin and destination and whose trip timing is compatible with the schedule-delay preferences of the rider. Finally, the fourth step consists in running another simulation of the traffic simulator to compute the new congestion level.

Whereas existing works have evaluated ride-sharing methods only on small scale networks, we evaluate the potential of our ride-sharing scheme for the Paris metropolitan area, under dynamic congestion. We consider different scenarios, with a different share of travellers willing to participate in the ride-sharing scheme. Scenarios with more than one rider in the car and scenarios where incentives are being proposed to riders are also considered.

In October 2022, the French government announced a subsidy of 100 € for new ride-sharing users, illustrating the willingness of governments to subsidize ride-sharing.<sup>1</sup> Our results show that a government-funded ride-sharing scheme can be beneficial for society. More precisely, we find that, even when only a small share of the population is willing to participate in the scheme, ride-sharing can significantly reduce congestion and CO<sub>2</sub> emissions. Additionally, as shown by Lian and Van Ryzin (2021), the optimal policy in two-sided markets like ride-sharing is to provide initially significant spending in order to reach, as early as possible, a critical mass of users. Therefore, a temporary government intervention, even if it is costly, can have a long-run impact on the modal share of ride-sharing and thus on congestion and CO<sub>2</sub> emissions.

The remainder of this paper is structured as follows. Section 2 reviews the ride-sharing literature and the ways it is modelled. Section 3 describes the proposed ride-sharing scheme, the dynamic traffic simulator, and the proposed driver-rider matching methodology. Section 4 presents the case study results for Île-de-France (Paris area) under three maximum walking time scenarios, and for various penetration rates. Section 5 concludes with the key results and explores further research steps needed to explore the feasibility of a real operational-system.

---

<sup>1</sup><https://www.service-public.fr/particuliers/actualites/A16012>

## 1.2 Literature Review

Sharing mobility is part of the global trend towards a sharing economy (Standing et al., 2019). Shaheen and Cohen (2019) provide an overview of the different shared-ride services. Ride-sharing, also known as carpooling, and ride-sourcing, also known as ride-hailing, are two of the main shared-ride services. Whereas the former is associated with many societal benefits, the latter is an on-demand transportation service similar to taxi service with privately owned vehicles. Ride-sourcing is often associated with an increased traffic congestion (Schaller, 2021). Ride-sharing is inherently a non-profit mode that brings together people with similar trip itineraries to share their trip. The body of literature on this topic has significantly increased in the last decade as it has become more convenient to plan, book, and pay for a ride (Shaheen and Cohen, 2019). Indeed, Transportation Network Companies (TNCs) such as Uber and Lyft offer online ride-sharing services (UberPool and Lyft Line) in addition to their standard ride-sourcing services.

The matching problem between the rider and the driver has been extensively studied. Matching problems can be either static (Yan and Chen, 2011; Herbawi and Weber, 2012; Ma et al., 2019a; Liu et al., 2020; Lu et al., 2020) or dynamic (Agatz et al., 2011; Kleiner et al., 2011; Di Febbraro et al., 2013). In static matching problems, all drivers and riders are known in advance and are matched at the same time. Dynamic matching problems consider that drivers and riders arrive gradually. In this case, partial matchings can be performed with a subset of the drivers and riders. This work uses static ride-sharing under dynamic congestion.

The main benefit of ride-sharing is that it eases traffic congestion (Cici et al., 2014; Xu et al., 2015). It hence offers a great potential for CO<sub>2</sub> emission reductions (Chan and Shaheen, 2012; Bruck et al., 2017). Furthermore, it offers more accessibility to public transit as a first/last mile solution (Teubner and Flath, 2015; Li et al., 2016; Kong et al., 2020). Ride-sharing may, however, increase the driver's trip time through detours to pick up and drop off riders (Diao et al., 2021). Schaller (2021) analyses extensive longitudinal data from TNCs in



American cities. He observes that ride-sharing services mainly draw people from transit as it is mostly popular in neighbourhoods with low incomes and low car ownership rates. This phenomenon has become even more evident since UberPool and Lyft Line stopped door-to-door services, with the aim to reduce detours. This finding is in line with many other studies concluding that an increase in the modal share of ride-sharing does not cause a significant reduction in the modal share of car (Xu et al., 2015; Li et al., 2016; Shaheen et al., 2016; Coulombel et al., 2019; Li et al., 2016).

Despite the many benefits of ride-sharing, it is still not widely used as a mode to commute (Liu et al., 2020). Amongst the challenges to have a successful ride-sharing system is the large population of drivers necessary to provide high-quality matches in terms of geographic and temporal proximity (Bahat and Bekhor, 2016). Substantial research has been conducted to understand the individual motivations behind ride-sharing in order to increase its popularity. Cost savings followed by environmental concerns are the main motivations reported both by the drivers and the riders (Delhomme and Gheorghiu, 2016; Neoh et al., 2017; Gheorghiu and Delhomme, 2018; Pinto et al., 2019). Malichová et al. (2020) observed through a pan-European survey that travellers prefer to adopt ride-sharing for work compared to other purposes.

Ride-sharing has been modelled alongside transit both as a complement providing a solution to the first/last mile problem (Masoud et al., 2017; Ma et al., 2019b; Kumar and Khani, 2020; Reck and Axhausen, 2020) and as a competitor (Qian and Zhang, 2011; Galland et al., 2014; Friedrich et al., 2018). Qian and Zhang (2011) use a theoretical bottleneck model where the modal choice between car, transit, and ride-sharing depends on the generalised travel time. They account for transit perceived-inconvenience depending on transit passenger-flow. Schedule delay is considered for the three modes. de Palma et al. (2022b) build on this framework to add dynamic congestion. Coulombel et al. (2019) use a transportation-integrated land use model to consider the impact of ride-sharing on car and transit ridership for the Paris region. Finally, Galland et al. (2014) propose an agent-based model for ride-sharing to

analyse individual mobility behaviour. Due to its computational complexity, they test their model on a population of only 1000 agents. To predict the route taken by the drivers in a large-scale scenario, this work uses the dynamic traffic-assignment simulator METROPOLIS (de Palma et al., 1997), which can account for the timings of the trips when matching riders with drivers. This traffic simulator computes, for each individual, the route, departure-time and mode choice, using a nested Logit model. The schedule-delay costs are based on idiosyncratic  $\alpha - \beta - \gamma$  preferences (Vickrey, 1969) and congestion is modelled with link-specific bottlenecks.

This work builds on the many ride-sharing models present in the literature. The methodology allows to assess the potential of ride-sharing in a large urban area under dynamic congestion, whereas previous models were either applied on simple bottleneck models (Qian and Zhang, 2011; Yu et al., 2019; de Palma et al., 2022b) or were too sophisticated to provide results for a large urban network (Galland et al., 2014). Alisoltani et al. (2021) also consider a medium sized urban network with traffic dynamics, but the main difference with our work is that their drivers are not commuters but employees of the company offering the ride-sharing service. Furthermore, we consider mode choice (including public transport) and departure time choices of commuters that consider scheduling delay preferences.

## 1.3 Methodology

### 1.3.1 Ride-sharing Scheme

This paper explores a ride-sharing scheme where ride-sharing drivers make no detour and keep the exact same schedule as when driving alone. Ride-sharing drivers simply pick up a passenger at a defined road intersection on their itinerary and drop off their passenger at another intersection on their itinerary. As for the riders, they need to walk from their origin to a pick-up point and from a drop-off point to their destination. They face schedule-delay costs if their arrival time does not match their desired arrival time. However, the trip is free

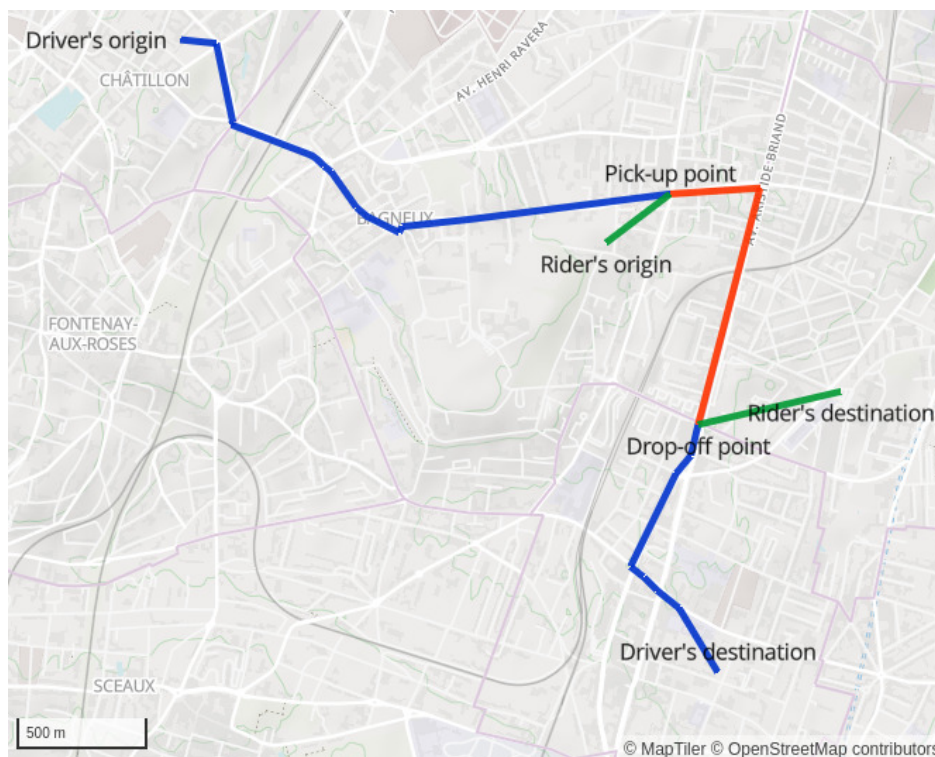


Figure 1.1: Ride-sharing trip example

Note. The blue lines represent the driver's trip alone, the red line represents the car trip shared between the driver and the rider and the green lines represent the walking trips of the rider.

of fare for them.

Figure 1.1 provides an example of a ride-sharing trip under this scheme. The driver's itinerary is composed of the two blue parts (where he / she is alone) and the red part (where he / she has a passenger). The rider's trip consists of a walking leg from his / her origin to the pick-up point (in green), a car leg with the driver (in red) and another walking leg from the drop-off point to his / her destination (in green).

We emphasize that in this framework, in order to match drivers and passengers and evaluate the corresponding costs, drivers announce their complete itinerary (i.e., exact route and departure time), whereas riders only communicate their origin, destination and desired arrival time. The exact route of the riders follows directly from the route of the matched driver.

Implementing such a ride-sharing scheme on a large scale would require some sort of state intervention, e.g., through subsidies, in order to convince enough drivers and riders to

subscribe to the scheme. As drivers do not deviate from their route nor from their desired departure time, drivers only need to be compensated for the inconvenience of having someone in their car (Li et al., 2020a).

Riders may also experience an inconvenience cost, arising from the discomfort of sharing a ride with a stranger. On top of this, they also have to walk and may incur additional schedule-delay costs. However, the scheme allows them to save money on gas, car wear and tear, parking, and car insurances. Moreover, they do not spend time driving around for a parking slot any more. To convince more individuals, additional subsidies may be offered to reduce the generalised costs of ride-sharing.

An increase in the modal share of ride-sharing can greatly reduce the congestion in a city. By increasing the occupancy of vehicles, the number of vehicles on the road decreases. In turn, the externalities associated with congestion (including air pollution, noise pollution and safety) also decrease. Public authorities may therefore be interested in subsidising ride-sharing to reduce congestion and its environmental cost.

### 1.3.2 Four-Step Procedure

We propose a four-step procedure to match drivers and riders in our ride-sharing framework and evaluate the impact of ride-sharing on congestion. A part of this procedure may be repeated to achieve convergence.

1. *Initialize:* The traffic simulator METROPOLIS is used to simulate a reference scenario without ride-sharing. This means all commuters can choose between driving solo or using public transport. The simulator is run up to a stationary regime such that the mode choice, routes and departures times of every individual is identified. A detailed explanation of the traffic simulator and our definition of stationary regime is given in Section 1.3.3.
2. *Cost computation:* The ride-sharing costs of any pair of commuters that is willing to participate in the ride-sharing scheme is computed. Section 1.3.4 describes how the

costs are evaluated for every pair. The reference simulation is used to obtain the exact route and departure time of a participating driver, as well as the origin, destination and desired arrival time of a participating rider.

3. *Optimal matching:* Based on the computed costs, the optimal matching of drivers to riders is obtained by solving an Integer Linear Programming (ILP) problem. This is described in Section 1.3.5.
4. *Ride-sharing simulation:* All riders that were matched to a driver in the previous step are excluded from the set of commuters and the traffic simulator is run to a new stationary regime with the remaining commuters.

Steps 2-4 can be repeated until convergence is observed. Then, the final simulation can be used to evaluate the effect of ride-sharing on congestion, mode share and total mileage. We emphasize that over the entire network, ride-sharing reduces congestion by reducing the total number of cars. However, due to mode changes to and from public transport, there may be a local increase in congestion in some parts of the network. Due to the change in congestion, additional commuters may choose to ride-share or individuals who chose to ride-share may regret this for the new congestion level. For this reason, an iterative framework using steps 2-4 can be used to evaluate this kind of behaviour, but this is omitted in this work due to its computational complexity.

### 1.3.3 Traffic Simulator

To assess how congestion evolves as the number of cars decreases and car occupancy increases, we use METROPOLIS, a mesoscopic dynamic traffic simulator developed by de Palma et al. (1997). Since then, it has mostly been used to estimate various transport policies, including different road pricing schemes (Saifuzzaman et al., 2016; de Palma et al., 2022b). The inputs of the simulator include a description of the road and public-transit networks, origin-destination matrices and travellers' preferences. The outputs are the choices of the travellers and the dynamic state of the road network (time-dependent congestion levels and travel

times). In METROPOLIS, congestion is represented using link-level bottlenecks.

The choices made by each traveller can be summarised as:

1. Mode choice (between car and transit): The generalised cost for transit is compared with the generalised cost for car. The public transit cost is function of the value of time of transit, the transit travel time, and the transit fare. In the current version, generalised public transport costs are exogenous. The generalised cost for car is a function of the value of time of car, the endogenous travel-time, and the schedule-delay cost. The mode choice is given by a nested Logit model.
2. Departure-time choice: The probability of choosing a departure time  $t$  is given by a continuous Logit model, according to the generalised cost for each possible departure time.
3. Route choice: Each day, at each intersection, travellers observe the congestion on upstream roads and choose a road in order to minimize their generalised cost (closed loop equilibrium).

Note that commuters only choose between car and transit, while ride-sharing is not explicitly modelled as a mode in the simulator. By assumption, drivers are always fully compensated for ride-sharing inconvenience and riders only accept matches that decrease their generalised travel cost, which justifies their mode choice for ride-sharing.

METROPOLIS uses a day-to-day iterative procedure. At each iteration, the travellers choose their mode, departure-time and route, given the expected dynamic congestion levels. At the end of each iteration (day), the expected congestion levels are updated using the observed congestion levels, according to a day-to-day adjustment process: the expected congestion for the next iteration is a weighted average of the current expected congestion and the observed congestion. The simulation stops when the two levels are close, i.e., when a stationary regime is reached.

### 1.3.4 Ride-sharing Cost

The cost of ride-sharing, for a rider, is the sum of walking cost, in-vehicle cost and schedule-delay cost. We consider the cost for a rider  $i$  when matched to a driver  $j$ .

The walking cost is the cost of walking from the origin to the pick-up point and from the drop-off point to the destination. Let  $v^{\text{walk}}$  be the walking speed. The duration of the walking trip from the rider's origin to the pick-up point is assumed to be  $d_{ij}^{\text{pick}}/v^{\text{walk}}$ , where  $d_{ij}^{\text{pick}}$  is the Euclidean distance between the rider's origin and the pick-up point, when matching rider  $i$  with driver  $j$ . The duration of the walking trip from the drop-off point to the rider's destination is assumed to be  $d_{ij}^{\text{drop}}/v^{\text{walk}}$ , where  $d_{ij}^{\text{drop}}$  is the Euclidean distance between the drop-off point and the rider's destination.

The time at which driver  $j$  picks up (resp. drops off) rider  $i$  is denoted  $t_{ij}^{\text{pick}}$  (resp.  $t_{ij}^{\text{drop}}$ ). It is equal to the time at which driver  $j$  is predicted to reach the pick-up point (resp. drop-off point). Then, the duration of the car trip for the rider is  $tt_{ij}^{\text{iv}} = t_{ij}^{\text{drop}} - t_{ij}^{\text{pick}}$  and the arrival time at destination is  $t_{ij}^a = t_{ij}^{\text{drop}} + d_{ij}^{\text{drop}}/v^{\text{walk}}$ .

Each rider  $i$  has a specific desired arrival time  $t_i^*$  and a tolerance for lateness or earliness  $\Delta_i$ . Riders who reach their destination within the  $t_i^* \pm \Delta_i$  window experience no schedule-delay penalty. Every minute outside this on-time window generates a schedule delay cost. The schedule delay cost of rider  $i$ , when matched with driver  $j$ , is

$$SD_{ij} = \beta_i [(t_i^* - \Delta_i) - t_{ij}^a]^+ + \gamma_i [t_{ij}^a - (t_i^* + \Delta_i)]^+,$$

where  $t_{ij}^a = t_{ij}^{\text{drop}} + d_{ij}^{\text{drop}}/v^{\text{walk}}$  is the arrival time of the rider at destination,  $\beta_i$  is the penalty associated to early arrival,  $\gamma_i$  is the penalty associated to late arrival, and  $[x]^+ = \max(0, x)$ .

To sum up, the generalised cost of ride-sharing experienced by rider  $i$ , when matched

with driver  $j$ , is

$$c_{ij}^R = \underbrace{\alpha_i^{\text{RS}} \cdot tt_{ij}^{\text{iv}}}_{\text{In-vehicle cost}} + \underbrace{\alpha_i^{\text{walk}} \cdot \left[ \frac{d_{ij}^{\text{pick}} + d_{ij}^{\text{drop}}}{v^{\text{walk}}} \right]}_{\text{Walking cost}} + \underbrace{SD_{ij}}_{\text{Schedule-delay cost}},$$

where  $\alpha_i^{\text{RS}}$  is the value of time of rider  $i$  during the ride and  $\alpha_i^{\text{walk}}$  is the value of time of rider  $i$  when walking. The pick-up and drop-off points are chosen so as to minimize the generalised cost  $c_{ij}^R$ .

### 1.3.5 Matching

The matching of drivers and riders is determined through an Integer Linear Programming (ILP) formulation. We define  $N$  the set of individuals that are willing to participate in the ride-sharing program. Individuals in  $N$  agree to be matched as driver or rider in the ride-sharing program; they travel alone in the absence of match. According to the proposed ride-sharing scheme, drivers will not deviate from their route nor will they change their arrival time to account for riders. Therefore, drivers will not have costs involved with ride-sharing and the costs for driver  $j \in N$  are equal to  $c_j^{\text{NR}}$ , independent of whether there are any riders on the car. We also emphasize that an individual can use public transport rather than drive. In this case,  $c_j^{\text{NR}}$  encompasses the costs of public transportation. Riders on the other hand, have a cost associated to ride-sharing, as defined in Section 1.3.4. We define  $c_{ij}^R$  the cost of rider  $i \in N$  when taking a ride from driver  $j \in N$ . The total number of riders that driver  $j \in N$  can take is equal to  $a_j$ . Here the driver is not accounted for, so for a car with 5 seats  $a_j$  would be equal to 4.

We define binary decision variable  $x_{ij}$  which is equal to 1 if rider  $i \in N$  is matched to driver  $j \in N$ , and 0 otherwise. Furthermore, we define binary decision variable  $y_i$  which is equal to 1 if rider  $i \in N$  is not matched to any driver. The objective is to minimize the total costs of all individuals that are willing to participate in the ride-sharing scheme. This



is equivalent to maximizing the total cost reduction associated to ride-sharing.

$$\min \sum_{i,j \in N} c_{ij}^R x_{ij} + \sum_{i \in N} c_i^{\text{NR}} y_i \quad (1.1)$$

s.t.

$$y_i + \sum_{j \in N} x_{ij} = 1 \quad \forall i \in N \quad (1.2)$$

$$\sum_{i \in N} x_{ij} \leq a_j y_j \quad \forall j \in N \quad (1.3)$$

$$x_{ij} \in \{0, 1\} \quad \forall i, j \in N \quad (1.4)$$

$$y_j \in \{0, 1\} \quad \forall j \in N \quad (1.5)$$

The objective in Equation (1.1) is to minimize the joint cost of ride-sharing and travelling by car for all individuals. Constraints (1.2) imposes that all individuals are either matched to a driver or driving themselves. Constraints (1.3) enforces that a rider can only be matched to an individual that is driving and it enforces the capacity of the vehicle. Constraints (1.4) and (1.5) define the binary range of the decision variables. We note that when individual  $j \in N$  uses public transit,  $y_j = 1$ . However, it is impossible to assign a rider to this individual as  $c_{ij}^R = \infty$  for all  $i \in N$  if  $j$  is a public transport user.

Note that the total number of riders is not known *a priori* as individuals in the set  $N$  can be either rider, driver (with or without passenger) or public-transit user. Instead, the number of riders depends on the quality of the matches. Also observe that an individual  $i$  can be matched with a driver  $j$  only if the ride-sharing cost is smaller than the non ride-sharing cost, i.e.,  $c_{ij}^R < c_i^{\text{NR}}$ . In this respect, the matching program proposes only Pareto-improving matches.

## 1.4 Case Study: Ride-sharing in Île-de-France

The Paris area, as many other large cities, experiences frequent heavy pollution episodes partly due to car emissions (Degraeuwe et al., 2017; Kumar et al., 2021). The regional government of Île-de-France created subsidy programs in 2017 to promote ride-sharing and address this issue. The programs include, *inter alia*, direct subsidies for ride-sharing drivers, the funding of ride-sharing companies so that they offer lower fares to riders, and two monthly free rides to frequent transit users. Drivers receive from the government 1.50 € per passenger plus 0.10 €/km up until a maximum of 3 € per trip. Moreover, the regional government has made ride-sharing completely free for riders during peak pollution episodes and during transit strikes.<sup>2 3</sup>

The ride-sharing scheme proposed in this research is tested on the Île-de-France region. Île-de-France accounts for nearly a fifth of France’s population with its 12 175 000 inhabitants in 2017. The region, mainly consisting of Paris and its suburbs, has a density of 1013 inhabitants per square kilometre. Region-wide, there are 43 million trips daily amongst which 42 % are made by foot or bicycle, 22 % by public transit, and 36 % by car. There are however wide disparities between the city of Paris, the inner and the outer suburbs (Île-de-France Mobilités, 2019).

### 1.4.1 Network Modelling

We use the calibration of METROPOLIS for Île-de-France from Saifuzzaman et al. (2012), which is based on demand data from the 2001 Paris origin-destination survey. The road network consists of 43 857 links, 18 584 intersections, and 1360 zones. Each link is unidirectional and represents a bottleneck with a link-specific capacity. The origin and destination of travellers is set to the centroid of their origin / destination zone. The centroids are con-

---

<sup>2</sup><https://www.iledefrance.fr/la-prime-au-covoiturage-prolongee-et-etendue>

<sup>3</sup><https://www.iledefrance.fr/covoiturage-jusqua-150-euros-par-mois-pour-les-conducteurs>

nected to the road network with uncongested links. Figure 1.2 is a visual representation of the network. Compared to the original calibration by Saifuzzaman et al. (2012), we enable mode choice, which requires recalibrating the road capacities.




Figure 1.2: Île-de-France road network

Note. Red lines are uncongested virtual roads connecting the centroids of the zones to the road network (in blue).

There is no public transit network *per se*, but rather exogenous travel times for each origin and destination pair of Île-de-France. The public-transit generalised costs are taken from the DRIEAT (*Direction régionale et interdépartementale de l'environnement, de l'aménagement et des transports d'Île-de-France*).

The pick-up and drop-off points can be any of the intersections of the road network. The computed walking distance is the Euclidean distance between the centroid of the zones and the intersections.

### 1.4.2 Travel Demand

We simulate the morning commute, which is the most congested period of the day in Île-de-France. The simulation starts at 6 a.m. and ends at 13 p.m.. Travel demand is represented as an origin-destination matrix for different traveller groups. All demand data are taken from

the calibration of METROPOLIS for Île-de-France (Saifuzzaman et al., 2012). Demand data is representative of a typical morning commute. A more thorough analysis would be needed to consider the impact of the day-to-day variations of demand on the ride-sharing matching.

Travel demand for the morning commute is divided in four traveller groups: workers going towards Paris, workers leaving Paris, and two groups of non-workers. Both the demand and the road capacity are scaled down to 50% to reduce computation time. All travellers are car owners and can choose between taking their car or taking the public transit. There is a total of 934 042 travellers.

In each group, the travellers have the same schedule-delay parameters and values of time but the desired arrival times are normally distributed. Figure 1.3 represents the desired arrival time distribution for the four groups of travellers. Workers coming from Paris are the ones with the narrowest distribution and the earliest desired arrival time. The workers originating from the suburbs and going towards Paris want to reach their destination, in average, a few minutes later. The desired arrival time of the non-workers is represented by two normal curves with a standard-deviation of 90 minutes. Non-workers have a later desired arrival time than commuters.




Figure 1.3: Desired arrival time distribution of the four traveller groups

All the preference parameters used in this research are presented in Table 1.1. The

values of time for car and transit as well as the early and late penalties come from the work of Saifuzzaman et al. (2012). The value of time for riders is assumed to be equal to the value of time of car (i.e.,  $\alpha_{\text{car}} = \alpha_{\text{RS}}$ ), which means that for riders the savings incurred by ride-sharing are completely offset by its inconvenience. Workers starting their journey in Paris are more inflexible in their desired arrival time as shown by their penalty for late arrival being more than twice the one of workers starting their journey in the suburbs. The value of time of walking is assumed to be  $\alpha_{\text{walk}} = 1.1 \cdot \alpha_{\text{RS}}$  (Wardman, 2001; Hensher and Rose, 2007). For all travellers, the walking speed is set to  $v_{\text{walk}} = 4 \text{ km/h}$  and the length of the on-time window is set to  $\Delta = 5 \text{ min}$ .

Traveller group	$\beta$	$\gamma$	$\alpha_{\text{car}}$	$\alpha_{\text{RS}}$	$\alpha_{\text{PT}}$	$\alpha_{\text{Walk}}$
Workers going towards Paris	6.09	7.53				
Workers coming from Paris	8.36	17.43				
Non-workers 1	5.24	10.64	12.96	12.96	13.24	14.26
Non-workers 2	5.24	10.64				

Table 1.1: Preference parameters for the four groups of travellers, in €/h

We assume that the ride-sharing cost  $c_{ij}^{\text{R}}$  is computed based on the realised travel time of the driver, i.e. his / her departure time on the previous day (or any announced departure time). One justification is that the driver has to announce beforehand his / her departure time in order to make the matching procedure feasible. On the other hand, the cost as a driver,  $c_i^{\text{NR}}$ , is based on anticipated travel times, i.e., it is the expected minimum cost over the departure-time period, computed from the log-sum formula given by METROPOLIS. This discrepancy between cost as a driver and as a rider can introduce a bias towards ride-sharing if the anticipated travel times are over-estimated within METROPOLIS (e.g., due to an imperfect convergence of the simulator).

Estimating the actual willingness of travellers to participate in the ride-sharing scheme is outside the scope of this paper. Instead, we study five scenarios with a different *participation rate*: 10 %, 20 %, 30 %, 40 % and 50 %. The travellers willing to participate are selected randomly among both car drivers and public-transit users. For simplicity, we assume that

the set of drivers who participate in the ride-sharing scheme coincides with the set of drivers who accept to have a passenger in their car. In practice, some travellers might be willing to do ride-sharing as a rider but not as a driver, or vice-versa.

### 1.4.3 Results

Table 1.2 presents the results from the METROPOLIS simulation for the reference scenario (with no ride-sharing), and the five ride-sharing scenarios. The individual surplus is the sum, over any individual, of the (opposite of the) individual's generalised travel cost (for car, public-transit or ride-sharing). The surplus variation represents its absolute variation, compared to the reference simulation. The *Car VKT* indicator represents the total distance (in thousands of kilometres) travelled by cars during the morning commute. Congestion is computed as

$$\frac{1}{|L|} \cdot \sum_{l \in L} \frac{tt_l^{avg} - tt_l^0}{tt_l^0},$$

where  $L$  is the set of all links in the network,  $|L|$  is its cardinality,  $tt_l^{avg}$  is the average travel-time on link  $l$  for the simulation period and  $tt_l^0$  is the free-flow travel-time of link  $l$ . The reduction in CO<sub>2</sub> (carbon dioxide equivalent) emissions is computed assuming average CO<sub>2</sub> emissions per car of 0.193 kg/km (Agence de la transition écologique, 2021).

In the ride-sharing scenarios, the individual surplus increases compared to the reference scenario for two reasons: (i) all the riders have now a lower generalised travel cost; (ii) there are less cars on the network and thus congestion and the average generalised cost of the drivers is smaller. In addition to the decrease in congestion, the reduction of the number of cars on the network also implies a significant decrease of CO<sub>2</sub> emissions, less noise, and improved air quality.

The public-transit modal share decreases from 25.5 % in the reference scenario to 23.5 % in the 50 % scenario. This is due to two different shifts: (i) public-transit users shifting to riders; (ii) public-transit users shifting to car drivers (more details in Section 1.4.4). As this

Scenario	Reference	10 %	20 %	30 %	40 %	50 %
<b>Shares</b>						
Transit modal share	25.5 %	25.3 %	24.8 %	24.3 %	23.9 %	23.5 %
Car modal share	74.5 %	73.9 %	73.2 %	72.4 %	71.5 %	70.5 %
Ride-sharing modal share	0.0 %	0.9 %	2.0 %	3.3 %	4.6 %	6.0 %
<b>Surplus</b>						
Individual surplus variation (€)	—	+20 326	+63 594	+104 897	+148 529	+248 744
<b>Road network</b>						
Car VKT (10 <sup>3</sup> km)	10 799	10 740	10 686	10 595	10 499	10 377
Congestion	22.1 %	21.7 %	21.4 %	20.6 %	19.8 %	19.2 %
CO <sub>2</sub> emissions reduction (tons)	—	11.387	21.809	39.372	57.900	81.446
<b>Drivers</b>						
Mean travel time	15' 32"	15' 31"	15' 32"	15' 27"	15' 22"	15' 19"
Mean travel cost (€)	6.03	6.02	6.02	6.00	5.97	5.95
Share of time spent with a passenger (for ride-sharing drivers only)	—	51.5 %	56.1 %	58.0 %	59.8 %	60.5 %
<b>Riders</b>						
Mean Euclidean distance (meters)	—	5491	5972	6205	6425	6539
Mean walking distance (meters)	—	383	347	325	310	303
Mean walking distance (if positive, meters)	—	470	451	438	432	430
Mean car travel time	—	7' 21"	8' 00"	8' 20"	8' 38"	8' 47"
Mean travel time	—	13' 06"	13' 12"	13' 13"	13' 17"	13' 20"
Mean travel cost (€)	—	3.26	3.24	3.22	3.22	3.22
Riders at their best match	—	76.7 %	69.3 %	65.0 %	62.2 %	59.1 %

Table 1.2: Comparison of results for the reference scenario and five scenarios with a different participation rate

Note. The surplus, car VKT and CO<sub>2</sub> emissions values are for a single representative morning commute.

work focuses on the impact of ride-sharing on congestion and CO<sub>2</sub> emissions,<sup>4</sup> neither of these shifts is desirable. First, the shift of some public-transit users to riders implies that some potential good matches are no longer available to other potential riders. Therefore, the number of car users becoming riders is smaller than what it would be if the matching algorithm ignored all public-transit users. Second, the shift of some public-transit users to car drivers implies that the decrease in congestion and CO<sub>2</sub> emissions is not as large as what it would be if public-transit users were forced to stick to public transit.

We can observe that the mean Euclidean distance between the origin and the destination

<sup>4</sup>In particular, we omit the impact of a decrease of public-transit use on in-vehicle congestion or service quality / frequency.

of the riders increases with the participation rate (from 5.5 km in the 10 % scenario to 6.5 km in the 50 % scenario). An explanation for this increase would be that, as the participation rate increases, it gets easier to find good matches between drivers and riders with a large O-D distance. Even though the mean O-D distance increases with the participation rate, the mean generalised travel cost for the riders decreases (from 3.26 € to 3.22 €), implying an increase of the match quality for the riders.

#### 1.4.4 Detailed Results for the 30 % Scenario

In this section, we look at detailed results on the matches for the scenario with a participation rate of 30 %. Figure 1.4 represents the mode shifts compared to the reference scenario without ride-sharing. It can be observed that the 3.34 % of riders were either former car drivers or former public-transit users. Some public-transit users are shifting to the car because road congestion is smaller.

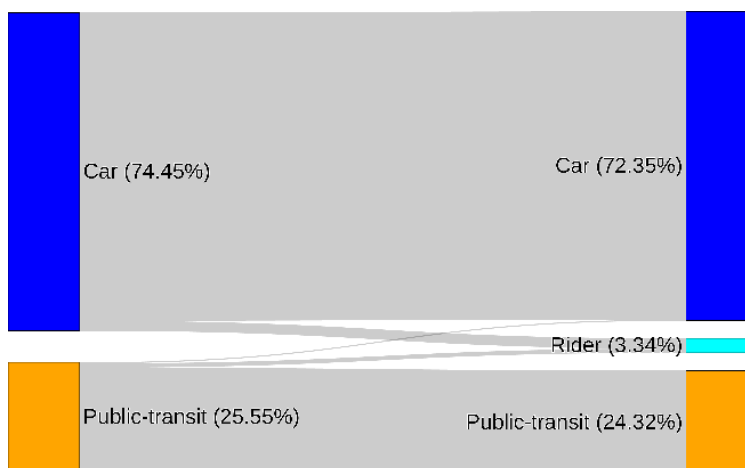
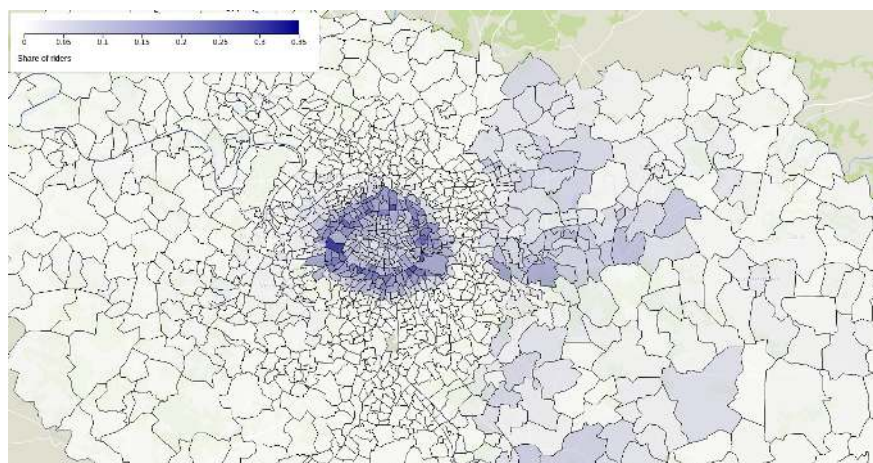


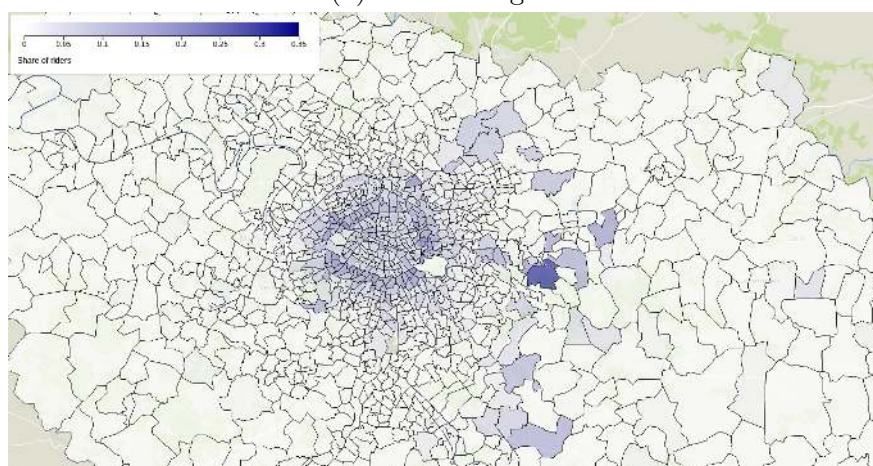
Figure 1.4: Mode shifts from the reference scenario (left) to the 30 % participation rate scenario (right)

Figure 1.5 represents the spatial partition of the origins and destinations of the matched riders. It shows that riders' origins are mainly located in the closest to Paris, where population density is the highest. The intuition is that it is easier to find a matching driver within reasonable walking distance in areas where population density is higher.





(a) Riders' origins



(b) Riders' destinations

Figure 1.5: Origins and destinations of matched riders in Île-de-France

The total walking distance displayed on Figure 1.6 represents the sum of the Euclidean distances from the rider's origin to the pick-up intersection and from the drop-off intersection to the rider's destination. The walking distance is zero for about 28% of matches, meaning that the rider and the driver have the same origin-destination pair. Almost all riders have a walking distance smaller than 1 km, which corresponds to a walking time of less than 15 minutes.

Figure 1.7 presents the schedule delay of riders. For around 66% of riders, the schedule delay is zero, i.e., they arrive at destination within their on-time window of 10 minutes. More riders are arriving early than late (arriving early is less costly than arriving late because

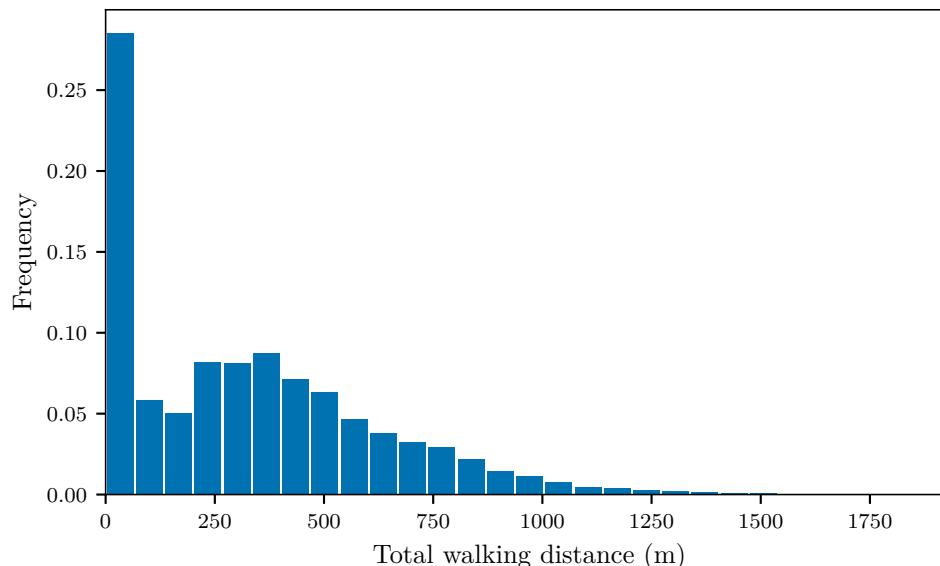


Figure 1.6: Distribution of walking distance for the matched riders

$\beta < \gamma$ ).

Figure 1.8 presents the distribution of the generalised cost savings for riders, i.e., the difference between their generalised travel cost as a rider and their generalised travel cost in the reference scenario. It is the sum of the schedule delay cost, the walking cost, and the in-vehicle travel cost. An analysis of the ride-sharing cost reveals that the main component is the in-vehicle travel cost: the mean cost of 5.09 € can be divided in 16 % of schedule delay cost, 10 % of walking cost, and 74 % of in-vehicle travel cost.

Recall that the drivers are perfectly compensated for the inconvenience cost of having someone in their car and that the riders do not receive any subsidy. The cost of implementing such a ride-sharing scheme is thus equal to the sum of the inconvenience cost of all the drivers. As the literature on ride-sharing is still lacking good estimates of the inconvenience cost of having someone in their car, we cannot estimate precisely the cost of the policy that we propose. Instead, we analyse the results with different values for the inconvenience cost, ranging from 2 €/h to 16 €/h.<sup>5</sup> The results are reported in Table 1.3. Even with a large

<sup>5</sup>Although inconvenience in public transit and with ride-sharing are not directly comparable, the value of crowding in public transit can be used as an approximation for the inconvenience cost of ride-sharing. Björklund and Swärdh (2017) estimates that, when seated, the value of time is multiplied by 1.48 when




Figure 1.7: Distribution of schedule delay for the riders

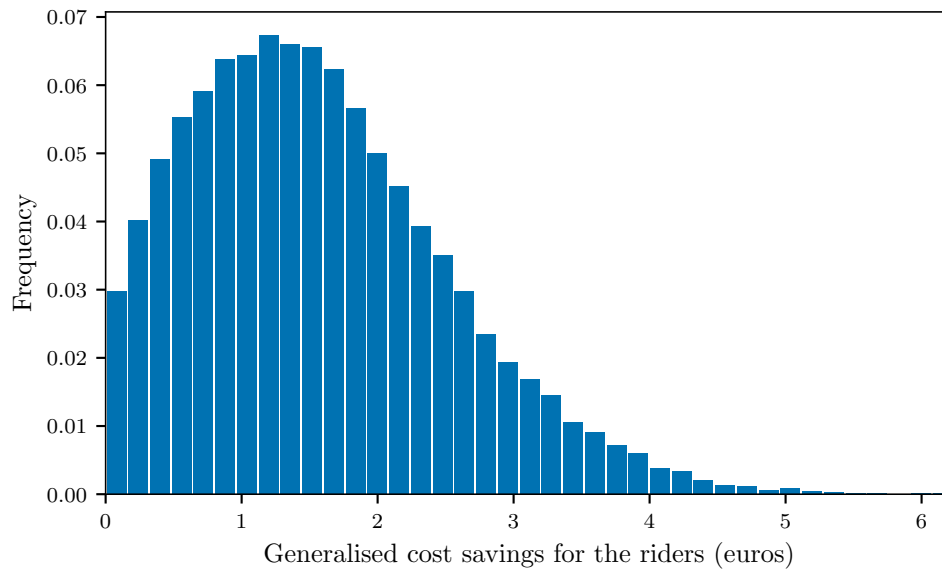


Figure 1.8: Distribution of the decrease of the riders' generalised cost, compared to the reference scenario

inconvenience cost of 16 €/h, the increase in individual surplus is larger than the cost of the policy. This suggests that the policy has a positive social impact, even before accounting for CO<sub>2</sub> emissions reduction and the long-run impact of ride-sharing modal share.

Inconvenience cost (€/ hour)	2	4	8	16
Average subsidy for drivers (€)	0.28	0.56	1.11	2.22
Total policy cost (€) [I]	8667.54	17 335.09	34 670.18	69 340.35
Individual surplus increase (€) [II]	104 896.96	104 896.96	104 896.96	104 896.96
Total surplus variation (€) [II – I]	96 229.42	87 561.87	70 226.79	35 556.61

Table 1.3: Average subsidy and cost of the policy as a function of the inconvenience cost for the drivers

### 1.4.5 Allowing Multiple Riders in the Same Car

The results presented so far assume that all the drivers accept at most one rider in their car. In practice, most cars can hold up to 4 passengers (excluding the driver). Allowing more than one rider in each car might increase the inconvenience cost for the driver but it should also increase the number of matches and match quality.

Table 1.4 presents the results of the simulations for the 30% participation rate scenario, assuming that each driver can accept 1, 2 or 3 riders in their car. The share of riders increases from 3.3% with at most 1 passenger to 4.0% with at most 2 passengers, implying a reduction in the number of vehicles on the road and thus a reduction in congestion and CO<sub>2</sub> emissions. With at most 3 passengers per car, the ride-sharing share increases again to 4.2% but the decrease of the number of cars is less significant.

Although the share of riders is increasing with the maximum number of passengers per car, the share of ride-sharing drivers (i.e., drivers with a least one passenger) is decreasing at the same time, because each driver is carrying more passengers. Also, match quality is increasing because more riders are matched with the best potential driver for them.

---

shifting from a situation with no crowding to a situation with overcrowding. In our model, the value of time is 12.96 €/h. A multiplier of 1.48 implies an inconvenience cost of 6.22 €/h.

<b>Passengers per driver</b>	1	2	3
<b>Shares</b>			
Transit modal share	24.3 %	24.1 %	24.0 %
Car modal share	72.4 %	71.9 %	71.8 %
Ride-sharing modal share	3.3 %	4.0 %	4.2 %
<b>Surplus</b>			
Individual surplus variation (€)	+104 897	+126 508	+137 374
<b>Road network</b>			
Car VKT (10 <sup>3</sup> km)	10 595	10 534	10 538
Congestion	20.6 %	20.1 %	19.6 %
CO <sub>2</sub> emissions reduction (tons)	39.372	51.145	50.373
<b>Drivers</b>			
Mean travel time	15' 27"	15' 22"	15' 20"
Mean travel cost (€)	6.00	5.98	5.96
<b>Ride-sharing drivers</b>			
Share of ride-sharing drivers	3.3 %	2.7 %	2.4 %
Average number of passengers	1.0	1.5	1.7
Share of time spent with a passenger	58.0 %	59.4 %	59.7 %
<b>Riders</b>			
Mean Euclidean OD distance (meters)	6205	6174	6164
Mean walking distance (meters)	325	325	327
Mean walking distance (conditional on it being positive, meters)	438	438	440
Mean car travel time	8' 20"	8' 23"	8' 23"
Mean travel time	13' 13"	13' 15"	13' 17"
Mean travel cost (€)	3.22	3.26	3.27
Riders at their best match	65.0 %	72.7 %	76.0 %

Table 1.4: Comparison of results when drivers can have at most 1, 2 or 3 passengers in their car, for 30 % of people willing to participate in the ride-sharing scheme

Note. The share of time spent with a passenger is the average share of time spent with each passenger, for ride-sharing drivers.

### 1.4.6 Proposing Incentives to the Riders

The analysis conducted so far shows that a larger ride-sharing share implies less congestion and CO<sub>2</sub> emissions. Therefore, the governments might be interested in subsidising ride-sharing in order to further increase the ride-sharing share. In this section, we assess the efficiency of proposing subsidies to travellers to induce them to switch to ride-sharing.

More formally, we assume that the government gives a fixed amount of money to each traveller, each time they travel by ride-sharing (as a rider) instead of taking their car or the public transit services. Table 1.5 shows the results of the simulations for a subsidy amount

of 0.5€, 1.0€ and 1.5€. The incentive is effective at increasing the share of riders, from 3.3% with no incentive to 4.4% with an incentive of 1.5€.

Compared to the scenario with no incentive, the individual surplus varies for three reasons (the first two have a positive impact, the third one has a negative impact):

1. Riders are receiving a subsidy which decreases their generalised travel cost.
2. There are less cars on the road, which decreases the generalised travel cost of all drivers.
3. Match quality decreases because the number of matched riders increases while the number of potential drivers stays constant.

Table 1.5 shows that the two first effects largely dominate the third one as the individual surplus is increasing by around 54 000€ from the scenario without incentive to the scenario with an incentive of 1.5€. However, the variation of the total surplus (which account for the amount of subsidies spent by the government) is more ambiguous. From the total surplus, it is unclear whether subsidies for ride-sharing have the desired positive effect. The main reason for this is that subsidies are awarded to ride-sharing participants coming from both private and public transportation. Whereas those coming from private transportation have a positive external effect in reducing congestion, those coming from public transportation may deteriorate the match quality and may therefore negatively influence the total surplus.

## 1.5 Conclusion

Ride-sharing is a tool with a great potential to reduce pollution and congestion in urban areas. It nevertheless remains unpopular amongst commuters, despite a growing number of ride-sharing apps. In this paper, we propose a ride-sharing scheme that any traveller can subscribe to and that allows them to be picked up by a driver for free or to pick up passengers (as a driver) in exchange for a subsidy. In this scheme, drivers are completely inflexible, i.e., they can keep the same route and schedule with and without passengers.

This study proposes a state-subsidised ride-sharing scheme to increase the modal share

Incentive amount per rider	0€	0.5€	1€	1.5€
<b>Shares</b>				
Transit modal share	24.3%	24.2%	24.1%	23.9%
Car modal share	72.4%	72.1%	71.8%	71.7%
Ride-sharing modal share	3.3%	3.7%	4.1%	4.4%
<b>Surplus</b>				
Individual surplus variation (€)	+104 897	+127 972	+142 444	+158 876
Expenses (€)	0	17 421	38 132	61 003
Total surplus variation (€)	+104 897	+110 551	+104 312	+97 873
<b>Road network</b>				
Car VKT (10 <sup>6</sup> km)	10 595	10 567	10 543	10 539
Congestion	20.6%	20.3%	20.2%	19.9%
CO <sub>2</sub> emissions reduction (tons)	39.372	44.776	49.408	50.180
<b>Drivers</b>				
Mean travel time	15' 27"	15' 25"	15' 26"	15' 23"
Mean travel cost (€)	6.00	5.99	5.99	5.98
Share of time spent with a passenger (for ride-sharing drivers only)	58.0%	55.1%	52.9%	51.1%
<b>Riders</b>				
Mean Euclidean OD distance (meters)	6205	6077	6010	5970
Mean walking distance (meters)	325	366	406	449
Mean walking distance (conditional on it being positive, meters)	438	473	508	547
Mean car travel time	8' 20"	8' 10"	8' 03"	7' 58"
Mean travel time	13' 13"	13' 39"	14' 9"	14' 41"
Mean travel cost (excluding the subsidy, €)	3.22	3.34	3.48	3.63
Riders at their best match	65.0%	60.7%	56.2%	52.9%

Table 1.5: Comparison of results with different incentive amount, for 30% of people willing to participate in the ride-sharing scheme (at most 1 passenger per car)

of ride-sharing. The potentials of this scheme are tested on the Île-de-France region to evaluate the individual and social benefits. Drivers and riders are matched through a linear-programming algorithm, based on their itineraries and preferences. The ride-sharing scheme induces a significant reduction of congestion and CO<sub>2</sub> emissions, due to a smaller number of vehicles on the roads. The results might be further improved by considering these externalities directly in the objective function of the matching algorithm.

The ride-sharing scheme we propose considers, since no detour nor extra schedule delays are involved, that the vast majority of drivers would be ready to pick up someone in their car in exchange for a small monetary incentive. This state subsidy then only compensates for the inconvenience cost of sharing a car. Riders need to walk, but benefit from a free ride.

Their individual savings are gasoline saving, time and monetary saving related to parking, and wear and tear (beside the reduced congestion).

In this paper, we consider the morning commute, which has to be seen as an intermediary step for two reasons. First, the evening commute is not a mirror case of the morning commute in dynamic models (as shown for example by de Palma and Lindsey, 2002). Second, if a user decides not to take his / her car the morning, he/she has to do ride-sharing or take public transport in the evening. Moreover, if the schedule preferences of two matched users are similar in the morning, this does not necessarily mean they will be similar in the evening for the same driver/rider couple. So, in general, the same match could not be arranged in the morning and in the evening. As a consequence, the riders are not guaranteed to find another convenient match for their return trip in the evening. Matching for round-trip commuting is a constraint that could be studied in future research. A mathematical formulation of this problem was proposed by de Palma and Nesterov (2006), in the case of stable-dynamic models.

Also, the analysis conducted only explored the benefits of ride-sharing due to reduced road traffic. The generalised travel cost for public transit was assumed to be independent of the number of public-transit users. In reality, a reduction of the number of public-transit users can have an impact on in-vehicle congestion and the operating costs of the public-transit services.

To keep the analysis simple, this research only considered that riders take one car to complete their trip. Allowing for riders to take multiple vehicles to make their trip can increase the potential of ride-sharing. This type of ride-sharing, referred to as multi-hop ride-sharing, has been recently investigated (Herbawi and Weber, 2012; Teubner and Flath, 2015). Even though multi-hop ride-sharing generates transfer penalties between cars, it could be interesting for some segments of a network, in particular for OD pairs with low demand. The combinatorial issues (the multiple matching problem) remain widely unexplored in the matching literature in economics (labour market and marriage market, for rather obvious



reasons).

Finally, it remains to be seen if riders and drivers are prepared to be involved in a more complex organization, with potential safety concerns. Empirical research is needed in order to evaluate the acceptance of such a system. A mobile application could be developed to mitigate the complexity and provide certifications guaranteeing the safety of the system. Such analysis is out of the scope of the present article.

## Chapter 2

# METROPOLIS2: Bridging Theory and Simulation in Agent-Based Transport Modeling

This chapter is available as a working paper:

Javaudin, L., & de Palma, A. (2024). METROPOLIS2: Bridging Theory and Simulation in Agent-Based Transport Modeling. Technical Report 2024-03, THEMA (Théorie Economique, Modélisation et Applications), CY Cergy Paris Université.

### Abstract

Transport simulators can be used to compute the equilibrium between transportation demand and supply within complex transportation systems. However, despite their theoretical foundations, there is a lack of comparative analysis between simulator results and theoretical models in the literature. In this chapter, we bridge this gap by introducing METROPOLIS2, a novel mesoscopic transport simulator capable of simulating agents' travel decisions (including mode, departure-time, and route choice), based on discrete-choice theory within a dynamic, continuous-time framework. We demonstrate METROPOLIS2's functionality through its application to the single-road bottleneck model and validate its ability to replicate analytical results. Furthermore, we provide a comprehensive overview of METROPOLIS2 in large-scale scenarios. Finally, we compare METROPOLIS2's results with those of the original METROPOLIS1 simulator in a simulation of Paris, highlighting its speed and ability to converge to an equilibrium.

**Keywords:** transport simulation; agent-based modeling; bottleneck; dynamic traffic assignment; discrete-choice models

**JEL Codes:** C63; R4

## 2.1 Introduction

Transportation science explores the complex interplay between demand and supply within transportation systems. Demand encompasses individuals, seeking to travel from one location to another, each with their own preferences such as value of time and scheduling constraints, along with some limitations imposed by factors like car ownership or public-transit subscription. Supply encompasses the transportation infrastructure, comprising a road network defined by regulations (e.g., speed limits, traffic signals, intersection priorities) and congestion dynamics (e.g., bottlenecks, shock waves, queueing), a public transit system with associated timetables, available parking spaces, etc.

A large portion of transportation science literature is dedicated to analyzing equilibrium properties and evaluating transport policies through analytical models. These models, while tractable, only capture certain facets of equilibrium dynamics. For example, the bottleneck model (Vickrey, 1969; Arnott et al., 1990, 1993) sheds light on the impact of schedule constraints and road tolls using a model featuring a single road and a continuum of identical individuals. With the advancement of computing power in recent decades, transportation simulators are being used to analyze specific scenarios or policies, considering the complex interplay of various effects (Nguyen et al., 2021). These simulators range from microscopic models examining traffic infrastructure impacts at the neighborhood level (Lopez et al., 2018) to macroscopic models investigating aggregated effects at city or national levels (McNally, 2007).

Mesoscopic models like METROPOLIS (de Palma et al., 1997) and MATSim (Horni et al., 2016) adopt an intermediate approach between microscopic and macroscopic simulations. Relying on an agent-based methodology, they account for heterogeneous effects while simplifying congestion modeling (e.g., omitting lane changing or car following behaviors) to enhance scalability for large-scale scenarios. This chapter introduces METROPOLIS2, a novel mesoscopic transport simulator that extends the legacy of its predecessor, METROPOLIS, by drawing inspiration from similar mesoscopic modeling frameworks.

METROPOLIS2's key characteristics include the simulation of agents' travel decisions (mode, departure-time, and route) within a dynamic, continuous-time framework. The simulator is adapted to simulate and evaluate transport policies at the city or national level, such as low-emission zones, infrastructure changes (e.g., alteration in speed limit, introduction of new public-transit lines) or pricing adjustments (e.g., fuel cost taxation, modification in public-transit fares). Grounded in economic and discrete-choice theories, METROPOLIS2 is well-suited for computing agents' surplus in cost-benefit analyses. Moreover, it can derive insights from analytical examples, as demonstrated in Section 2.3 and Appendix 2.A.

METROPOLIS2 draws inspiration from and extends the transport simulator METROPOLIS (de Palma et al., 1997; de Palma and Marchal, 1999, 2002). To avoid confusion, the original simulator METROPOLIS is referred to as METROPOLIS1 throughout the chapter. Compared to its predecessor, METROPOLIS2 offers additional features including trip chaining, more flexible utility specifications, various vehicle types and explicit bottleneck queues. Moreover, it relies on state-of-the-art routing algorithms (Geisberger and Sanders, 2010; Batz et al., 2013), resulting in improved speed.

The primary difference between METROPOLIS2 and MATSim (Horni et al., 2016), a popular agent-based transport simulator with an activity-based approach, lies in the convergence algorithm used to reach an equilibrium. While MATSim employs a co-evolutionary algorithm where agents select the best alternative from a small agent-specific choice set at each iteration, with a low probability of switching to a different random alternative outside of this choice set, METROPOLIS2 agents choose the best alternative from the complete choice set at each iteration. Consequently, MATSim tends to produce a more stable system with fewer changes between iterations but converges slower to an equilibrium, compared to METROPOLIS2. MATSim may also become trapped in suboptimal equilibria if the parameters of the co-evolutionary algorithm are not properly adjusted. For an in-depth comparison between MATSim, METROPOLIS1 and METROPOLIS2, refer to Section 2.2.

The chapter is structured as follows. Section 2.2 conducts a literature review on meso-

scopic transport simulators and similar models. Section 2.3 introduces a simplified version of METROPOLIS2, tailored to simulate the single-road bottleneck model, highlighting its foundation in economic and discrete-choice theory and its ability to replicate results from analytical models. Section 2.4 provides an overview of the general version of METROPOLIS2, including a presentation of the convergence algorithm, an explanation of the features modeled on the demand and supply side, and a discussion of the equilibrium concept. Section 2.5 delves into implementation details for the simulator, such as programming language, code parallelization and testing. Section 2.6 presents a comparison with METROPOLIS1 using a simulation of Paris’s urban area. Finally, Section 2.7 offers concluding remarks.

## 2.2 Literature Review

This section starts with a comparison of METROPOLIS2 with two other transport simulators: METROPOLIS1 and MATSim. Additional related works are discussed at the end of the section.

METROPOLIS1 is a dynamic mesoscopic transport simulator, whose original version was presented by de Palma et al. (1997). METROPOLIS1 is the main source of inspiration for METROPOLIS2. METROPOLIS1 can simulate mode choice (between car and public transit), departure-time choice (with a Continuous Logit) and route choice for a population of agents all performing a single trip. In METROPOLIS1, congestion is simulated using an event-based model relying on speed-density functions. Many applications of METROPOLIS1 have been proposed in the last decades (e.g., road pricing with de Palma et al. 2005, vehicle-emission pricing with Vosough et al. 2022, ride-sharing with de Palma et al. 2022a).

MATSim is defined as “an open-source framework for implementing large-scale agent-based transport simulations.”<sup>1</sup> A detailed description of this simulator is provided in Horni et al. (2016). MATSim has been used in many cities and regions around the world and its architecture is close to the one of METROPOLIS2, which justify including it in the detailed

---

<sup>1</sup><https://matsim.org/>, last accessed on 23d January 2024.

comparison.

Tables 2.1, 2.2, 2.3 and 2.4 provide a comparison between the three simulators METROPOLIS2, METROPOLIS1 and MATSim on various areas. Table 2.1 compares the simulators with respects to some general and software-related characteristics. While METROPOLIS2's development started in 2022, METROPOLIS1 was initially released around 1997 and the Java version of MATSim was released in 2010 (with preliminary versions dating back from the 1990s, see Nagel 1996). The three simulators each use different programming languages and different input / output formats.

Table 2.1: Comparison of 3 transport simulators: General and software

	METROPOLIS2	METROPOLIS1	MATSim
Reference	This chapter	de Palma et al. (1997)	Horni et al. (2016)
Initial release	2022 (version 0.1.0)	Around 1997	2010 (version 0.1.0)
Language	Rust	C++	Java
Open source	Undecided	No	Yes
Input / output format	CSV or Parquet	MySQL	XML

Table 2.2 compares how demand is defined in the simulators. METROPOLIS1 generates the population from origin-destination matrices, while METROPOLIS2 and MATSim represent the population as a list of agents. Therefore, METROPOLIS1 defines trips between zones, while METROPOLIS2 and MATSim define them between network nodes or links (but zone-to-zone trips can also be simulated as a special case). The latter two models differ in that, for each agent, METROPOLIS2 considers a list of trips, while MATSim considers a list of activities. However, since converting between trips and activities is straightforward (activities represent the time spent between trips and trips represent the way to connect activities), this difference has no practical consequences. One limitation is that METROPOLIS2 assumes that the activity duration is exogenous.

With regards to the definition of the utility function, METROPOLIS2 uses a polynomial function of travel time, with a linear schedule-delay cost at origin and at destination (see

Appendix 2.E.1). Other specifications could be easily added in the future. METROPOLIS1' uses the  $\alpha$ - $\beta$ - $\gamma$  model (Arnott et al., 1990), with a schedule-delay cost either at origin or at destination (this is a special case of METROPOLIS2). MATSim's definition of utility (referred to as "score") is based on Charypar and Nagel (2005): utility depends on activities' duration, activities' schedule delays, trips' travel time and trips' euclidean distance. The parameters of these utility functions can be specific to each agent in METROPOLIS2, while they are defined at the class- (or subpopulation-) level in METROPOLIS1 and MATSim.

Table 2.2: Comparison of 3 transport simulators: Demand definition

	METROPOLIS2	METROPOLIS1	MATSim
Population	Agents with 1+ trips	Origin-destination matrix by subpopulation (single trip)	Agents with 1+ activities
Origin / destination definition (for road modes)	Node	Zone centroid	Link
Activity duration	Exogenous	N/A	Endogenous
Utility / generalized cost	Function of travel time and schedule-delay costs <sup>a</sup>	Generalized cost (Vickrey model)	Charypar-Nagel utility function <sup>b</sup>
Preference classes	Parameters defined at the agent-level	Distribution of parameters at the class-level	Constant parameters at the class-level

<sup>a</sup> The default utility function is a polynomial function of travel time with early / late schedule-delay costs at origin and at destination. Other specifications can be easily added.

<sup>b</sup> Utility of activities (logarithmic function of duration with early and late penalties) and utility of trips (linear function of travel time and euclidean distance).

Table 2.3 provides a comparison of the three simulators in terms of their choice models. A fundamental difference lies in the fact that both METROPOLIS2 and METROPOLIS1 are grounded in discrete-choice theory, whereas MATSim employs a co-evolutionary algorithm. In METROPOLIS2 and METROPOLIS1, agents select the alternative that maximizes their utility among all available alternatives (given the expected network conditions). In contrast, MATSim maintains a list of plans (typically 5) in the "memory" of each agent. Each plan represents an alternative (a combination of modes, departure times and routes). Agents in MATSim then select one of these plans based on their perceived utility, or they "evolve" by



choosing a new plan (usually generated by applying a random perturbation to an existing plan). While the co-evolutionary algorithm in MATSim offers a more realistic depiction of agent behavior (although this is not the primary rationale behind its use, as noted by Flötteröd 2016), it comes with certain drawbacks: (i) convergence to an equilibrium may be slower compared to METROPOLIS1 and METROPOLIS2 due to agents' slower adaptation and (ii) computing agent-level surplus for cost-benefit analysis is not as straightforward (see Kickhöfer and Nagel 2016).

METROPOLIS1 and METROPOLIS2 are very similar, with two notable exceptions. First, METROPOLIS2 offers more flexibility than METROPOLIS1: while METROPOLIS1 is limited to a binary Logit mode choice and a Continuous Logit departure-time choice, METROPOLIS2 allows for various discrete-choice models. Second, route selection in METROPOLIS1 involves making decisions at every intersection, while, in METROPOLIS2, the complete route is selected prior to departure from the origin (which enables the use of more efficient routing algorithms).

Table 2.3: Comparison of 3 transport simulators: Choice models

	METROPOLIS2	METROPOLIS1	MATSim
Basic principle	Discrete-choice theory	Discrete-choice theory	Co-evolutionary algorithm
Choice of mode	Any discrete-choice model (random or deterministic)	Binary Logit model	Random shifts (when mode innovation is triggered) <sup>a</sup>
Choice of departure time	Various specifications <sup>b</sup>	Continuous Logit	Random shifts (when time innovation is triggered)
Choice of route	Fastest path at the time of departure	Choice at each intersection (deterministic or stochastic)	Fastest path at the time of departure (when route innovation is triggered)
Choice of destination	Yes (with exogenous destinations)	No	With a module (Horni et al., 2011)

<sup>a</sup> An alternative mode-choice model based on discrete-choice theory is available as a MATSim module (see Hörl et al. 2018 and Hörl et al. 2019).

<sup>b</sup> The choice models include Multinomial Logit and Continuous Logit. Departure-time choice can also be deactivate (i.e., departure time is exogenous).

Table 2.4 compares how congestion is modeled and simulated in the three simulators. In

this comparison, we only consider the QSim mobsim<sup>2</sup> for MATSim (the default and most used mobsim). MATSim, with QSim, uses a time-step based model, which means that at each time step (1 second by default) all the vehicles currently on the network get an “update” to move them. In contrast, METROPOLIS1 and METROPOLIS2 use an event-based model, which means that vehicles are “updated” only when an event happens (e.g., they reach the end of their current link). Apart from the typical road vehicles (such as car and trucks), MATSim has the capability to simulate public-transit vehicles. This means that (i) buses can contribute traffic congestion and are also susceptible to being affected by congestion caused by other vehicles and (ii) agents are forced to wait when a train vehicle is full. Contrarily to METROPOLIS1 and MATSim, METROPOLIS2 represents time as a continuous variable, which allows for greater accuracy in the computation.<sup>3</sup> Congestion is simulated in METROPOLIS1 via speed-density functions (the speed of a vehicle entering a road is a function of the density of vehicles currently on this road) and in MATSim via bottleneck queues (the outflow of vehicles exiting a road is limited by the road’s capacity; a queue builds up if the flow is larger than the capacity). METROPOLIS2 can make use of either speed-density functions or bottleneck queues, or a combination of both. The three models support queue propagation: when the available space on a road is smaller than the incoming vehicle length, the vehicle is stuck at his current location until enough space is freed. Both METROPOLIS2 and MATSim support different vehicle types, with different length, passenger car equivalent, speed limits and road restrictions. Finally, both METROPOLIS1 and MATSim support tolls and time-varying networks (e.g., capacity of a road is decreased at an exogenous time).

We now discuss SimMobility (Adnan et al., 2016), an open-source C++ agent-based simulator that shares many similarities with METROPOLIS2. SimMobility is divided in three main components: short-, mid- and long-term. While the short-term component mir-

---

<sup>2</sup>A MATSim mobsim, or mobility simulation, is a program responsible for the network loading of the model, i.e., it is a program which simulates the movement of vehicles on the network.

<sup>3</sup>In METROPOLIS2, for a road with a capacity of 1440 vehicles per hour, at most 1 vehicle each  $3600/1440 = 2.5$  seconds is allowed to enter / exit the road. This is not easy to simulate in a discrete-time model without introducing some randomness or rounding errors.

Table 2.4: Comparison of 3 transport simulators: Supply

	METROPOLIS2	METROPOLIS1	MATSim (QSim)
Model	Event-based	Event-based	Time-stepping (1s step by default)
Simulated modes	Road and teleport <sup>a</sup>	Road and teleport <sup>a</sup>	Road, teleport <sup>a</sup> and public-transit
Time	Continuous	Discrete (1s step)	Discrete (1s step)
Speed-density functions	Yes	Yes	No
Bottleneck queues	Yes	No	Only for outflow
Queue propagation (spillback)	Yes (configurable backward propagation speed)	Yes (instantaneous backward propagation)	Yes (backward propagation speed fixed to 15 km/h)
Vehicle types	Yes (with different length, PCE, speed limit)	No	Yes (with different length, PCE, speed limit)
Road restrictions	Yes	No	Yes
Tolls	Partial <sup>b</sup>	Yes	Yes
Network time variations	No	Yes	Yes

<sup>a</sup> A “teleport” mode is a mode without any interaction with the other agents. For example, walking and bicycling are usually considered as teleport modes since the travel time is constant. Public transit can also be considered as a teleport mode if one assumes that the schedules are not impacted by car traffic or in-vehicle congestion.

<sup>b</sup> In METROPOLIS2, simulating a single tolled road is possible with a strategy employing road restrictions and vehicle types: agents can select between a vehicle restricted from the tolled road or a vehicle able to take any road but incurring a fixed monetary penalty (equivalent to the toll amount). Expanding to multiple tolled roads involves an increasing complexity and running time, scaling combinatorially with the number of tolls.

rors a microscopic model and the long-term component represents a land-use model, the mid-term component, SimMobilityMT (Lu et al., 2015), adopts a mesoscopic approach akin to that of METROPOLIS2. The research focus of SimMobilityMT has primarily revolved around mobility-on-demand applications (Basu et al., 2018; Oh et al., 2020). Similar to METROPOLIS2, the demand model of SimMobilityMT employs discrete-choice methods. SimMobilityMT uses a fully econometric activity-based model to simulate the daily activity patterns of agents, relying on a hierarchy of discrete-choice models. In contrast to METROPOLIS2’s deterministic route choice model, SimMobilityMT employs a path-size Logit model

for route choice. Moreover, it allows for within-day decisions, enabling agents to adjust their routes or activities based on observed network conditions throughout the (simulated) day. Regarding the supply model, both METROPOLIS2 and SimMobilityMT simulate congestion using a combination of speed-density functions and bottlenecks, incorporating queue propagation. Both simulators are also able to simulate different vehicle types on the road network. However, METROPOLIS2 uses a continuous-time event-based model, whereas SimMobilityMT relies on a time-stepping model. Finally, while METROPOLIS2 represents the demand-supply equilibrium as a solution to a fixed-point problem, the literature concerning SimMobilityMT cited so far does not explicitly mention the notion of equilibrium.

For an extended literature review on agent-based transport simulators, the interested readers can refer to the work of Nguyen et al. (2021). METROPOLIS2 also share similarities with commercial transport software like Aimsun (Barceló and Casas, 2005) or PTV Visum.

Our work shares similarities with the study conducted by Otsubo and Rapoport (2008), who propose a discrete version of the bottleneck model and observe that the solution to the symmetric mixed-strategy Nash equilibrium diverges from the equilibrium solution of the continuous model. In Section 2.3, we investigate a discretized version of the bottleneck model with a finite number of agents. Our model differs from theirs in two key aspects: (i) we maintain time as a continuous variable, and (ii) we introduce stochasticity in the departure-time decision process. Moreover, our findings shows that the results remain consistent whether considering a finite number of agents or a continuum of agents (see Appendix 2.A).

Finally, Guo et al. (2018) investigate the bottleneck model and conclude that an equilibrium cannot be attained through a day-to-day learning process like proportional swap. However, our findings present a more optimistic outlook, demonstrating that we can achieve numerical proximity to an equilibrium with METROPOLIS2, using the same bottleneck model with three differences: (i) a finite number of agents, (ii) the introduction of stochasticity in the departure-time choice, and (iii) discretization of the travel-time function.

## 2.3 From Models to Simulations

This section explains the challenges encountered in transport simulations and provides an overview of how METROPOLIS2 operates on a basic example. All the analysis in this section rely on the bottleneck model.

### 2.3.1 The Bottleneck Model

The so-called bottleneck model is a transport model which was derived analytically by Vickrey (1969) and Arnott et al. (1990). The model consists in a single road, a single mode (car) and identical individuals whose utility is a linear function of travel time and schedule delay. Numerous extensions of the bottleneck model have been proposed in the literature (Li et al., 2020b). We provide below a brief description of this bottleneck model.

There is a continuum of  $N$  individuals traveling from an origin  $A$  to a destination  $B$ , via a single road. This road consists in a free-flow travel time  $t^f$ , followed by a bottleneck of capacity  $s$  with a vertical queue, as illustrated in Figure 2.1. The travel time from  $A$  to  $B$  can thus be expressed as

$$T(t) = t^f + \frac{Q(t + t^f)}{s},$$

where  $t$  is the departure time from  $A$  and  $Q(t)$  is the length of the bottleneck's queue at time  $t$ .




Figure 2.1: Illustration of the road network in the bottleneck model

Individuals choose their departure time from the origin to maximize their utility, given by

$$V(t) = -\alpha \cdot T(t) - \beta \cdot [t^* - t - T(t)]_+ - \gamma \cdot [t + T(t) - t^*]_+, \quad (2.1)$$

where  $\alpha$  is the value of travel time,  $\beta$  is the value of time for being early at destination,  $\gamma$  is the value of time for being late at destination,  $t^*$  is the desired arrival time at destination and  $[x]_+ = \max(x, 0)$ .

Departure time is the only decision variable in the bottleneck model, as there is neither mode choice (all individuals travel by car) nor route choice (only one road connects  $A$  to  $B$ ). An equilibrium is reached when no individual has an incentive to unilaterally change their departure time. The analytical solution to the equilibrium of the model is found by relying on the following property:

$$\begin{cases} V(t) = \bar{V} & \text{if } \bar{r}^d(t) > 0, \\ V(t) \leq \bar{V} & \text{if } \bar{r}^d(t) = 0, \end{cases}$$

where  $\bar{r}^d(t)$  is the equilibrium rate of departures at time  $t$  and  $\bar{V}$  is the equilibrium utility level. This property implies that at equilibrium, all individuals get the same utility level. The analytical solution of the model is presented in Arnott et al. (1990). The equilibrium rate of departures and arrivals and the equilibrium utility function are represented on Figure 2.2.

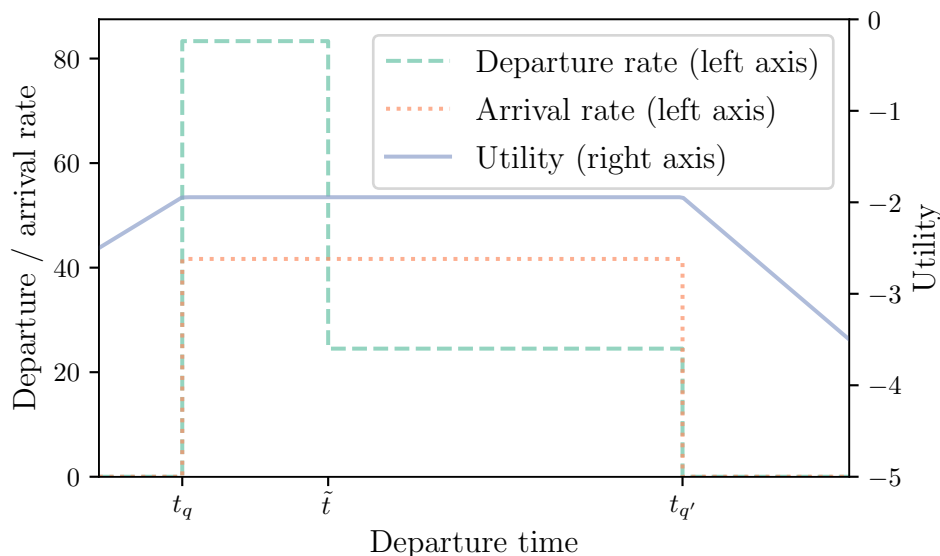


Figure 2.2: Analytical solution of the bottleneck model

Note:  $t_q$  is the time at which congestion starts to appear,  $\tilde{t}$  is the departure time such that the individual arrives on time,  $t_{q'}$  is the time at which congestion ends.

### 2.3.2 Challenges of Simulation

Simulating a model, even as basic as the bottleneck model, presents several challenges that requires specific adaptations and approximations. First, the analytical bottleneck model assumes a continuum of identical individuals, allowing for departure times to be represented as a continuous distribution. In practice, discretization of the model is required for simulations. Indeed, simulations are limited to a finite number of distinct individuals, typically referred to as “agents”. Each agent is treated as a separate entity with their own departure time. Consequently, departure-time distribution in simulations becomes discrete, whereas the analytical model assumes a continuous distribution. See Otsubo and Rapoport (2008) for a numerical solution of the bottleneck model with discrete agents. Note that considering a discrete number of separate individuals is actually more realistic than a continuum of individuals and it allows to more easily deal with individual heterogeneity.

Second, the analytical model assumes a population of perfectly identical individuals all choosing a departure time that maximizes their deterministic utility. Equilibrium conditions, as demonstrated by Arnott et al. (1990) and illustrated on Figure 2.2, reveal that there is a continuous range of departure times that yield maximum utility. While the analytical solution of the model provides a single solution for the rate of departures at equilibrium, there exists an infinity of solutions regarding the order in which agents depart. In simulations, whenever there are multiple departure times maximizing utility, there is no easy way to spread the individuals over these departure times so that the equilibrium conditions are satisfied. To address this, randomness is introduced in the simulation’s decision-making process. This stochastic element encourages individuals to spread their departure times across the equilibrium departure-time range. One approach to introduce randomness is by adding a random component to the utility, similar to the methods used in discrete-choice models. In this vein, de Palma et al. (1983) analyze the bottleneck model where departure-time choice is modeled as a Continuous Logit model.

Third, travel-time function  $T(t)$  and utility function  $V(t)$  are continuous functions that

may not always have explicit analytical forms. In such cases, simulations need to resort to numerical approximations. One common approach is to represent these continuous functions as piecewise linear functions. This approximation involves dividing the function into linear segments, with approximation errors depending on the length of these linear segments.

### 2.3.3 Simulating the Bottleneck Model

METROPOLIS2 is able to simulate various small-scale and large-scale models with the goal of finding an equilibrium of the simulated model. For the bottleneck model, achieving equilibrium entails finding a departure time  $t_n^d$ , for each agent  $n$ , such that no agent has an incentive to choose a different departure time, given the travel-time function  $T(t)$  derived from the agents' collective departure times. We now describe how the bottleneck model presented in Section 2.3.1 can be simulated.

We assume a population of  $N$  discrete agents, all traveling from origin  $A$  to destination  $B$ . The utility of agent  $n \in \{1, \dots, N\}$  is given by

$$U_n(t) = V(t) + \varepsilon_n(t),$$

where  $t$  is the departure time from the origin,  $V(t)$  is the deterministic utility function, defined as in equation (2.1), and  $\varepsilon_n(t)$  is an agent-specific random component of utility. All agents are identical except for the random components  $\varepsilon_n(t)$ . The single road connecting origin  $A$  to destination  $B$  has a free-flow travel time  $t^f$  and a congested travel time described by a bottleneck with capacity  $s$ . This means that, at most  $s$  cars can cross the bottleneck during each time unit, or, equivalently, that two cars cannot follow each other in less than  $1/s$  time units. As discussed in Section 2.3.2, the  $N$  agents are treated as separate entities and the travel-time function  $T$  and the utility functions  $U_n$  are both represented as a piecewise linear functions.

The model simulated is split in a demand model, corresponding to the travel behavior of



the agents simulated, and a supply model, corresponding to the road-network infrastructure operations.

**Demand model** The demand model consists in the agents choosing a departure time for their trip, given the expected travel-time function from origin to destination.

Following de Palma et al. (1983), we assume that the distribution of  $\varepsilon_n(t)$  is such that the probability that agent  $n$  chooses departure time  $t$  is given by the Continuous Logit model (Ben-Akiva and Watanatada, 1981; Ben-Akiva et al., 1985):

$$p_n^d(t) = \frac{e^{V(t)/\mu}}{\int_{t^0}^{t^1} e^{V(\tau)/\mu} d\tau}, \quad (2.2)$$

where  $\mu$  represents the randomness of departure-time choice (with  $\mu \rightarrow 0$  corresponding to a deterministic choice and  $\mu \rightarrow \infty$  corresponding to a pure random choice) and  $[t^0, t^1]$  represents the period of possible departure times.

The departure time of an agent  $n$  can be drawn from the probability distribution defined by  $p_n^d(t)$  by relying on inverse sampling, which works as follows. Let  $F_n^d$  be the cumulative distribution function of the chosen departure time, i.e.,

$$F_n^d(t) = \int_{t^0}^t p_n^d(\tau) d\tau,$$

and let  $u_n \sim \mathcal{U}_{[0,1]}$  (uniform random number between 0 and 1). The chosen departure time is then simulated as  $(F_n^d)^{-1}(u_n)$ . Figure 2.3 illustrates how the departure time of an individual is drawn from the distribution  $F_n^d$  using inverse sampling. Note that, even though the utility function is represented as a piecewise linear function, time itself is not discretized and thus the departure time is a continuous variable.

The logsum formula of the departure-time choice is defined as

$$\text{logsum} = \mu \cdot \ln \int_{t^0}^{t^1} e^{V(\tau)/\mu} d\tau. \quad (2.3)$$




Figure 2.3: Illustration of the use of inverse sampling to draw a departure time from a cumulative distribution function

It represents the expected maximum utility that the agent can get from the trip, which can be used as a measure of the surplus of the agents.

We denote by

$$f^{\text{demand}} : T \mapsto f^{\text{demand}}(T) = \{(F_n^{\text{d}})^{-1}(u_n)\}_{1 \leq n \leq N}, \quad (2.4)$$

the function that returns the departure times chosen by all agents, given the expected travel-time function  $T$ .

**Supply model** The supply model consists in the road infrastructure which constrains the arrival times of the agents at destination. In the model investigated here, the road infrastructure consists in a single road with a free-flow travel time  $t^f$  and a road bottleneck, as depicted in Figure 2.1.

In METROPOLIS2, road bottlenecks can be in two states: open or closed. When a vehicle reaches an open bottleneck, it can cross it instantaneously and then the bottleneck is closed for a period of  $1/s$  time units, where  $s$  is the bottleneck's capacity. When a vehicle

reaches a closed bottleneck, the vehicle is pushed to the back of a point queue.<sup>4</sup> When the bottleneck opens again, the vehicle at the front of the queue can cross instantaneously and the bottleneck closes again.

The bottleneck thus defined satisfies the following two properties. First, two cars cannot cross the bottleneck in less than  $1/s$  time units. This implies that no more than  $s$  cars can cross the bottleneck in 1 time unit, which is consistent with the definition of the bottleneck capacity  $s$ . Second, the first-in-first-out property is satisfied as the next car to exit the bottleneck queue is the car that entered the queue the earliest.

In METROPOLIS2, the trips are simulated using an event-based model, which executes all the events occurring during the simulated period in chronological order. The events to be executed include, e.g., “the car of agent  $n$  is reaching the bottleneck at time  $t$ ” or “the bottleneck is re-opening at time  $t'$ ”. This event-based model will yield the arrival times of all the agents, denote by  $\mathbf{t}^a = \{t_n^a\}_{1 \leq n \leq N}$ , given their departure times, denoted by  $\mathbf{t}^d = \{t_n^d\}_{1 \leq n \leq N}$ .

In METROPOLIS2, the travel-time function,  $T$ , is derived from the simulated departure times,  $\mathbf{t}^d$ , and arrival times,  $\mathbf{t}^a$ . This function is represented as a sequence of  $M$  breakpoints, denoted by  $\{(x_1, y_1), \dots, (x_m, y_m), \dots, (x_M, y_M)\}$ . The  $x$ -values, corresponding to departure times, are defined as  $\{t^0, t^0 + \delta, t^0 + 2\delta, \dots, t^0 + (M - 1)\delta\}$ , where  $\delta$  denotes the length of the breakpoint intervals. The parameter  $M$  is defined as  $\lfloor (t^1 - t^0)/\delta \rfloor + 1$  to ensure coverage of the entire simulated period  $[t^0, t^1]$ . The  $y$ -values, corresponding to travel times, are computed as weighted averages of the simulated travel times within the corresponding departure-time interval. Formally, for  $1 \leq m \leq M$ , the value  $y_m$  is computed as

$$y_m = \sum_n w_n(x_m) \cdot (t_n^a - t_n^d),$$

---

<sup>4</sup>Horizontal queues with queue propagation can also be simulated in METROPOLIS2, see Section 2.4.

where the weights are given by

$$w_n(x_m) = \max\left(0, 1 - \frac{|x_m - t_n^d|}{\delta}\right).$$

Here, the weight of agent  $n$  decreases linearly with the distance between the departure-time breakpoint  $x_m$  and the agent's departure time. The weight is equal to 1 when  $t_n^d = x_m$  and it becomes zero when  $t_n^d$  is outside the interval  $[x_m - \delta, x_m + \delta]$ . Given the sequence of breakpoints  $(x_m, y_m)$  representing the travel-time function, METROPOLIS2 then relies on linear interpolation to compute the travel-time for any departure time  $t \in [t^0, t^1]$ . Figure 2.4 provides an illustration of how the travel-time function is computed from the departure times and travel times of the agents.

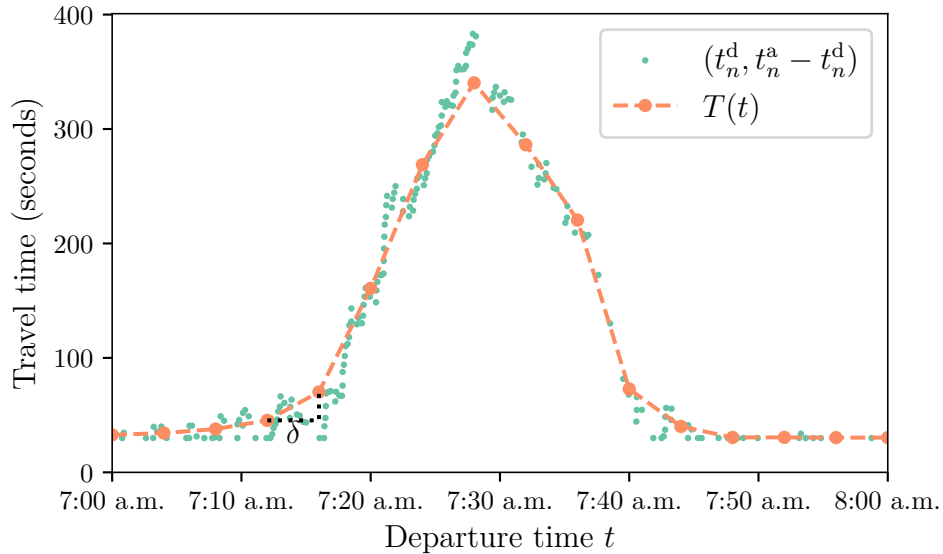


Figure 2.4: Agents' travel times and interpolated travel-time function (with  $N = 200$  and  $\delta = 4$  minutes)

We denote by

$$f^{\text{supply}} : \mathbf{t}^d \mapsto f^{\text{supply}}(\mathbf{t}^d) \quad (2.5)$$

the function that returns the travel-time function resulting from the departure times chosen by the agents.

**Equilibrium** The simulator METROPOLIS2 reaches an equilibrium of the model when it finds a set of departure times,  $\bar{\mathbf{t}}^d$ , which result, in the supply model, in the same travel-time function that was used, in the demand model, to simulate these departure times, i.e.,

$$f^{\text{demand}}(f^{\text{supply}}(\bar{\mathbf{t}}^d)) = \bar{\mathbf{t}}^d, \quad (2.6)$$

where  $f^{\text{demand}}$  and  $f^{\text{supply}}$  are the functions defined in equations (2.4) and (2.5), respectively.

We denote by  $\bar{T}$  the equilibrium travel-time function associated to  $\bar{\mathbf{t}}^d$ , i.e.,

$$f^{\text{supply}}(f^{\text{demand}}(\bar{T})) = \bar{T}. \quad (2.7)$$

The goal of METROPOLIS2 is therefore to find a set of departure times which solves the fixed-point problem of equation (2.6), or, equivalently, to find a travel-time function which solves the fixed-point problem of equation (2.7).

An iterative procedure to find such an equilibrium is depicted in Algorithm 1. The algorithm starts from an initial value for the travel-time function,  $T_0$  (e.g., the no-congestion travel-time function). Each iteration consists in simulating the departure times chosen by the agents, given the current value of the travel-time function (demand model, line 4), then simulating the travel-time function resulting from the current value of the departure times (supply model, line 5). The iterative process stops when a fixed point is reached (line 6).

---

**Algorithm 1:** Pseudo-code of a naive algorithm to find an equilibrium of the bottleneck model

---

**Input:**  $T_0$

```

1  $\kappa \leftarrow 0$ ;
2 repeat
3    $\kappa \leftarrow \kappa + 1$ ;                                /* Update the iteration counter */
4    $\mathbf{t}_\kappa^d = f^{\text{demand}}(T_{\kappa-1})$ ;                /* Run the demand model */
5    $T_\kappa = f^{\text{supply}}(\mathbf{t}_\kappa^d)$ ;                    /* Run the supply model */
6 until  $\mathbf{t}_\kappa^d \neq \mathbf{t}_{\kappa-1}^d$  and  $T_\kappa \neq T_{\kappa-1}$ ;
Output:  $\mathbf{t}_\kappa^d, T_\kappa$ 

```

---

Algorithm 1 is “naive” in the sense that, at each iteration, the travel-time function used

in the demand model is equal to the travel-time function simulated during the previous iteration. However, this naive algorithm will usually never converge, i.e., it will never reach a point where the departure times and travel-time function do not vary from one iteration to another (see Figure 2.19 in Appendix 2.A for a numerical illustration).

Instead, METROPOLIS2 introduces a learning model so that the expected travel-time function used to simulate the departure times is a weighted average of the simulated travel-time functions from the previous iterations. This approach is closely related to the Frank-Wolfe algorithm or the method of successive averages, both commonly applied in traffic assignment models (Patriksson, 2015).

Algorithm 2 describes the procedure used in METROPOLIS2, with the changes compared to the naive algorithm highlighted in blue.

---

**Algorithm 2:** Pseudo-code of the METROPOLIS2 algorithm, adapted to the bottleneck model

---

**Input:**  $\hat{T}_1$

```

1  $\kappa \leftarrow 0$ ;
2 repeat
3    $\kappa \leftarrow \kappa + 1$ ;                                /* Update the iteration counter */
4    $\mathbf{t}_\kappa^d = f^{\text{demand}}(\hat{T}_\kappa)$ ;                    /* Run the demand model */
5    $T_\kappa = f^{\text{supply}}(\mathbf{t}_\kappa^d)$ ;                      /* Run the supply model */
6    $\hat{T}_{\kappa+1} = f^{\text{learning}}(\kappa, T_\kappa, \hat{T}_\kappa)$ ;    /* Run the learning model */
7 until  $\mathbf{t}_\kappa^d \neq \mathbf{t}_{\kappa-1}^d$  and  $T_\kappa \neq \hat{T}_\kappa$ ;
Output:  $\mathbf{t}_{\kappa}^d, T_\kappa, \hat{T}_\kappa$ 

```

---

Algorithm 2 introduces a new variable for each iteration  $\kappa$ : the expected travel-time function,  $\hat{T}_\kappa$ , used in the demand model instead of the previously simulated travel-time function,  $T_{\kappa-1}$  (line 4). In the added learning model (line 6), the next expected travel-time function,  $\hat{T}_{\kappa+1}$ , is computed based on the simulated travel-time function,  $T_\kappa$ , and the current expected travel-time function,  $\hat{T}_\kappa$ . The stopping criteria (line 7) are updated compared to the naive algorithm so that the current simulated travel-time function,  $T_\kappa$ , is compared to the current expected travel-time function,  $\hat{T}_\kappa$ .<sup>5</sup>

---

<sup>5</sup>This change is required to prevent the algorithm from converging to a non-equilibrium. This would be

**Learning model** The learning model can be interpreted as a representation of how agents update their expectations regarding the travel time they anticipate based on the observed travel times in the past, although, it does not aim to be behaviorally accurate in modeling how agents form their anticipations. It is assumed that both the simulated and expected travel-time functions,  $T_\kappa$  and  $\hat{T}_\kappa$  respectively, are common knowledge among all agents. The selection of the learning model, specifically the choice of the function  $f^{\text{learning}}$ , plays a crucial role in the convergence of the algorithm toward an equilibrium. In the numerical applications, we consider the following learning function corresponding to the exponential smoothing method:

$$f^{\text{learning}}(\kappa, T_\kappa, \hat{T}_\kappa) = \frac{\lambda}{a_{\kappa+1}} T_\kappa + (1 - \lambda) \frac{a_\kappa}{a_{\kappa+1}} \hat{T}_\kappa, \quad (2.8)$$

where  $\lambda \in (0, 1]$  is the smoothing factor,  $\kappa$  is the iteration counter and  $a_\kappa = 1 - (1 - \lambda)^\kappa$  is a normalization term to correct for the weight of the initial value  $\hat{T}_1$ , such that, by recurrence,<sup>6</sup>

$$\hat{T}_{\kappa+1} = \frac{1}{a_\kappa} \lambda \sum_{i=0}^{\kappa} (1 - \lambda)^i T_\kappa,$$

with the convention that  $T_0 = \hat{T}_1$ .

With  $\lambda = 1$ , the learning function is  $f^{\text{learning}}(\kappa, T_\kappa, \hat{T}_\kappa) = T_\kappa$  and we are back to the naive algorithm (Algorithm 1). With  $\lambda \rightarrow 0$ , the learning function goes to  $f^{\text{learning}}(\kappa, T_\kappa, \hat{T}_\kappa) = 1/(\kappa + 1) \sum_{i=0}^{\kappa} T_\kappa$  and the expected value is simply an arithmetic mean of the past simulated values.

**Numerical applications** In Appendix 2.A, we present the results of various simulations of METROPOLIS2 for the bottleneck model considered in this section. The results show that,

---

the case whenever  $f^{\text{learning}}(\kappa, T_\kappa, \hat{T}_\kappa) \equiv \hat{T}_\kappa$ , i.e., the expected travel-time function is never updated, which would imply the chosen departure times and thus the simulated travel-time function to be the same from one iteration to another, despite equations (2.6) and (2.7) not being necessarily satisfied.

<sup>6</sup>In the standard exponential smoothing method, the weight of the initial value,  $\hat{T}_1$ , can get abnormally high, compared to the other values,  $\hat{T}_\kappa$ ,  $\kappa > 1$ , for a small  $\lambda$  value.

when the smoothing factor  $\lambda$  is not too large, the METROPOLIS2 algorithm is “almost” converging to an equilibrium. The simulations replicate the analytical results from de Palma et al. (1983) very accurately, suggesting that the discretization of the model does not impact its results.

We also perform many sensitivity analyses. We find that the simulations converge even with a small number of agents and a small number of breakpoints in the travel-time function. We find that when  $\mu$  is too small (the model is close to the deterministic model of Arnott et al., 1990), the simulation does not converge to an equilibrium.

## 2.4 Simulator Overview

The previous section described the METROPOLIS2 simulator on the simple single-road bottleneck model. In this section, we present an overview of METROPOLIS2 in the general case, which can be obtained by extending the bottleneck model with an arbitrary road network, heterogeneous agents, various modes, different vehicle types, etc.

**Algorithm** Algorithm 3 defines a pseudo-code of the process followed by the METROPOLIS2 simulator. It is a generalization of Algorithm 2 used for the bottleneck model (Section 2.3). The simulator starts with some input data  $\mathbf{D}$  and initial expected network conditions  $\hat{\mathbf{T}}_1$  (by default equal to free-flow conditions). The *input data*  $\mathbf{D}$  encompasses a population (agents with their preferences and assigned trips), a road network (roads and their attributes), vehicle types and technical parameters. Further details on the input data of METROPOLIS2 are provided in Appendix 2.B. The term *network conditions* denotes a data structure representing the road-level time-dependent travel times, which are common to all agents sharing the same vehicle type. See Appendix 2.D for more details on these network conditions.

At each iteration, the simulator executes three models: the demand model, the supply model and the learning model. Termination of the simulator occurs upon meeting one of



---

**Algorithm 3:** Pseudo-code of the METROPOLIS2 algorithm
 

---

**Input:**  $D, \hat{T}_1$   
**1**  $\kappa \leftarrow 0$ ;  
**2 repeat**  
**3**      $\kappa \leftarrow \kappa + 1$ ;                                     /\* Update the iteration counter \*/  
**4**      $\mathbf{z}_\kappa = f^{\text{demand}}(\hat{T}_\kappa; D)$ ;                             /\* Run the demand model \*/  
**5**      $\mathbf{T}_\kappa = f^{\text{supply}}(\mathbf{z}_\kappa; D)$ ;                             /\* Run the supply model \*/  
**6**      $\hat{T}_{\kappa+1} = f^{\text{learning}}(\kappa, \mathbf{T}_\kappa, \hat{T}_\kappa; D)$ ;             /\* Run the learning model \*/  
**7 until** *One of the stopping criteria is met*;  
**Output:**  $\mathbf{z}_\kappa, \mathbf{T}_\kappa, \hat{T}_\kappa$

---

the stopping criteria, which are discussed at the end of this section. Subsequently, the equilibrium travel decisions, along with the expected and simulated network condition, are returned. *Travel decisions*, denoted by  $\mathbf{z} \equiv \{z_n\}_{1 \leq n \leq N}$ , represent the mode, departure time and route chosen by agents for each of their trips. For further details regarding the output data of METROPOLIS2, refer to Appendix 2.C. Table 2.5 summarizes the primary notations used in this section.

Table 2.5: Main notations used in METROPOLIS2

Name	Notation
Input data (population, network, parameters)	$D$
Iteration counter	$\kappa$
Simulated network conditions at iteration $\kappa$ (vector of road-level travel-time functions)	$\mathbf{T}_\kappa$
Expected network conditions for iteration $\kappa$ (vector of road-level travel-time functions)	$\hat{T}_\kappa$
Travel decisions of agent $n$ at iteration $\kappa$	$z_{n,\kappa}$
Travel decisions at iteration $\kappa$ (vector of agent-level travel decisions)	$\mathbf{z}_\kappa$

Figure 2.5 illustrates the fundamental operational flow of METROPOLIS2. We now provide a brief description of the three models. More details on the demand and supply models are provided in Appendix 2.E.

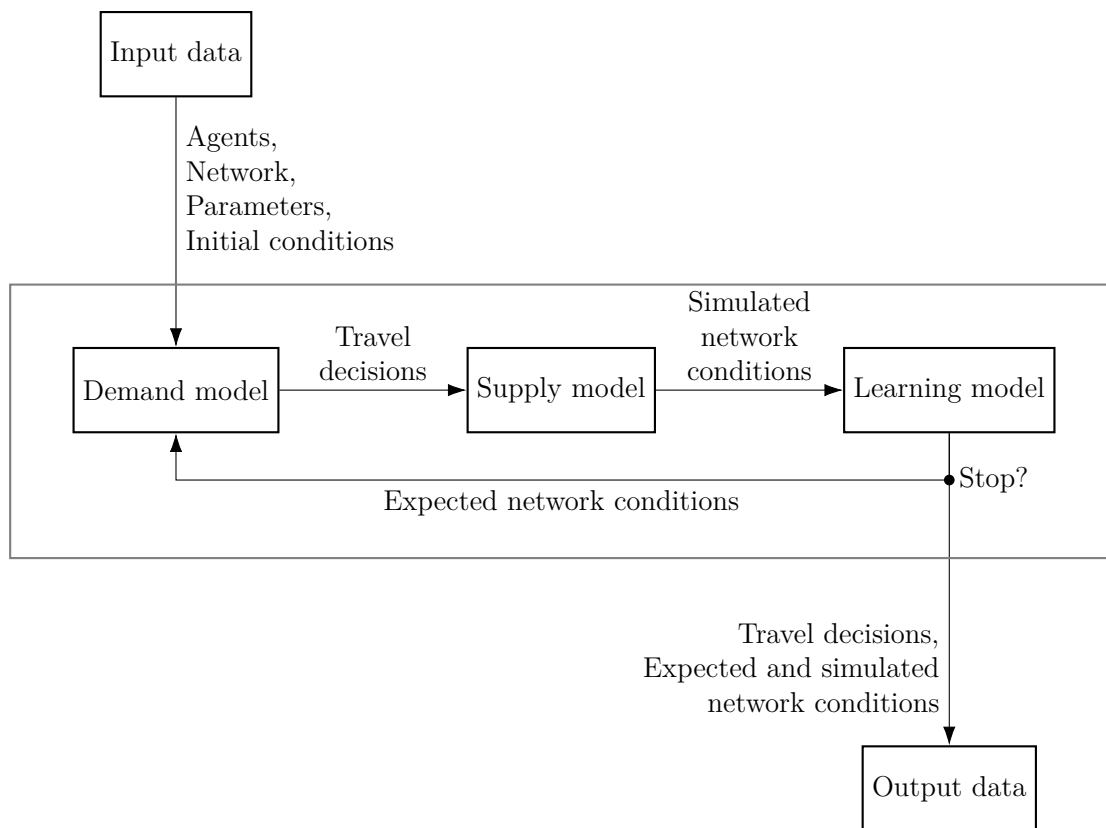


Figure 2.5: Overview of METROPOLIS2 process

**Demand model** The demand model simulates the travel decisions made by all agents, given the expected network conditions for the current iteration. Initially, agents choose between various *travel alternatives*. In its simplest form, a travel alternative represents a single trip with a given origin, destination and mode. However, it can also encompass chains of trips (e.g., a car trip from home to grocery store, followed by 15 minutes of shopping, followed by another car trip from grocery store to workplace). In most cases, the choice of a travel alternative is akin to a mode choice, although it can also represent a destination choice. Following the selection of a travel alternative, agents make decisions regarding their departure time and the route, if applicable. The agents' choices (travel alternative, departure time and route) are referred to as the travel decisions of the agents. The choices are taken based on utility maximization and the expected network conditions are used to predict the trips' travel time. In the general case, utility is expressed as a polynomial function of travel

time, with a linear penalty for early or late arrivals at destinations and for early or late departures from origins (see Appendix 2.E.1).

METROPOLIS2 follows a specific choice order, coherent with the decision-making process of agents in real world scenarios: agents select their travel alternative, then their departure time, and finally, their route. By relying on backward induction, METROPOLIS2 ensures that the optimal travel alternative is chosen knowing that the optimal departure time and route for that alternative will subsequently be determined. The chosen route is the fastest route at the time of departure. Departure-time choice can be represented by a Multinomial Logit model or a Continuous Logit model. Mode choice can be represented by a Multinomial Logit model or by a deterministic choice.<sup>7</sup> The output of the demand model are the travel decisions that the agents are planning to do during the simulated period.

In METROPOLIS2, a routing algorithm determines the route chosen by each agent based on their departure time from origin and the expected network conditions. The algorithm selects the fastest available route for the entire network. The underlying assumption mirrors real-life behavior, where individuals rely on car navigation system to plan their itinerary. This approach differs from many other simulators, which often constrain agents to choose their route from a limited choice set.

When using discrete-choice models, the expected maximum utility of the travel decision is computed by METROPOLIS2 using the logsum formula (See Appendix 2.E.1 for a formal definition). This value reflects the range of choices available to the agent in the complete choice hierarchy. It can be interpreted as the agent's surplus for cost-benefit analysis.

The demand model can be represented mathematically by a function  $f^{\text{demand}}$ , taking as an argument the expected network conditions  $\hat{\mathbf{T}}$  and the input data  $\mathbf{D}$ , and returning the utility-maximizing travel decisions of all agents  $\mathbf{z} = \{z_n\}_{1 \leq n \leq N}$ , i.e.,

$$f^{\text{demand}} : (\hat{\mathbf{T}}, \mathbf{D}) \mapsto f^{\text{demand}}(\hat{\mathbf{T}}; \mathbf{D}) = \mathbf{z}.$$

---

<sup>7</sup>Any discrete-choice models can be used for departure-time choice and mode choice, such as Probit model, as long as the random perturbations can be simulated.

**Supply model** The supply model simulates the movements of vehicles on the road network.<sup>8</sup> These movements are a direct outcome of the travel decisions made by agents in the demand model. The simulation operates in continuous time using an event-based model, where events can represent a vehicle departing from its origin or a vehicle reaching the end of its current link, for example. These events are executed in chronological order.

Congestion is simulated using bottleneck queues and speed-density functions. Bottleneck queues limit the inflow and outflow of vehicles at the link-level, based on the link’s capacity. The simulation of bottleneck queues is described in details in Section 2.3.3. With speed-density function, the speed at which a vehicle will travel on a link is determined by the density of vehicles currently on that link when the vehicle enters. The density on a link is computed as the sum of the headway length of vehicles currently on the link, divided by the total length of the link (product of its length and its number of lanes). Additionally, when queue propagation is enabled, vehicles are constrained from entering a link unless enough space is available for their headway length. Following a vehicle’s departure from a link, the liberated space it creates propagates backward through the link at a predetermined speed until it reaches the link’s origin. Only when this free space has reached the beginning of the link can entering vehicles utilize it.<sup>9</sup> At the end of the supply model, the simulated travel-time functions are computed at the link-level, based on the process described in Section 2.3.3 and Figure 2.4. These travel-time functions form the *simulated network conditions*.

The supply model can be represented mathematically by a function  $f^{\text{supply}}$ , taking as an argument the travel decisions of all agents  $\mathbf{z}$  and the input data  $\mathbf{D}$ , and returning the

---

<sup>8</sup>In the current version of METROPOLIS2, there exists a one-to-one correspondence between agents and vehicles, allowing these terms to be used interchangeably. However, it is worth noting that in future versions, this relationship might no longer hold, with the inclusion of features such as ride-sharing or public transit, where multiple agents might share a single vehicle.

<sup>9</sup>When queue propagation is enabled, gridlocks can appear. A gridlock is a situation in which traffic comes to a complete standstill because no vehicle can move until the vehicle in front of it moves. To address these gridlocks, METROPOLIS2 employs a mechanism where a vehicle is forced to move to the next link if it remains in a pending state for a specific duration (configurable). This approach is similar to the one taken in MATSim, while, in METROPOLIS1, the route choice at each intersection prevent gridlocks from occurring.

simulated network conditions  $\mathbf{T}$ , i.e.,

$$f^{\text{supply}} : (\mathbf{z}; \mathbf{D}) \mapsto f^{\text{supply}}(\mathbf{z}; \mathbf{D}) = \mathbf{T}.$$

**Learning model** The learning model computes the expected network conditions to be used in the next iteration by combining the current expected network conditions with the simulated network conditions obtained from the supply model. These expected network conditions are then used in the demand model’s routing algorithm to predict travel times. The learning model mirrors how car navigation systems leverage historical data from all users to forecast travel times. Different methods can be employed for the learning model, such as the exponential smoothing method, presented in Section 2.3.3.

The learning model can be represented mathematically by a function  $f^{\text{learning}}$ , taking as arguments the iteration counter  $\kappa$ , the simulated network conditions for the current iteration  $\mathbf{T}_\kappa$ , the expected network conditions for the current iteration  $\hat{\mathbf{T}}_\kappa$  and the input data  $\mathbf{D}$ , and returning the expected network conditions for the next iteration  $\hat{\mathbf{T}}_{\kappa+1}$ , i.e.,

$$f^{\text{learning}} : (\kappa, \mathbf{T}_\kappa, \hat{\mathbf{T}}_\kappa; \mathbf{D}) \mapsto f^{\text{learning}}(\kappa, \mathbf{T}_\kappa, \hat{\mathbf{T}}_\kappa; \mathbf{D}) = \hat{\mathbf{T}}_{\kappa+1}.$$

**Nash equilibrium** The goal of METROPOLIS2 is to find a Nash equilibrium of the simulation, defined by the input data  $\mathbf{D}$ . Here, a Nash equilibrium is defined as a state where no agent can improve their utility by unilaterally changing their travel decisions, considering the travel decisions made by the other agents. In the context of transport simulations, such a Nash equilibrium is also called an *agent-based User Equilibrium* (Nagel and Flötteröd, 2016). More formally, a Nash equilibrium can be defined as the following fixed-point problem. We want to find equilibrium travel decisions  $\bar{\mathbf{z}}$  such that

$$f^{\text{demand}}(f^{\text{supply}}(\bar{\mathbf{z}}; \mathbf{D}); \mathbf{D}) = \bar{\mathbf{z}}.$$

or, alternatively, we want to find equilibrium network conditions  $\bar{T}$  such that

$$f^{\text{supply}}(f^{\text{demand}}(\bar{T}; \mathbf{D}); \mathbf{D}) = \bar{T},$$

Intuitively, these equations mean that the expected network conditions used by the agents when taking their decisions are identical to the simulated network conditions resulting from these agents' decisions. This ensures that all agents choose the travel decisions maximizing their utility and that they thus cannot improve their utility by unilaterally changing these decisions. The definition of an equilibrium with the bottleneck model, in Section 2.3, equations (2.6) and (2.7), is a special case with departure time as the only decision variable and with a single travel-time function (corresponding to the single road).

**Properties of the Nash equilibrium** In dynamic traffic assignment models, there is no proof of the existence, uniqueness and stability of a Dynamic User Equilibrium in the general case (see Iryo 2013). A Nash equilibrium in METROPOLIS2 is a generalization of a standard Dynamic User Equilibrium (with the addition of mode choice, departure-time choice, trip chaining, etc.) so the existence, uniqueness and stability of a Nash equilibrium in METROPOLIS2 cannot be proven either. Algorithm 3 is a heuristic method to approximate a Nash equilibrium. Large-scale analytical applications (see Section 2.6) reveal that the algorithm does not perfectly converge to a Nash equilibrium. Therefore, it becomes essential to assess the quality of the solution by measuring its distance to a Nash equilibrium. One approach to evaluate the solution's quality is by calculating the magnitude of change in travel decisions or network conditions relative to the previous iteration. Another option is to compute the distance between the simulated and expected network conditions. See Section 2.6 for example of indicators computing distance to a Nash equilibrium.

These indicators of a distance to a Nash equilibrium can be used as stopping criteria to ensure that the simulator only stops when the quality of the solution is good enough. Alternatively, one can fix the number of iterations to run and check a posteriori the quality

of the solution.

## 2.5 Implementation

METROPOLIS2 is implemented as a Rust program. Rust was chosen for its excellent performance and safety guarantees. These safety guarantees ensure that code parallelization can be utilized with minimal risk, enhancing the efficiency and reliability of the program.

To enhance its performance, METROPOLIS2 employs time-dependent contraction hierarchies, a state-of-the-art routing algorithm (Batz et al., 2013). The simulator utilizes two types of routing queries: (i) profile queries, which determine the fastest travel time for any departure time, given an origin-destination pair (without returning corresponding routes), and (ii) earliest-arrival queries, which provide the earliest arrival time along with the corresponding route, for a given origin-destination pair and departure time from the origin. In the demand model, profile queries efficiently retrieve the travel-time function for all trips, required to compute the utility function used for departure-time choice. Earliest-arrival queries are used to retrieve routes once agents have selected their mode and departure time. The adoption of contraction hierarchies involves a preprocessing phase executed during each iteration of METROPOLIS2, enabling the acceleration of subsequent queries. Moreover, code parallelization is integrated into the routing algorithm to speed up the processes, such as retrieving routes for all agents. Additionally, code parallelization is also employed in the demand model, where agents' decisions can be computed independently.

The choice of CSV and Parquet formats for input and output files in METROPOLIS2 offers a combination of usability, flexibility, and performance. CSV offers simplicity and compatibility with a wide range of tools and platforms. Parquet, on the other hand, is a columnar storage format that is optimized for performance and efficiency, making it suitable for handling large datasets efficiently. This combination of formats satisfies different use cases and preferences, ensuring that users can interact with METROPOLIS2 data in a way

that best suits their needs.

Given the complexity of the simulator, issues or unexpected behaviors are likely to appear. To mitigate this risk, we have implemented a rigorous testing framework. The tests evaluate individual components (unit tests) as well as the interactions between various modules (integration tests). The example applications below (Section 2.6) also serve as validation tests for the correctness of the source code.

## 2.6 Comparison with METROPOLIS1

As discussed in Section 2.2, METROPOLIS2 replicates and improves most of the methodology used in METROPOLIS1. In this section, we compare the results of the two simulators when running the same simulation. We rely on the work of Saifuzzaman et al. (2012) which propose a calibrated simulation of METROPOLIS1 for Paris' urban area.<sup>10</sup>

**Demand input** The demand side of the simulation is composed of 477 067 agents, all performing a single trip between 4:00 a.m. and 1:00 p.m. There is no mode choice, car is the only available mode. The agents are divided in five categories. In each category, the agents have the same parameters  $\alpha$ ,  $\beta$ ,  $\gamma$ ,  $\Delta$  and  $\mu$  and the desired arrival times  $t^*$  are independent and identically distributed with a Normal or Uniform distribution. The studied area is divided in 1326 zones and there is one origin-destination matrix for each category.

**Supply input** The road network of the simulation is limited to the main roads of the area, which represent 17 987 nodes and 39 395 links. In addition, there are 4462 connectors (virtual links used to connect the zones' starting and ending nodes to the road network).

In METROPOLIS1, there is no explicit bottlenecks. Instead, the simulator uses the

---

<sup>10</sup>Because we are missing some data, we were not able to replicate exactly their results. However, we did run the two simulators with the exact same input (same population and preference parameters, same road network).



following speed-density function, imitating a road bottleneck:<sup>11</sup>

$$\text{speed}(\text{density}) = \min \left( \text{free-flow speed}, \frac{\text{vehicle headway} \times \text{capacity}}{\text{density}} \right). \quad (2.9)$$

Figure 2.6 proposes a representation of this speed-density function. This speed-density func-




Figure 2.6: Bottleneck speed-density function on an example road

Road characteristics: 200 meters, free-flow speed of 50 km/h, 1 lane, capacity of 1000 vehicles per hour.

tion is used in both METROPOLIS1 and METROPOLIS2, using the same road capacities. Queue propagation is enabled, with a vehicle headway of 8 meters. With this configuration, the congestion model of METROPOLIS2 is an exact replication of the congestion model of METROPOLIS1, with one difference: in METROPOLIS1, link travel times are truncated,

<sup>11</sup>To understand equation (2.9), observe that, using

$$\text{density} = \frac{\text{vehicle headway} \times \text{vehicle count}}{\text{length} \times \text{lanes}},$$

the road travel time can be written as

$$\text{travel time}(\text{vehicle count}) = \frac{\text{length}}{\text{speed}(\text{density})} = \min \left( \frac{\text{length}}{\text{free-flow speed}}, \frac{\text{vehicle count}}{\text{capacity} \times \text{lanes}} \right).$$

Assume that there are  $n$  vehicles on a road with a bottleneck of capacity  $s$ . Because of the bottleneck, there must be at most  $s$  vehicles traveling through the road during one hour. This implies that the  $n$  vehicles on the road cannot all cross the bottleneck in more than  $n/s$  hours. Therefore, the road travel-time is upper bounded by  $n/s$ , which corresponds to the second term in the “min” operator of the previous equation.

while, in METROPOLIS2, time is continuous (there is no rounding).

**Technical parameters** With both METROPOLIS1 and METROPOLIS2, we run 200 iterations using a linear learning model (the expected travel time for the next iteration is the arithmetic mean of the simulated travel time of all the previous iterations). Following Saifuzzaman et al. (2012), the interval length between two breakpoints is set to 20 minutes.

**Output indicators** To measure the convergence of the simulations, we use four different indicators. The two first indicators are used to measure by how much the choices of the agents (departure time and route) change from one iteration to another. The change in departure time is measured by the root of the mean squared difference between the departure time chosen in the previous and current iteration, for each agent. This indicator is denoted by  $\text{RMSE}_\kappa^{\text{dep}}$ , for iteration  $\kappa > 1$ , and is defined as

$$\text{RMSE}_\kappa^{\text{dep}} = \sqrt{\frac{1}{N} \sum_n (t_{n,\kappa}^{\text{d}} - t_{n,\kappa-1}^{\text{d}})^2},$$

where  $t_{n,\kappa}^{\text{d}}$  is the departure time chosen by agent  $n$  at iteration  $\kappa$ . The indicator is measured in seconds, so a value  $\text{RMSE}_\kappa^{\text{dep}} = 5$ , for example, means that, on average, the agents shifted their departure time by plus or minus 5 seconds from iteration  $\kappa - 1$  to iteration  $\kappa$  (note that, because the differences are squared, the larger shifts contribute more to the RMSE than the smaller shifts). The second indicator measures how much the routes taken by the agents change from one iteration to another. Let's denote by  $\rho_{n,\kappa}$  the share of the length of the route taken by agent  $n$  during iteration  $\kappa$  that was not taken during iteration  $\kappa - 1$ , i.e., a value  $\rho_{n,\kappa} = 5\%$ , for example, means that 5% of the length traveled by agent  $n$  during iteration  $\kappa$  was made on roads that were not taken during iteration  $\kappa - 1$ . The indicator  $\text{RMSE}_\kappa^{\text{route}}$  is defined as

$$\text{RMSE}_\kappa^{\text{route}} = \sqrt{\frac{1}{N} \sum_n (\rho_{n,\kappa})^2}.$$

The third indicator measures how the simulated travel times (i.e., the simulated network

conditions) differ from the expected travel times (i.e., the expected network conditions). To do so, we use the root of the mean squared difference between the simulated and expected travel time for iteration  $\kappa$ , for each link of the road network and for each departure time. This indicator is denoted by  $\text{RMSE}_{\kappa}^T$  and is defined as

$$\text{RMSE}_{\kappa}^T = \sqrt{\frac{1}{R} \sum_r \frac{1}{t^1 - t^0} \int_{t^0}^{t^1} [T_{r,\kappa}(t) - \hat{T}_{r,\kappa-1}(t)]^2 dt},$$

where  $T_{r,\kappa}(t)$  (resp.  $\hat{T}_{r,\kappa}(t)$ ) is the simulated (resp. expected) travel time on link  $r$  at time  $t$  during iteration  $\kappa$ ,  $R$  is the number of links and  $[t^0, t^1] = [4 \text{ a.m.}, 1 \text{ p.m.}]$  is the simulated period.

The fourth indicator measures how well the agents predict the travel time that they will face. It is based on the travel-time expectation error: the difference between the actual and expected travel time. The indicator is denoted by  $\text{RMSE}_{\kappa}^{\text{expect}}$  and is defined as

$$\text{RMSE}_{\kappa}^{\text{expect}} = \sqrt{\frac{1}{N} \sum_n (\text{tt}_{n,\kappa} - \hat{\text{tt}}_{n,\kappa})^2},$$

where  $\text{tt}_{n,\kappa}$  is the travel time of agent  $n$  for iteration  $\kappa$  and  $\hat{\text{tt}}_{n,\kappa}$  is the expected travel time of agent  $n$  for iteration  $\kappa$ .

**Output comparison** A comparison of the results of the simulations with METROPOLIS1 and METROPOLIS2 is provided in Table 2.6.<sup>12</sup> The technical results show that METROPOLIS2 runs slightly faster than METROPOLIS1 when using a single thread. However, as METROPOLIS2 can make use of parallel computing, the running time is around 9 times faster when using 16 threads compared to using a single thread (or around 10 times faster compared to the single-threaded METROPOLIS1). METROPOLIS2 peak memory use is almost 30 % larger than METROPOLIS1's. One reason for this difference is that METRO-

<sup>12</sup>All simulations are run on a machine with a 8-core Intel Xeon CPU (3.20GHz) and with 128GiB of RAM. The source code to run the simulations and generate the graphs is available at <https://github.com/LucasJavaudin/Metropolis1-2Comparison/>. The simulations are run with version 1.0.0 of METROPOLIS2.

POLIS1 uses single-precision floating-point format (32 bits), while METROPOLIS2 uses double-precision (64 bits).

Table 2.6: Comparison of the aggregate results for a Paris simulation with METROPOLIS1 and METROPOLIS2

	METROPOLIS1	METROPOLIS2
<i>Technical output</i>		
Running time (1 thread)	18h 29m 12s	16h 12m 48s
Running time (16 threads)	18h 29m 12s	1h 49m 17s
Peak memory use	2.74 GiB	3.52 GiB
<i>Convergence output</i>		
$\text{RMSE}_{200}^{\text{dep}}$	2m 5s	5s
$\text{RMSE}_{200}^{\text{route}}$	18.07 %	5.49 %
$\text{RMSE}_{200}^T$	1m 40s	46s
$\text{RMSE}_{200}^{\text{expect}}$	10m 28s	6m 05s
<i>Traffic-related output</i>		
Average utility	-5.58 €	-5.32 €
Average departure time	09:10:44	09:11:58
Average arrival time	09:27:07	09:27:31
Average travel time	16m 23s	15m 33s
Average free-flow travel time	13m 24s	13m 1s
Average route length	15.51 km	15.05 km
Average number of links taken	37.28	35.33

The output related to the convergence of the simulations show that METROPOLIS2 achieves a solution that is more stable, in all aspects. This reveals that METROPOLIS2's solution is closer to a Nash equilibrium than METROPOLIS1' solution.<sup>13</sup>

Figure 2.7 compares the convergence of the  $\text{RMSE}_k^{\text{dep}}$  metric across iterations, between METROPOLIS1 and METROPOLIS2. Initially, both simulators exhibit similar convergence patterns for the first 15 to 20 iterations. Then, METROPOLIS2 continues to steadily decrease, while METROPOLIS1 appears to reach a plateau.

<sup>13</sup>Observe that the  $\text{RMSE}_{200}^{\text{expect}}$  is still quite large in both METROPOLIS1 and METROPOLIS2. This is due to threshold effects that can appear with the speed-density function, where the road travel time can jump from a few seconds to a few minutes with just one more car being on the road segment. These threshold effects are less prominent when using a larger population size (the simulation presented uses 10 % of total population), a smaller breakpoint interval and explicit bottleneck queues instead of speed-density functions.




Figure 2.7: Comparison of the convergence of departure times over iterations, between METROPOLIS1 and METROPOLIS2

Figure 2.8 compares the convergence of the  $\text{RMSE}_{\kappa}^{\text{expect}}$  metric across iterations, between METROPOLIS1 and METROPOLIS2. Initially, METROPOLIS2 shows higher levels of  $\text{RMSE}_{\kappa}^{\text{expect}}$ . As iterations progress, METROPOLIS2 achieves lower levels of the metric compared to METROPOLIS1. On the other hand, METROPOLIS1 reaches a plateau after about 50 iterations but exhibits less variation of the metric from one iteration to another. This difference in behavior can be attributed to the adaptability of the agents in METROPOLIS1: they can adjust their routes during the trip when encountering congestion, thereby avoiding extreme travel times; while, in METROPOLIS2, the route is chosen before leaving origin (in the demand model) and cannot be changed afterward. However, this adaptability results in more frequent route changes between iterations, leading to less predictable travel times overall.

The traffic-related output on Table 2.6 shows that METROPOLIS2 reaches a situation where agents are slightly better-off compared to the situation returned by METROPOLIS1 (utility is larger by 26 cents on average, travel time is smaller by 50 seconds on average, etc.). As we mentioned previously, the congestion is modeled in the same way in the two




Figure 2.8: Comparison of the convergence of travel-time expectation error over iterations, between METROPOLIS1 and METROPOLIS2

simulations, which means that these differences are explained by the choices of the agents only.

Figure 2.9 shows that the departure-time distributions are very similar for the two simulations. The peak around 8:00 a.m. is slightly more pronounced in METROPOLIS2. Figure 2.10 shows that the travel-time distributions are also very similar. Figure 2.11 shows that the link-level flows are about the same between the two simulations: the correlation coefficient between link-level flows in METROPOLIS1 and METROPOLIS2 is 95.37%.

In conclusion, the results from METROPOLIS1 and METROPOLIS2 are similar but the latter runs faster and provides a closer approximation to equilibrium compared to METROPOLIS1. Further improvements in convergence could be achieved in METROPOLIS2 by using explicit bottleneck queues, defined in Section 2.3.3.<sup>14</sup> However, the use of bottleneck

<sup>14</sup>The use of speed-density functions can lead to a simulation that exhibits less stability compared to simulations using explicit bottleneck queues. This instability comes from a discontinuity in the measure of density. Specifically, when a vehicle enters a road at approximately the same time another vehicle exits the same road, the entering vehicle's travel time can experience a rapid jump from small to large value, conditional on whether the exiting vehicle left just before or just after the entry. This discontinuity is more prominent on shorter roads with smaller capacities. Conversely, with explicit bottleneck queues, there is no such discontinuities. The waiting time for a vehicle in the queue decreases linearly with the time that elapsed since the preceding vehicle exited, eventually reaching zero.

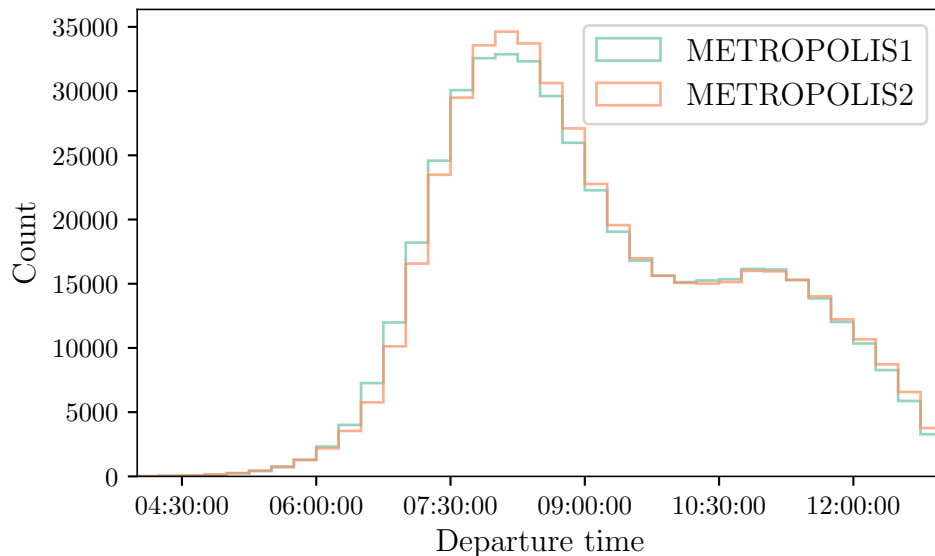


Figure 2.9: Comparison of the departure-time distribution in the Paris simulation with METROPOLIS1 and METROPOLIS2




Figure 2.10: Comparison of the travel-time distribution in the Paris simulation with METROPOLIS1 and METROPOLIS2




Figure 2.11: Comparison of the link-level flows for a Paris simulation with METROPOLIS1 and METROPOLIS2

Note: Each dot represents a link of the road network, with the observed flows in the METROPOLIS1 simulation on the x-axis and the observed flows in the METROPOLIS2 simulation on the y-axis.

queues leads to a distinct congestion model,<sup>15</sup> complicating comparison between convergence with speed-density functions and bottleneck queues.

## 2.7 Conclusion

This chapter has detailed the methodology and implementation details of the METROPOLIS2 simulator. Beginning with the foundational single-road bottleneck model, we provided insights into METROPOLIS2’s functionality and demonstrated its ability to replicate key results from the analytical model. We then proposed a comprehensive overview of ME-

<sup>15</sup>Consider a road with a free-flow travel time of 10 seconds and a bottleneck capacity of 300 vehicles per hour (i.e., a bottleneck closing time of 12 seconds). If there is a vehicle entering the road every 11 seconds, the congestion differs significantly between the bottleneck speed-density function and the bottleneck queue. When using the bottleneck speed-density function, equation (2.9), there is no congestion. This is because the density remains null when a new vehicle enters the road, as the preceding vehicle has already left. However, when using bottleneck queues, congestion does appear. Despite the first vehicle exiting the road before the second one enters, the second vehicle is forced to wait until the bottleneck closing time elapses. Consequently, a queue forms with an increasing waiting time as subsequent vehicles are arriving at a faster rate than the bottleneck’s output rate.



TROPOLIS2's capabilities in general scenarios, discussed some implementation details and conducted a comparative analysis with the METROPOLIS1 simulator using a Paris case study.

METROPOLIS2 exhibits promising potential for efficiently and accurately simulating traffic equilibrium in large-scale scenarios, making it well-suited for cost-benefit analysis. However, the discussion on generating a synthetic population and calibrating the simulator for real-world scenarios, though crucial, remains unexplored in this chapter, leaving avenues for future research.

Furthermore, future work may explore additional features for the simulator. For example, integrating a public-transit model could enhance METROPOLIS2's capabilities, enabling consideration of in-vehicle congestion in route choices for public-transit trips.

# Appendix

## 2.A Bottleneck Model: Numerical Application

This appendix presents the results of simulations of METROPOLIS2 for the bottleneck model presented in Section 2.3. It contains some definitions (Section 2.A.1), a comparison to the analytical results from de Palma et al. (1983) (Section 2.A.2) and some sensitivity analysis (Section 2.A.3).

### 2.A.1 Definitions and Notations

Table 2.7 shows the default values considered for the various parameters of the bottleneck model investigated. We consider a population of 100 000 agents who can choose a departure time during the time period 7:00 a.m. to 8:00 a.m. The desired arrival time of the agents at destination is  $t^* = 7:30$  a.m. The travel time from origin to destination when the bottleneck is not congested is  $t^f = 30$  seconds. The bottleneck capacity is set to  $s = 1.5 \times N$  cars per hour, which means that all agents can cross the bottleneck in 40 minutes. In the simulations, the interval between two breakpoints for the piecewise-linear travel-time functions is set to  $\delta = 1$  minute. The initial value for travel-time function,  $\hat{T}_1$ , is set to the no-congestion travel-time function (a constant function equal to the free-flow travel-time  $t^f$ ).

**Measuring convergence of the simulation** When running the simulations, the stopping criteria, as defined in Algorithm 2, line 7, are never satisfied. Instead, we run the simulations for a fixed number of iterations and check how far the results are from an equilibrium using the two indicators presented below.

First, we use the root of the mean squared difference between the departure time chosen in the previous and current iteration, for each agent. This indicator is denoted by  $\text{RMSE}_\kappa^{\text{dep}}$ , for iteration  $\kappa > 1$ , and is defined as

$$\text{RMSE}_\kappa^{\text{dep}} = \sqrt{\frac{1}{N} \sum_n (t_{n,\kappa}^{\text{d}} - t_{n,\kappa-1}^{\text{d}})^2},$$

Table 2.7: Default values for the parameters of the model

Parameter	Notation	Value
Number of agents	$N$	100 000
Time period	$[t^0, t^1]$	[7:00 a.m., 8:00 a.m.]
Desired arrival time at destination	$t^*$	7:30 a.m.
Marginal disutility of travel time	$\alpha$	10 \$/h
Marginal disutility of early arrivals	$\beta$	5 \$/h
Marginal disutility of late arrivals	$\gamma$	7 \$/h
Scale of the utility's random component	$\mu$	1
Free-flow origin-to-destination travel time	$t^f$	30 seconds
Bottleneck flow rate	$s$	$1.5 \times N$ cars / h
Travel-time function breakpoints interval	$\delta$	1 minute

where  $t_{n,\kappa}^d$  is the departure time chosen by agent  $n$  at iteration  $\kappa$ . Second, we use the root of the mean squared difference between the expected travel-time function and the simulated travel-time function, for each departure time. This indicator is denoted by  $\text{RMSE}_\kappa^T$ , for iteration  $\kappa > 1$ , and is defined as

$$\text{RMSE}_\kappa^T = \sqrt{\frac{1}{t^1 - t^0} \int_{t^0}^{t^1} [T_\kappa(t) - \hat{T}_\kappa(t)]^2 dt}.$$

For an iteration  $\kappa$ ,  $\text{RMSE}_\kappa^{\text{dep}} \approx 0$  implies that the agents did not change much their departure time compared to the previous iteration and  $\text{RMSE}_\kappa^T \approx 0$  implies that the travel-time function was correctly anticipated. Therefore, these indicators can be interpreted as a distance to an equilibrium.

**Measuring distance to the analytical results** The analytical results of the bottleneck model we consider are demonstrated in de Palma et al. (1983). From their findings, we compute the equilibrium rate of departures from origin,  $\bar{r}^d(t)$  (see Figure 2.12), and the equilibrium travel-time function,  $\bar{T}(t)$  (see Figure 2.13).

To compare the simulation results with the analytical results, we use the following distance:<sup>16</sup>

$$D = \max_t |R^d(t) - \bar{R}^d(t)|,$$

with

$$\bar{R}^d(t) = \frac{1}{N} \int_{t^0}^t \bar{r}^d(t) dt$$

<sup>16</sup>Observe that this measure corresponds to the Kolmogorov-Smirnov statistic for comparing a sample of values to a probability distribution, which justify using it here.




Figure 2.12: Analytical result for the rate of departures from origin

Note:  $t_q$  is the time at which congestion starts to appear,  $\tilde{t}$  is the departure time such that the individual arrives on time,  $t_{q'}$  is the time at which congestion ends.




Figure 2.13: Analytical result for the travel-time function

Note:  $t_q$ ,  $\tilde{t}$ ,  $t_{q'}$  are defined as in Figure 2.12.

the cumulative rate of departures from the analytical results and

$$R^d(t) = \frac{1}{N} \sum_n \mathbf{1}_{t_n^d \leq t}$$

the cumulative rate of departures from the simulation.

**Learning model** For the learning model, we use the exponential smoothing defined by equation 2.8. Further research is required to determine the optimal value of the smoothing factor  $\lambda$  to use. We give here some guidance for the choice of  $\lambda$  resulting from our experiments. With smaller values of  $\lambda$ , a smaller weight is put on the simulated travel-time function for the current iteration so the model learns slower but the convergence is more stable (in the sense that it is less sensitive to what happened during a single iteration). Conversely, with larger values of  $\lambda$ , convergence is usually faster but less stable.

In the simulations presented below, we use a value of  $\lambda = 0.4$ , unless specified otherwise. The impact of  $\lambda$  on the convergence of the simulations is further discussed in Section 2.A.3, using numerical results.

## 2.A.2 Comparison to the Analytical Results

We simulate 200 iterations of METROPOLIS2 using the parameters defined in Table 2.7.<sup>17</sup> Some results for the simulation are presented in Table 2.8. First, observe that, the simulation is essentially reaching an equilibrium as, at the last iteration, both  $\text{RMSE}_\kappa^{\text{dep}}$  and  $\text{RMSE}_\kappa^T$  are practically zero ( $3 \times 10^{-12}$  and  $2 \times 10^{-12}$  seconds respectively). This means that, from one iteration to another, the chosen departure times stay almost identical and that the simulated travel-time function in the supply model is almost identical to the expected travel-time function in the demand model. One reason the simulation is not converging to exactly zero might be related to the approximation errors of floating-point arithmetic. The maximum distance between the cumulative rate of departures in the analytical results and in the simulation results is  $D = 0.19\%$ , which indicates that the simulation replicates quite well the analytical results. We also find that the expected maximum utility, computed using the logsum formula (2.3), is at 7.188, which is close to the equilibrium value of 7.276 from the analytical results.

Figures 2.14 and 2.15 show a comparison of the rate of departures from origin and the travel-time function between the simulation and the analytical results. For the rate of

---

<sup>17</sup>All simulations are run on a machine with a 8-core Intel Xeon CPU (3.20GHz) and with 128GiB of RAM. The source code to run the simulations and generate the graphs is available at <https://github.com/LucasJavaudin/MetropolisBottleneck/>. The simulations are run with version 1.0.0 of METROPOLIS2.

Table 2.8: Results of running METROPOLIS2 to simulate the bottleneck model

Variable	Value
Running time	3m 24s
Expected max. utility	7.188
Average travel time	1m 57s
Distance to analytical results $D$	0.19 %
$\text{RMSE}_{200}^{\text{dep}}$ (seconds)	$3 \times 10^{-12}$
$\text{RMSE}_{200}^T$ (seconds)	$2 \times 10^{-12}$

departures (Figure 2.14), we observe some noise in the curve corresponding to the simulation, due to the discretization of the agents, but the two curves are always close to each other. For the travel-time function (Figure 2.15), there is less noise and the curves follows each other very closely. The maximum travel time is about 3 seconds larger in the simulation results than in the analytical results.

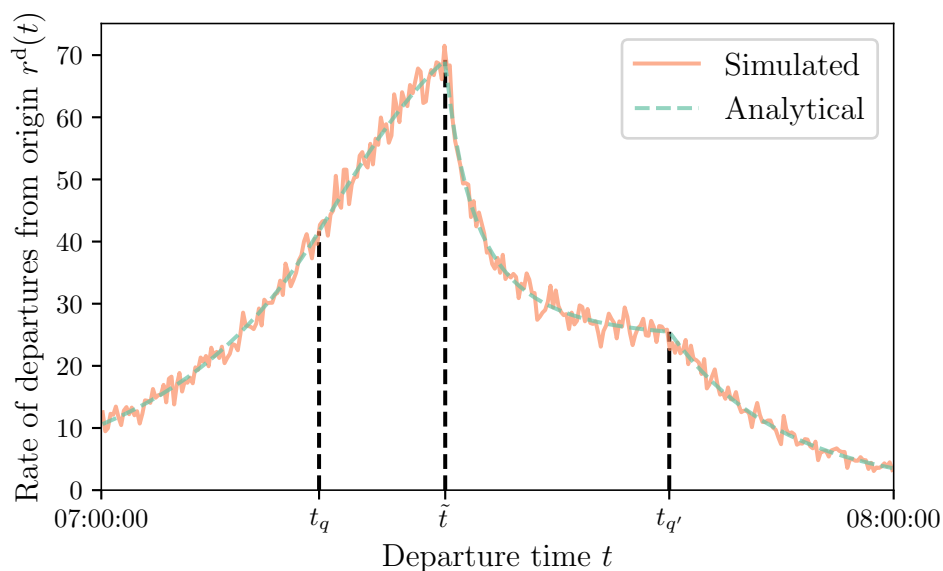


Figure 2.14: Comparison of the rate of departures from origin in the simulation and in the analytical results

Note:  $t_q$ ,  $\tilde{t}$ ,  $t_{q'}$  are defined as in Figure 2.12.

Figure 2.16 shows the convergence of  $\text{RMSE}_{\kappa}^{\text{dep}}$  and  $\text{RMSE}_{\kappa}^T$  from iteration to iteration. Both indicators steadily decrease for the first 120 iterations. Then, the indicators stay stable for the remaining iterations.




Figure 2.15: Comparison of the travel-time function from origin to destination in the simulation and in the analytical results

Note:  $t_q$ ,  $\tilde{t}$ ,  $t_{q'}$  are defined as in Figure 2.12.




Figure 2.16: Convergence of  $\text{RMSE}_{\kappa}^{\text{dep}}$  and  $\text{RMSE}_{\kappa}^T$

**Running with different random seeds** The only randomness in the simulation is due to the draw of the uniform random numbers  $u_n$  to simulate the departure time of each agent using inverse sampling. We run 10 simulations with different random seeds to investigate how the values of the draws affect the results. Table 2.9 shows how the results vary from one simulation to another. The results are roughly similar for all 10 simulations. For example, the distance to analytical results  $D$  varies between 0.19 % and 0.38 %.

Table 2.9: Average, minimum and maximum values for the results with 10 different random seeds

Variable	Average	Minimum	Maximum
Running time	3m 33s	3m 30s	3m 42s
Expected max. utility	7.190	7.187	7.194
Average travel time	1m 56s	1m 55s	1m 57s
Dist. analytical res. $D$	0.26 %	0.19 %	0.38 %
RMSE <sub>200</sub> <sup>dep</sup> (seconds)	$3 \times 10^{-12}$	$2 \times 10^{-12}$	$3 \times 10^{-12}$
RMSE <sub>200</sub> <sup>T</sup> (seconds)	$1 \times 10^{-12}$	$8 \times 10^{-13}$	$2 \times 10^{-12}$

When averaging the departure rates for the 10 simulations, we find a distance to the analytical results of  $D = 0.10$  %, which is almost twice smaller than the smallest distance for any individual simulation. The reason is that averaging over multiple simulations is akin to simulating more individuals, which reduces the noise created by drawing random values.<sup>18</sup> Figure 2.17 shows a comparison of the analytical departure rate and the departure rate resulting from the average of the 10 simulations. Compared to Figure 2.14, we observe much less noise for the simulated departure rate.

**Using systematic sampling** We now try to use systematic sampling (Madow and Madow, 1944) to draw the uniform random numbers  $u_n$ . Concretely, the  $u_n$  values are set to be evenly spaced between 0 and 1, with an interval of  $1/N$  between two values. Figure 2.18 compares the simulated and analytical equilibrium rate of departures, when running 200 iterations of METROPOLIS2. With systematic sampling, we reach a smaller distance to the analytical results of  $D = 0.06$  %. The reason is that, by using evenly spaced values of  $u_n$ , we reduce greatly the noise induced by drawing random values.

<sup>18</sup>As shown in Section 2.A.3, the distance to the analytical results gets smaller as the number of agents simulated increases.






Figure 2.17: Comparison of the rate of departures when averaging over 10 simulations  
 Note:  $t_q$ ,  $\tilde{t}$ ,  $t_{q'}$  are defined as in Figure 2.12.




Figure 2.18: Comparison of the rate of departures when using systematic sampling  
 Note:  $t_q$ ,  $\tilde{t}$ ,  $t_{q'}$  are defined as in Figure 2.12.

### 2.A.3 Sensitivity Analysis

**Impact of the smoothing factor** The smoothing factor  $\lambda$  is used in the exponential smoothing, equation 2.8. It represents the weight of the simulated travel-time function of the current iteration over the weight of the expected travel-time function of the previous iteration. We run a simulation for 1000 iterations with 6 different values of  $\lambda$ : 0.05, 0.2, 0.4, 0.6, 0.8 and 1. Table 2.10 shows the main results of the simulations for the different values of  $\lambda$ . Figure 2.19 shows the convergence of  $\text{RMSE}_\kappa^{\text{dep}}$  for the different values of  $\lambda$  (we do not show here the convergence of  $\text{RMSE}_\kappa^T$  as it is roughly the same). For a value of  $\lambda = 1$ , there is no convergence:  $\text{RMSE}_\kappa^{\text{dep}}$  stays at high values. For the other values of  $\lambda$ , the simulation always converges to a very low value of  $\text{RMSE}_\kappa^{\text{dep}}$  and  $\text{RMSE}_\kappa^T$  and the distance to the analytical results is the same. The convergence is the fastest with  $\lambda = 0.4$ , which justify using this value in the previous section. We do not observe any major difference with regards to the running time of the simulation between the different values of  $\lambda$ .

Table 2.10: Sensitivity of the results to the smoothing factor

$\lambda$	0.05	0.2	0.4	0.6	0.8	1
Running time	17m 30s	17m 35s	17m 38s	17m 35s	17m 33s	18m
Expected max. surplus	7.188	7.188	7.188	7.188	7.188	6.921
Average travel time	1m 57s	1m 57s	1m 57s	1m 57s	1m 57s	2m 13s
Dist. analytical res. $D$	0.19 %	0.19 %	0.19 %	0.19 %	0.19 %	13.67 %
$\text{RMSE}_{200}^{\text{dep}}$	$2 \times 10^{-12}$	$4 \times 10^{-12}$	$3 \times 10^{-12}$	$3 \times 10^{-12}$	$7 \times 10^{-12}$	$2 \times 10^2$
$\text{RMSE}_{200}^T$	$2 \times 10^{-12}$	$1 \times 10^{-12}$	$3 \times 10^{-12}$	$2 \times 10^{-12}$	$4 \times 10^{-12}$	$2 \times 10^2$

**Impact of the number of agents** We run a simulation for 200 iterations with 3 different values of the number of agents  $N$ : 1000, 10 000 and 100 000.<sup>19</sup> Note that, since the bottleneck capacity  $s$  is proportional to  $N$ , the equilibrium travel-time function from the analytical results is not impacted by a variation in  $N$  and the equilibrium rate of departures is proportional to  $N$ . The results are shown in Table 2.11. First, we observe that the running time of the simulation is roughly linear in the number of agents. Then, we observe that the simulation results are closer to the analytical results as the number of agents increases, as shown by the distance to the analytical results  $D$ . However, the convergence of  $\text{RMSE}_\kappa^{\text{dep}}$  is roughly the same for all values of  $N$ , as shown on Figure 2.20. We conclude that, for smaller values of  $N$ , the discretization of the agents has a larger impact which explains why the comparison to the analytical results gets worse, but this discretization does not prevent the simulations from converging.

<sup>19</sup>We also run simulations with fewer agents but these simulations dependent heavily on the draws of the random numbers so we decided not to report their results.



Figure 2.19: Impact of the smoothing factor on  $\text{RMSE}_{\kappa}^{\text{dep}}$  convergence

Table 2.11: Sensitivity of the results to the number of agents

$N$	1000	10 000	100 000
Running time	1.70s	16.53s	3m 28s
Expected max. utility	7.147	7.176	7.188
Average travel time	2m 13s	2m 1s	1m 57s
Dist. analytical res. $D$	3.37 %	0.83 %	0.19 %
$\text{RMSE}_{200}^{\text{dep}}$	$2 \times 10^{-12}$	$3 \times 10^{-12}$	$3 \times 10^{-12}$
$\text{RMSE}_{200}^T$	$8 \times 10^{-13}$	$3 \times 10^{-12}$	$3 \times 10^{-12}$

**Impact of the travel-time function breakpoints interval** The parameter  $\delta$  represents the time interval between two breakpoints in the travel-time function. We run a simulation for 200 iterations with 4 different values of  $\delta$ : 1 min, 5 min, 10 min and 15 min. Note that, as  $\delta \rightarrow 0$ , the travel-time function converges to a continuous function. The results are shown in Table 2.12 and Figure 2.21. We observe that the running time slightly decreases as  $\delta$  increases. The simulations converge to an equilibrium for all values of  $\delta$ , convergence is faster for larger values of  $\delta$ . The convergence of the simulation is satisfactory even with  $\delta = 15$  min, i.e., with only 5 breakpoints in total (1 each 15 minutes). Still, the distance to the analytical results  $D$  is smaller for smaller values of  $\delta$  because agents can better predict the travel time that they will face for any departure time.

Figure 2.22 shows the simulated travel-time function at the last iteration for the different values of  $\delta$ . With larger values of  $\delta$ , the simulated travel-time function cannot match well the

Figure 2.20: Impact of the number of agents on  $\text{RMSE}_{\kappa}^{\text{dep}}$  convergenceTable 2.12: Sensitivity of the results to the breakpoint interval  $\delta$ 

$\delta$	1 min	5 min	10 min	15 min
Running time	3m 24s	3m 34s	3m 16s	3m 11s
Expected max. utility	7.188	7.181	7.160	7.176
Average travel time	1m 57s	2m 2s	2m 14s	2m 16s
Dist. analytical res. $D$	0.19 %	1.14 %	2.39 %	3.08 %
$\text{RMSE}_{200}^{\text{dep}}$	$3 \times 10^{-12}$	$3 \times 10^{-12}$	$2 \times 10^{-12}$	$2 \times 10^{-12}$
$\text{RMSE}_{200}^T$	$2 \times 10^{-12}$	$2 \times 10^{-12}$	$4 \times 10^{-13}$	$3 \times 10^{-13}$

analytical results because there is not enough breakpoints. Note that using a larger value for  $\delta$  is akin to assuming that agents have limited information on the travel-time function. This might give more realistic results than the analytical model which assumes perfect information on the travel times.

**Impact of the scale of the utility’s random component** The parameter  $\mu$  represents the scale of the utility’s random component. Figure 2.23 shows how the probability distribution of the departure-time choice varies with  $\mu$ . As  $\mu \rightarrow 0$ , the probability distribution converges to a distribution with a probability 1 to choose the best departure time (i.e.,  $t^* - t^f$  under free-flow conditions) and the model converges to the deterministic bottleneck model from Arnott et al. (1990). As  $\mu \rightarrow +\infty$ , the probability distribution converges to a uniform distribution.




Figure 2.21: Impact of the breakpoint interval on  $\text{RMSE}_{\kappa}^{\text{dep}}$  convergence




Figure 2.22: Travel-time function for different values of the breakpoint interval

Note:  $t_q, \tilde{t}, t_{q'}$  are defined as in Figure 2.12.




Figure 2.23: Probability distribution of the departure-time choice for different values of the scale of the utility’s random component (assuming a constant travel time)

We run 1000 iterations of METROPOLIS2 with 5 different values of  $\mu$ .<sup>20</sup> The main results are shown in Table 2.13 and Figure 2.24. From the figure, observe that as  $\mu$  decreases, convergence is slower. With  $\mu = 0.2$ , the simulation does not converge ( $\text{RMSE}_{1000}^{\text{dep}} = 30$ ,  $\text{RMSE}_{1000}^T = 100$ ) and the distance to the analytical results is large ( $D = 20.10\%$ ). For small values of  $\mu$ , the stochastic model that we simulate is close to the deterministic one so the non-convergence is likely related to the challenges identified in Section 2.3.2 which prevent simulations from solving the deterministic bottleneck model. For the simulations which converge ( $\mu \geq 0.5$ ), the distance to the analytical results is similar (from 0.15% to 0.34%). The running times are also very similar for all values of  $\mu$ .

Table 2.13: Sensitivity of the results to the scale of the utility’s random component

$\mu$	0.2	0.5	1	2	5
Running time	17m 40s	17m 48s	17m 32s	17m 14s	17m 1s
Average travel time	6m 7s	3m 30s	1m 57s	44s	30s
Dist. analytical res. $D$	20.10%	0.15%	0.19%	0.30%	0.34%
$\text{RMSE}_{1000}^{\text{dep}}$ (seconds)	$3 \times 10^1$	$3 \times 10^{-11}$	$3 \times 10^{-12}$	$2 \times 10^{-12}$	$3 \times 10^{-12}$
$\text{RMSE}_{1000}^T$ (seconds)	$1 \times 10^2$	$3 \times 10^{-11}$	$1 \times 10^{-12}$	$2 \times 10^{-12}$	$1 \times 10^{-13}$

<sup>20</sup>For each value of  $\mu$ , we choose the value of the smoothing factor  $\lambda$  which gives the best convergence. The values used are:  $\lambda = 0.01$  for  $\mu = 0.2$ ,  $\lambda = 0.1$  for  $\mu = 0.5$ ,  $\lambda = 0.4$  for  $\mu = 1.0$ ,  $\lambda = 0.8$  for  $\mu = 2.0$  and  $\lambda = 1.0$  for  $\mu = 5.0$ . Observe that for smaller values of  $\mu$ , the convergence is better with a smaller value of  $\lambda$ .




Figure 2.24: Impact of the scale of the utility’s random component on  $\text{RMSE}_{\kappa}^{\text{dep}}$  convergence

#### 2.A.4 Considering Heterogeneity

So far, we have considered a population of homogeneous agents. The only difference between agents, and the reason their departure times differed, was the different values of  $u_n$ , which we interpret as unobserved heterogeneity. In this section, we include some observed heterogeneity to the model by assuming that the  $t^*$  values are distributed. More specifically, we assume that the desired arrival times of the agents,  $\{t_n^*\}_{1 \leq n \leq N}$ , are independent and identically distributed with

$$t_n^* \sim \mathcal{N}(\bar{t}^*, \sigma_{t^*}),$$

where  $\bar{t}^* = 7:30$  a.m. is the center of the desired-arrival-time distribution and  $\sigma_{t^*} = 5$  min is the variance of the distribution. The distribution of the  $t_n^*$  values drawn is presented on Figure 2.25.

We run two different simulations, with  $\mu = 1$  and  $\mu = 0.6$ . The main results are shown in Table 2.14. Observe that the running time of METROPOLIS2 is similar between the three scenarios. This is because agents are always simulated as separate entities so there is no speed gains when simulating homogeneous agents.

In the heterogeneous case, the desired arrival times are spread over the simulated period so the distribution of departure times is more spread than in the homogeneous case, as shown on Figure 2.26. This explains why, with  $\mu = 1$ , the average travel time is smaller in the heterogeneous case (1m 8s) than in the homogeneous case (1m 57s). In the heterogeneous




Figure 2.25: Distribution of the  $t_n^*$  values for the simulations with observed heterogeneity

case, when  $\mu$  decreases from 1 to 0.6, the average travel time increases from 1m 8s to 2m 1s (consistently with the results from the sensitivity analysis to  $\mu$ , in Section 2.A.3). Then the homogeneous simulation with  $\mu = 1$  and the heterogeneous simulation with  $\mu = 0.6$  result in a similar travel time. Finally, in the heterogeneous simulations,  $\text{RMSE}_{200}^{\text{dep}}$  and  $\text{RMSE}_{200}^T$  are larger than in the homogeneous simulation but the values are still practically zero.

Table 2.14: Comparison between the homogeneous and heterogeneous bottleneck models

Model Parameters	Homogeneous $\mu = 1, \sigma_{t^*} = 0$	Heterogeneous A $\mu = 1, \sigma_{t^*} = 5 \text{ min}$	Heterogeneous B $\mu = 0.6, \sigma_{t^*} = 5 \text{ min}$
Running time	3m 24s	3m 25s	3m 27s
Average travel time	1m 57s	1m 8s	2m 1s
$\text{RMSE}_{200}^{\text{dep}}$ (seconds)	$3 \times 10^{-12}$	$7 \times 10^{-6}$	$6 \times 10^{-4}$
$\text{RMSE}_{200}^T$ (seconds)	$2 \times 10^{-12}$	$2 \times 10^{-5}$	$4 \times 10^{-4}$






Figure 2.26: Comparison of the rate of departures between homogeneous and heterogeneous simulations

## 2.B Input Data

The input data of METROPOLIS2 is composed of the following categories:

- The *road network*, which defines the infrastructure through which the road trips are performed (see Section 2.B.1).
- The *agents*, which are the entities whose travel behaviors are simulated (see Section 2.B.2).
- The *parameters*, which govern various technical aspects of the simulation, such as stopping criteria (see Section 2.B.3).
- The *initial network conditions*, which are the expected conditions used in the first iteration of the simulator (see Section 2.D). When omitted the free-flow conditions are used.

### 2.B.1 Road-Network Input

In METROPOLIS2, a road network is defined by a set of edges and a set of vehicle types.

Each edge is characterized by its fundamental attributes, including its source node, target node, base speed, length, number of lanes, bottleneck capacity and speed-density function parameters.

Each vehicle type is defined by the following characteristics:

- headway: typical distance between two vehicles, measured from tip-to-tip;

- passenger car equivalent: a measure of the congestion impact of the vehicle type, in comparison to a standard passenger car;
- speed restrictions: speed limitations that apply to this vehicle type.

Additionally, it is possible to specify road restrictions, i.e., edges which are forbidden to some vehicle types.

## 2.B.2 Agents

METROPOLIS2 can simulate an arbitrary number of agents. Each agent chooses between various *travel alternatives*, where a travel alternative can represent a single trip (with given mode, origin and destination), a chain of trips, or no trip (e.g., work-at-home). Figure 2.27 represents the structure of the data for a single agent. An agent is defined by the travel alternatives from which they can choose and the choice model defining how the chosen travel alternative is determined (e.g., Multinomial Logit model, deterministic model).




Figure 2.27: Structure of an agent

A travel alternative can be either a constant utility amount (e.g., to represent the utility of work-at-home) or the description of a trip chain. The structure of trip chains is represented in Figure 2.28. A single trip can be represented as a special case of a trip chain.

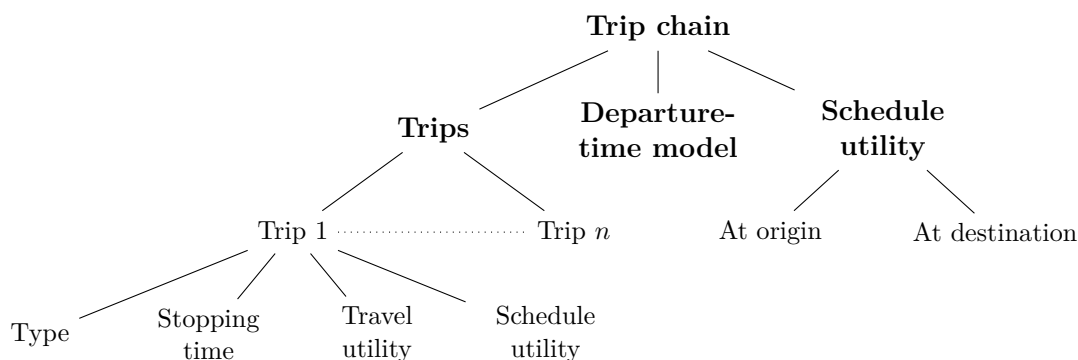


Figure 2.28: Structure of a trip chain

A trip chain is constructed from an arbitrary number of trips, each of which can belong to one of two distinct types, defined below.

- *Road trip*: a trip defined by an origin and destination node on the road network with a specific vehicle type.
- *Virtual trip*: a trip without any interaction with the network infrastructure (e.g., walk trip, which neither impacts nor is impacted by congestion), defined by either a constant travel time or a travel-time function.

These types can be combined within the same trip chain to represent inter-modal journeys.

Apart from the trip type and its associated data, each trip is characterized by the following elements.

- Stopping time: waiting time at the end of the trip before the start of the next trip, or, equivalently, duration of the activity performed at the trip's destination.
- Travel utility: utility as a function of the trip's travel time.
- Schedule utility: utility as a function of the arrival time.

The sequence followed by an agent  $n$  with a chain of multiple trips works as follows. Let  $t_{n,k}^\Delta$  denote the stopping time at the end of the  $k$ -th trip of agent  $n$ . The agent's departure time,  $t_n^d$ , coincides with the start of their first trip,  $t_{n,1}^d$ . Subsequently, the agent reaches the destination of their first trip upon the completion of the trip's travel time,  $tt_{n,1}$ , resulting in their arrival time at the destination of the first trip as  $t_{n,1}^a \equiv t_{n,1}^d + tt_{n,1}$ . The second trip starts after the end of the first trip's stopping time,  $t_{n,1}^\Delta$ , occurring at time  $t_{n,2}^d \equiv t_{n,1}^a + t_{n,1}^\Delta$ . This alternation between trips and stopping times continues, until the agent arrives, upon the conclusion of the stopping time of their last trip, at time  $t_n^a$ . Figure 2.29 illustrate the timing dynamics of a two-trip chain.




Figure 2.29: Timing dynamics of a trip chain with two trips

Finally, a trip chain is characterized by a *departure-time model* indicating how the departure time from origin is computed (e.g., Continuous Logit, Multinomial Logit), and a *schedule utility* at origin and at destination (a function of  $t_n^d$  and  $t_n^a$ ).

### 2.B.3 Parameters

The parameters of METROPOLIS2 include:

- the time window within which the trips take place,
- the learning model used to update the expected network conditions at each iteration,
- stopping criteria defining when the simulation should stop.

## 2.C Output Data

Output data for METROPOLIS2 can be divided in three categories: agent-specific output (Section 2.C.1), output specific to the road network (Section 2.C.2) and aggregate results (Section 2.C.3).

### 2.C.1 Agent Output

For each simulated agent, the output of METROPOLIS2 includes

- the travel alternative chosen;
- the expected utility (pre-trip) and the actual utility (post-trip);
- the departure time, arrival time and travel time of all the trips taken;
- the route taken for all the trips (if applicable), with the detailed timings.

These values correspond to the choices and actions taken for the last simulated iteration.

### 2.C.2 Road-Network Output

METROPOLIS2 returns the current value of the simulated and expected network conditions at the end of the simulation. These network conditions contain, for each edge of the road network, the expected travel time as a function of time of day. The simulated network conditions represent the simulated travel time in the supply model, for the last iteration. The expected network conditions represent an average of the simulated travel time over the past iterations, as per the learning model of the simulation.

These values can be analyzed to identify the main congested edges of the network. They can also be used to compute the time-dependent travel time from any origin to any destination of the network.

### 2.C.3 Aggregated Results

The aggregated results of METROPOLIS2 include:

- aggregate results specific to the agents (e.g., mean departure time, arrival time, travel time, expected utility, actual utility);
- aggregate results specific to road trips (e.g., mean travel time, mean route length, mean uncongested travel time);
- aggregate results related to convergence (e.g., departure-time and route variation compared to the previous iteration, difference between actual and expected travel times).

These values are returned at the end of each iteration of the simulation.

## 2.D Network Conditions

Network conditions are a critical component of METROPOLIS2. They represent the perceptions of agents on the travel time that they will face during their trips. These network conditions are common knowledge, i.e., all agents have the same perceptions. The network conditions are updated at each iteration of the simulator based on the simulated travel times on the network in the supply model.

The road-network conditions are composed of a travel-time function for each edge of the road network. These travel-time functions represent the expected time required to traverse the edge for any time of the day. They are used in the routing algorithm during the demand model to compute the fastest route from origin to destination.

In METROPOLIS2, the travel-time functions are represented as piecewise linear functions so that they can simply be stored as a list of breakpoints. If the simulation contains multiple vehicle types, one travel-time function is saved for each vehicle type because these vehicle types can experience different speed limits.

In the supply model, the occupied length and bottleneck queue are recorded for each edge at any time. These values allow to compute the simulated travel time for any vehicle type and any edge, i.e., the simulated road-network conditions. Then, the road-network conditions are used in the learning model to update the expected conditions for the next iteration.

## 2.E Description of the Models

### 2.E.1 Demand Model

The demand model simulates the travel decisions of all the agents, given the expected network conditions (which are common knowledge). The travel decisions consist of choosing one of the available travel alternatives and, if applicable, a departure time or route. These choices are based on utility maximization. The network conditions are used to compute the expected travel time, for road trips.

Backward induction is used to select the travel alternative, departure time and route in this order. The demand model thus works as follows for each agent:

1. The selected route is computed for each travel alternative and each possible departure time.
2. Based on the routes chosen in Step 1, the departure time maximizing utility and the expected value of the departure-time choice are computed for each travel alternative.

3. A travel alternative is chosen based on its expected value, which is given by the departure-time choice from Step 2.

The remainder of this section is organized as follows. We provide an explanation of how utility is defined in METROPOLIS2. Then, the route, departure time and travel alternative choices are presented, in this order.

**Utility** The utility of trip  $k$  of agent  $n$  is defined by

$$V_{n,k}^{\text{trip}}(t_{n,k}^d, t_{n,k}^a) = V_{n,k}^{\text{so}}(t_{n,k}^d) + V_{n,k}^{\text{tt}}(t_{n,k}^a - t_{n,k}^d) + V_{n,k}^{\text{sd}}(t_{n,k}^a),$$

where  $t_{n,k}^d$  is the trip's departure time,  $t_{n,k}^a$  is the trip's arrival time,  $V_{n,k}^{\text{so}}$  is the schedule utility at origin,  $V_{n,k}^{\text{sd}}$  is the schedule utility at destination and  $V_{n,k}^{\text{tt}}$  is the travel-time utility. In the current version of METROPOLIS2,  $V_{n,k}^{\text{so}}$  is limited to a linear schedule-delay penalty:

$$V_{n,k}^{\text{so}}(t_{n,k}^d) = \beta_{n,k} \cdot [t_{n,k}^* - \Delta_{n,k} - t_{n,k}^d]_+ + \gamma_{n,k} \cdot [t_{n,k}^d - t_{n,k}^* - \Delta_{n,k}]_+,$$

with  $\beta_{n,k}$  (resp.  $\gamma_{n,k}$ ) the penalty for early (resp. late) departures,  $t_{n,k}^*$  the desired departure time,  $\Delta_{n,k}$  the length of the desired departure time period and  $[x]_+ = \max(x, 0)$ . The same specification can be used for  $V_{n,k}^{\text{sd}}$ , with (desired) arrival time instead of (desired) departure time. Any other parametric specifications could be implemented easily. The travel-time utility  $V_{n,k}^{\text{tt}}$  is limited to a polynomial function of travel time, with up to degree 4, including a constant. Again, any other parametric specification could be implemented easily. To evaluate numerically the utility function, METROPOLIS2 relies on linear approximation.

**Route choice** In the case of road trips,<sup>21</sup> given a travel alternative and a departure time, agents choose the route that minimizes travel time.<sup>22</sup> These fastest routes are computed using a time-dependent routing algorithm directly in the demand model (there is no en-route choice in METROPOLIS2). They are computed using the expected network conditions which give the time-dependent travel time for each edge of the road network. The relevant information for departure-time choice are the travel time of these fastest routes, for all possible departure times of each travel alternative.

<sup>21</sup>There is no route to choose in the case of virtual trips as travel time is exogenous.

<sup>22</sup>The assumption that agents always take the fastest route allows the use of very efficient routing algorithm. In some special cases, however, the fastest route is not actually the one which maximizes utility. For example, in the  $\alpha$ - $\beta$ - $\gamma$  model, agents can increase their utility by taking a longer route whenever  $\beta > \alpha$ , i.e., the value of travel-time savings is smaller than the penalty for being early at destination. In such rare cases, the travel decisions returned by the demand model do not strictly adhere to the utility-maximization principle.

**Departure-time choice** Departure-time choice is continuous for trip chains composed of road or virtual trips because utility is a continuous function of time. In the case of multiple trips, only the departure time for the first trip is chosen. The departure times for the subsequent trips are then derived based on the departure time of this first trip, along with the corresponding travel times and stopping times for each trip within the chain.

The way the departure time is selected depends on the departure-time choice model of the agent given as input. For example, when the departure-time choice model is a Continuous Logit model, the probability to choose departure time  $t$  is given by:

$$p_{n,j}^d(t) = \frac{e^{V_{n,j}(t)}}{\int_{t^0}^{t^1} e^{V_{n,j}(\tau)} d\tau},$$

where  $V_{n,j}(t)$  is the utility as a function of departure time for travel alternative  $j$  of agent  $n$  and  $t^0$  (resp.  $t^1$ ) is the earliest (resp. latest) possible departure time. Inverse transform sampling is used to draw the selected departure time from the probability distribution given above. Then, the expected maximum utility of the travel alternative is given by the logsum formula:

$$V_{n,j} = \ln \int_{t^0}^{t^1} e^{V_{n,j}(\tau)} d\tau.$$

Consider the case of an agent  $n$  with a chain of  $K$  trips,  $1, \dots, k, \dots, K$ . From route choice, METROPOLIS2 gets the expected travel-time function  $T_{n,k}$  from origin to destination for each trip  $k$  of agent  $n$ . Then, for any departure time  $t^d$  from the first origin, the departure times and arrival times of any trip  $k$  can be computed recursively from these expected travel-time functions  $\{T_{n,k}\}_{1 \leq k \leq K}$  and the trips' stopping times  $\{t_{n,k}^\Delta\}_{1 \leq k \leq K}$ . For example, the arrival time at destination of trip 1, is  $t_{n,1}^a \equiv t_n^d + T_{n,1}(t_n^d)$  the departure time from origin of trip 2 is  $t_{n,2}^d \equiv t_{n,1}^a + t_{n,1}^\Delta$ , and the arrival time at destination of trip 2 is  $t_{n,2}^a \equiv t_{n,2}^d + T_{n,2}(t_{n,2}^d)$ . Therefore, from these travel-time functions and the stopping times, it is possible to compute the utility of the trip chain for any departure time from the first origin.

**Choice of a travel alternative** The choice between the travel alternatives available to an agent is based on the *choice model* of the agent. The choice model specifies how the travel alternative is selected given the *expected maximum utility* of all the travel alternatives, determined by their departure-time choice. For example, with a deterministic choice model, the travel alternative with the largest expected utility is chosen.

With a Multinomial Logit choice model, the probability that agent  $n$  chooses travel

alternative  $j$  is given by the Logit formula:

$$p_{n,j} = \frac{e^{V_{n,j}}}{\sum_{j'} e^{V_{n,j'}}},$$

where  $V_{n,j}$  is the expected maximum utility of travel alternative  $j$  of agent  $n$ . Like with the departure-time choice model, the travel alternative can be drawn from such a probability distribution using inverse transform sampling. The expected maximum utility, or surplus, that the agent can get from their available travel alternatives is computed as the logsum formula:

$$V_n = \ln \sum_j e^{V_{n,j}}.$$

## 2.E.2 Supply Model

The supply model of METROPOLIS2 uses an event-based model to simulate the movements of agents and vehicles on the network, based on the travel decisions given by the demand model. An event is defined as an action taken by an agent or a vehicle (e.g., a car reaches an intersection, an agent reaches the destination of a virtual trip), together with the time at which the action happens (its *execution time*). The events are stored in a priority queue, based on their execution time, where the earliest events to be executed are on top of the queue. The supply models work as follows:

- When the supply model starts, the event queue is loaded with the initial event of the agents' trips and public-transit trips, which corresponds to the departure from origin.
- The supply model then proceeds in executing all events in chronological order.
- Events can have an impact on the network-infrastructure state when being executed. For example, when a car enters the next edge on its route, vehicle density is reduced on the previous edge and increased on the new one.
- Events can push new events to the priority queue while being executed. For example, when a car enters an edge, an event is created to be executed at the time the car will reach the end of the edge (where the travel time on the edge depends on the current density on the edge).
- The supply model stops when there is no more event to execute in the priority queue.

While the events are being executed, the simulated network conditions are recorded to be used in the learning model.



## Chapter 3

# Impact of Low Emission Zones on Spatial and Economic Inequalities using a Dynamic Transport Simulator

### Abstract

Low Emission Zones (LEZs) are widely implemented in European cities to improve air quality by restricting access for the most polluting vehicles. The effects of LEZs on air quality, travel behavior, road congestion, and inequalities are complex and challenging to predict. This chapter evaluates the LEZ planned for the Greater Paris area, set for implementation in January 2025.

We use METROPOLIS2, a dynamic agent-based transport simulator that estimates equilibrium travel decisions – mode, departure time and route choice – across large-scale road networks for millions of agents. A novel calibration methodology, integrating machine-learning techniques, is developed to adjust model parameters using real-world data, ensuring that the baseline simulation closely mirrors observed behaviors.

Our analysis estimates the LEZ’s impact on travel surplus and pollution exposure for individuals both inside and outside the zone. The findings suggest that while the LEZ significantly improves air quality and reduces CO<sub>2</sub> emissions, it also creates disparities: a small segment of the population bears most of the travel costs, while others benefit greatly from reduced congestion. Those most penalized are owners of banned cars living within the LEZ, particularly in areas with poor public transit access, while owners of authorized cars in heavily congested areas see the greatest benefits.

This study provides important insights into the wider effects of LEZs on urban mobility and offers a robust framework for evaluating transportation policies.

**Keywords:** transport simulation; low emission zone; road traffic emissions; air quality; calibration methodology; policy evaluation.

**JEL Codes:** C63; Q53; R4

## 3.1 Introduction

A Low Emission Zone (LEZ) is a designated area, typically in or around a city center, where the most polluting vehicles are restricted from entry. Since the 2010s, LEZs have become an increasingly popular tool for policymakers aiming to improve urban air quality. Hundreds of European cities have already implemented LEZs. According to the European Environmental Agency, an estimated 40 400 premature deaths in 2019 were attributed to chronic exposure to nitrogen dioxide (NO<sub>2</sub>).<sup>1</sup> Road transport accounts for 37 % of nitrogen oxides emissions in Europe, making it a key target for emissions reduction policies.<sup>2</sup>

Individual travel decisions – such as route choice, departure time, transportation mode, and vehicle ownership – are influenced by a complex interplay of factors, including time-dependent road congestion, scheduling constraints and the value of time. These factors are further complicated by the interdependence of individuals’ travel decisions, which affect one another through transportation networks. As a result, determining equilibrium travel conditions and evaluating the impact of policies that influence these decisions is a highly complex task.

The immediate effect of a LEZ is that individuals with restricted vehicles may choose to reroute around the zone, switch to another mode of transportation, or buy a new vehicle. This can reduce road congestion and emissions within the LEZ, but it could also increase congestion and pollution outside the zone if some drivers take long detours. Furthermore, some individuals with authorized vehicles may switch from public transit to driving, taking advantage of reduced congestion inside the LEZ – a phenomenon known as “rebound effect”. Consequently, predicting the detailed effects on congestion and air quality requires sophisticated model.

LEZs are expected to bring significant benefits, such as improved air quality as well

---

<sup>1</sup>Source: <https://www.eea.europa.eu/publications/air-quality-in-europe-2021/health-impacts-of-air-pollution> [accessed 2024/09/11]

<sup>2</sup>Source: <https://www.eea.europa.eu/publications/air-quality-in-europe-2022/sources-and-emissions-of-air> [accessed 2024/09/11]

as reduced road congestion, CO<sub>2</sub> emissions, and noise pollution. However, they often face public resistance due to their potential to limit mobility and exacerbate spatial and economic inequalities (see e.g., Tarrío-Ortiz et al. 2021; Player et al. 2023). A common criticism is that the LEZs affect mostly people owning banned vehicles who regularly travel within the LEZ, especially those living in low-density areas with limited access to public transit. They can thus negatively impact lower-income individuals, who are more likely to rely on older vehicles and live farther from city centers, while benefiting wealthier populations who live in areas with better public transit access and already own cleaner, authorized vehicles.

Estimating the impact of LEZs at both global and individual levels requires detailed transport simulations. In this study, we use METROPOLIS2, a dynamic agent-based transport simulator presented in Chapter 2. METROPOLIS2 simulates the equilibrium arising from the travel decisions of millions of individuals (referred to as *agents*) on large-scale road networks composed of up to hundreds of thousands of road segments. As an agent-based model, METROPOLIS2 provides detailed results into how each agent react to the LEZ policy (e.g., switching mode or route) and how they are affected (e.g., changes in travel surplus). METROPOLIS2 may also be coupled with METRO-TRACE (Le Frioux et al., 2024) to analyze the emissions of global and local air pollutants generated by road traffic and the exposure of the population to these pollutants, making it suitable for evaluating LEZs.

In this study, we evaluate the LEZ being implemented in the Greater Paris area, which is scheduled to expand in January 2025. This LEZ is substantial, covering 367 km<sup>2</sup>, encompassing approximately 5 million inhabitants, and affecting 19.7% of the region’s predicted vehicle fleet for 2025. The policy has been the subject of intense political debates.<sup>3</sup> To evaluate the LEZ’s impact, we compare a baseline simulation (without the LEZ) to a counterfactual simulation (with the LEZ as planned for 2025). To ensure the validity of the results, the baseline simulation must be correctly calibrated, meaning it accurately replicates observed

---

<sup>3</sup>See e.g. <https://www.40millionsdautomobilistes.com/articles/zones-a-faibles-emissions-z-fe-la-france-a-rebours> [accessed 2024/10/11] or <https://www.assemblee-nationale.fr/dyn/media/16/organes/commissions-permanentes-legislatives/developpement-durable/communication-miflash-zfem> [accessed 2024/10/11].

data, such as mode shares, departure-time distribution and road congestion levels.

We propose a novel and efficient methodology for calibrating METROPOLIS2 (or other transport simulators) to observed data. This four-step calibration procedure leverages the hierarchical nature of individual travel decisions (mode choice, followed by departure time and route choice) to fine-tune parameters without requiring computationally expensive simulations at each step. The first two steps rely on data from the TomTom API, processed with map-matching algorithms and linear regressions, to calibrate both free-flow and congested travel times. The third step uses Bayesian optimization with Gaussian processes to calibrate departure-time distributions based on travel survey data. Finally, the fourth step calibrates mode shares using Random Forest regressions.

In both the baseline and LEZ simulations, five modes of transportation are considered: car driver, car passenger, public transit, bicycle and walking. We adopt a short-term perspective, assuming fixed household locations, activity locations, and vehicle ownership. In practice, the policy could lead to longer-term behavior changes, such as households relocating closer to transit hubs or purchasing authorized vehicles.

The aggregate results from the METROPOLIS2 simulations suggest that the LEZ will decrease the mode share of car trips by 1.9 percentage point (p.p.) – from 36.6 % to 34.7 % – reducing total vehicle-kilometers by 3.9 %. Pollution analysis indicates larger reductions in emissions:  $-4.5\%$  for  $\text{CO}_2$ ,  $-9.2\%$  for  $\text{NO}_x$ , and  $-7.6\%$  for  $\text{PM}_{2.5}$ , demonstrating the policy’s effectiveness in targeting the most polluting vehicles. Additionally, the number of premature deaths due to  $\text{NO}_2$  and  $\text{PM}_{2.5}$  exposure is expected to decrease by 9.9 % and 13.0 %, respectively, due to the greatest air quality improvements occurring inside the LEZ, where population density is higher.

At the individual level, the health benefits of the LEZ are distributed relatively equitably, with changes in individual health exposure ranging from 0.00 € to +0.35 € daily. However, the distribution of travel surplus changes shows greater disparities, with 3.3 % of the population losing more than 1 € daily, while 1.2 % gain more than 1 €. These disparities may

explain the LEZ's low public acceptance.

Further analysis reveals that the LEZ disproportionately impacts owners of banned vehicles living within the LEZ, who face longer travel times when switching to alternative modes. Conversely, the biggest beneficiaries are authorized vehicle owners living outside the LEZ and traveling in highly congested areas who gain from reduced congestion. However, no clear pattern emerges regarding the policy's effect on economic inequalities, as both winners and losers are spread across high- and low-income municipalities.

The chapter is organized as follows. Section 3.2 provides a brief literature review. Section 3.3 describes the theoretical framework used to model travel decisions (demand side) and road congestion (supply side). Section 3.4 outlines the data sources and the preprocessing steps. Section 3.5 details the calibration methodology. Section 3.6 discusses the Paris LEZ policy. Section 3.7 presents the simulation results and analyzes the LEZ's impact. Finally, Section 3.8 summarizes the findings and suggests extensions for future research.

## 3.2 Literature Review

The literature review for this chapter is divided in two sections. The first section focuses on previous works related to the calibration of agent-based transport simulators. The second section reviews studies evaluating Low Emission Zones.

### 3.2.1 Calibration of agent-based transport simulators

Zhao and Sadek (2012) discuss the calibration of the TRANSIM agent-based model in the Buffalo-Niagara metropolitan area. They focus on adjusting the departure-time distribution (exogenous in their model) and total demand to match traffic count data from 162 stations. Due to the long runtime of the simulation, calibration was conducted manually through trial and error.

More recently, Ziemke et al. (2019) describe the calibration of the MATSim simulator for

Berlin. Our approach improves upon their calibration process at each step. For instance, while we calibrate road free-flow speeds using origin-destination free-flow travel times from the TomTom API data and applying a Lasso regression, they instead reduce inner-city free-flow speeds by half to account for factors like traffic lights or acceleration / deceleration at intersections.

For congestion calibration, Ziemke et al. employ the CaDyTS calibration procedure (Flötteröd, 2009; Flötteröd et al., 2012), which leverages MATSim’s co-evolutionary algorithm to match traffic count data. However, their approach assumes fixed road capacities and instead adjusts agents’ activity patterns to better reflect the observed traffic counts. By contrast, our calibration uses origin-destination travel time data and iteratively combines simulations with linear regressions to adjust road capacities.

Departure-time calibration also differs significantly. They assume that agents may randomly shift departure times within a two-hour window, without any specific calibration. In contrast, our method calibrates purpose-specific parameters to replicate the observed departure-time distribution, divided into clusters.

Finally, while Ziemke et al. manually calibrate a mode-specific constant for each mode to match observed mode shares, our approach calibrates both utility constants and values of time for each mode across twelve population segments, based on gender and socio-professional category, to match observed mode shares, divided into clusters.

Unlike the approaches in these studies, our calibration does not rely on traffic count data, though such data could be valuable as validation in future work.

### 3.2.2 Evaluation of Low Emission Zones

Numerous studies have explored the environmental, social, and economic impacts of Low Emission Zones (LEZs), offering valuable insights into their design and implementation. A central issue in LEZ design is the need to balance local and broader societal interests. De Borger and Proost (2013) develop an analytical model that compares the implementation

of LEZs and other traffic policies by local versus federal governments. Their analysis shows that local governments, focusing primarily on the welfare of residents, tend to impose more stringent emission standards compared to federal governments, which must consider the welfare of non-resident commuters.

In terms of costs – drivers having to adapt to the restrictions –, Börjesson et al. (2021) analyze the impact of Stockholm’s LEZ using two methodologies: (i) measuring user cost increases related to the observed reduction in traffic volumes, and (ii) estimating driver losses based on price changes in the used car market. Their findings suggest that the costs of implementing Stockholm’s LEZ outweigh its benefits, underlining the importance of a thorough cost-effectiveness evaluation when designing LEZ policies.

The environmental benefits of LEZs are well-documented. Chamberlain et al. (2023) conducted a systematic review of empirical studies, concluding that LEZs generally reduce air pollution and improve public health. Similarly, Holman et al. (2015) found that LEZs in Germany led to significant reductions in pollution. However, they note that isolating the effect of LEZs from other policies or from the natural renewal of the vehicle fleet can be challenging. This emphasizes the need for sophisticated tools, such as transport simulators, to assess LEZ impacts accurately.

Regarding the case of Paris, Poulhès and Proulhac (2021) propose a methodology to assess the health benefits of the LEZ at an individual level. By combining pollutant concentration data, a household travel survey, and a road traffic model, they estimate population exposure to pollution. While we adopt a similar approach to compute pollution exposure – by considering the spatial and temporal location of individuals over the course of a day – their methodology is limited to an *a priori* evaluation. This is because they rely on observed pollutant concentrations before and after the policy’s implementation. In contrast, we employ an *ex-ante* approach, simulating the impact of the LEZ policy on trips, pollutant emissions, and concentrations, prior to its implementation.

Host et al. (2020) provide an *ex-ante* evaluation of Paris’ LEZ using a complete chain



of model to estimate pollutant emissions and population exposure, similar to the METRO-TRACE framework we employ (Le Frioux et al., 2024). However, unlike our work, which simulates how agents adapt their routes, departure times, and modes in response to the LEZ, their analysis relies on predefined assumptions about the share of agents who would switch to a newer vehicles, modify their itineraries, or shift to public transit.

Yin et al. (2024) present a socioeconomic evaluation of Paris’ driving restriction zone using the MATSim transport simulator (for a comparison between METROPOLIS2 and MATSim, see Chapter 2). Their study evaluates a driving restriction that applies to all vehicles equally, unlike the LEZs, which target only the most polluting vehicles. They explore changes in traffic behaviors, including rerouting and mode shifts, and analyze emission variations inside and outside the restricted zone, with a particular focus on intermodality. While both policies share similarities, LEZs typically target only the most polluting vehicles and cover larger areas. Our study extends their work by considering vehicle heterogeneity, in terms of age and fuel type, which can have a significant impact on emissions. Additionally, we perform a more granular analysis at the agent level, evaluating individual travel surplus and pollution exposure.

Other recent evaluations of transport policies in Paris metropolitan area have been presented by Durrmeyer and Martinez (2022) and Bou Sleiman (2023). Durrmeyer and Martinez (2022) assess the impact of driving restrictions – banning a fraction of cars randomly – and road tolls using a structural model. This model allows them to compute the welfare consequences of the policy, including the impact on inequalities. Although this approach is more tractable than agent-based simulations, it simplifies the spatial analysis by dividing the region into only five zones and assuming homogeneous congestion within each zone. In contrast, transport simulators like METROPOLIS2, though computationally complex, allow for detailed simulations of traffic dynamics across extensive road networks, providing richer insights into traffic assignment and equilibrium.

Bou Sleiman (2023) take an ex-post approach, using a difference-in-difference methodol-

ogy to evaluate the impact of a road closure in Paris. While this technique efficiently isolates policy effects, it is limited to post-implementation analysis, whereas this study focuses on an ex-ante assessment of the LEZ policy.

Although the environmental benefits of LEZs are well-established, their broader impacts on road traffic, and spatial and socioeconomic inequalities require further investigations. By using agent-based simulations, and employing advanced modeling techniques, this study contributes to a more comprehensive evaluation of LEZ policies, shedding light on both their effectiveness and equity implications.

### 3.3 Definitions and Theoretical Foundations

The transport simulator METROPOLIS2 (see Chapter 2) is a dynamic mesoscopic agent-based model designed to compute a Nash equilibrium of the interaction between supply (network infrastructure) and demand (agents traveling). While METROPOLIS2 is highly flexible, this chapter focuses on a specific model specification described below.

#### 3.3.1 Demand Side

The demand side of the model is characterized by a population of agents, indexed by  $n$ , who travel between various activities. These agents and their activities (including the activities' locations) are generated to be representative of an average weekday (see Section 3.4.2).

The activity purposes considered are *home*, *work*, *education*, *leisure*, *shopping* and *other*. All agents begin and end the day at home, with an arbitrary number of activities to be completed outside of home during the day. The location and duration of all outside activities are treated as exogenous.

To travel between activities, agents perform *trips*. Let  $\mathcal{K} = \{1, \dots, k, \dots, K_n\}$  denote the set of trips performed by an agent to complete all their activities, where  $K_n$  is the number of trips.

An agent's trips are partitioned into one or more *tours*. Formally,  $\mathcal{K}_q \subseteq \mathcal{K}$  denotes the set of trips for tour  $q$ , with  $\cup_q \mathcal{K}_q = \mathcal{K}$ . A tour  $q$  is defined as a sequence of successive trips where the purpose of the activity at the origin of the first trip and at the destination of the last trip is always *home*, with no *home* activity in between. Figure 3.1 illustrates an example agent with two tours. As shown in the figure, all activities have fixed duration, except for *home* activities, which serve as the default.




Figure 3.1: Example activity pattern of an agent with two tours

For each tour, the agent selects a transport mode and the departure time of the first trip, maximizing their utility. The utility of a tour is defined as the sum of utilities of its trips, where the utility of trip  $k$  with mode  $j$ , departure time  $t^d$  and arrival time  $t^a$  is given by (the agent index is omitted for readability)

$$V_{k|j}(t^d, t^a) = c_j - m_{k|j} \underbrace{-\alpha_j \cdot [t^a - t^d]}_{\text{travel utility}} \underbrace{-\beta_k [t_k^* - t^a]_+ - \gamma_k [t^a - t_k^*]_+}_{\text{schedule-delay utility}}, \quad (3.1)$$

where  $c_j$  is the mode-specific utility constant,  $m_{k|j}$  are the monetary costs (fuel cost for car trips),  $\alpha_j$  is the value of time,  $\beta_k$  is the penalty for early departure,  $\gamma_k$  is the penalty for late departure,  $t_k^*$  is the desired start time of the activity at destination of trip  $k$  – or equivalently, the desired arrival time at destination for trip  $k$  –, and  $[x]_+ = \max(x, 0)$ . The value of time depends on the selected mode, while the schedule-delay penalties depend on the purpose of the activity. These preference parameters are agent-specific. The activities' desired start times,  $t_k^*$ , are determined when generating the synthetic population (see Section 3.4.2). The utility function (3.1) is derived from the alpha-beta-gamma model (Vickrey, 1969; Arnott et al., 1990).

For a given tour  $q$ , with mode  $j$  and departure time for the first trip  $t$ , the departure times

$(t_k^d)_k$  and arrival times  $(t_k^a)_k$  for all tour trips  $k \in \mathcal{K}_q$  can be inferred from the mode-specific travel-time function of each trip. The deterministic utility of tour  $q$  when selecting mode  $j$  and departure time  $t$  is given by the sum of the trips' utilities:

$$V_{q|j}(t) = \sum_{k \in \mathcal{K}_q} V_{k|j}(t_k^d, t_k^a). \quad (3.2)$$

Agents are assumed to choose their mode first and then their departure time. They can choose from five transport modes: car as a driver, car as a passenger, public transit, bicycle and walking.<sup>4</sup> All trips in a tour must be completed using the same mode but different tours can involve different modes for the same agent. The mode choice follows a Multinomial Logit model and the departure-time choice follows a Continuous Logit model (Ben-Akiva and Watanatada, 1981).

For each tour, the agent selects the departure time  $t$  for the first trip by maximizing their utility  $V_{q|j}(t) + \varepsilon_q(t)$ , where the deterministic utility  $V_{q|j}(t)$  is given by Equation (3.2) and  $\varepsilon_q(t)$  is an idiosyncratic random component. With the Continuous Logit formula, the probability that departure time  $t$  is chosen for tour  $q$ , given that mode  $j$  is selected, is

$$p_{q|j}^d(t) = \frac{e^{V_{q|j}(t)/\mu_1}}{\int_{t^0}^{t^1} e^{V_{q|j}(\tau)/\mu_1} d\tau},$$

where  $\mu_1$  is the scale parameter of the Continuous Logit model and  $[t^0, t^1]$  is the feasible departure-time period. The expected utility that the agent derives from the departure-time choice is given by the ‘‘logsum’’ formula:

$$V_{q|j} = \mu_1 \ln \int_{t^0}^{t^1} e^{V_{q|j}(t)/\mu_1} dt + \mu_1 \cdot \gamma, \quad (3.3)$$

where  $\gamma$  is Euler-Marscheroni constant.

The agent then selects the mode  $j$  that maximizes  $V_{q|j} + \varepsilon_{q,j}$  where  $V_{q|j}$  is the deterministic

---

<sup>4</sup>These five modes encompass more than 98% of trips made in the study area.

utility of choosing mode  $j$ , defined by Equation (3.3), and  $\varepsilon_{q,j}$  is an idiosyncratic random component with Gumbel distribution of scale  $\mu_2$ . The probability to choose mode  $j$  for tour  $q$  is then given by the Multinomial Logit formula:

$$p_{q|j} = \frac{e^{V_{q|j}/\mu_2}}{\sum_{j'} e^{V_{q|j'}/\mu_2}}.$$

Given the mode and departure time selected, agents choose the fastest route connecting origin to destination for all their trips.

For each tour performed, the agent obtains a *travel surplus*,  $\bar{V}_q$ , defined as the expected utility that they will get from the combined mode and departure-time choice. It is defined as the “logsum” of the upper-level decision (the mode choice):

$$\bar{V}_q = \mu_2 \ln \sum_j e^{V_{q|j}/\mu_2} + \mu_2 \cdot \gamma, \quad (3.4)$$

where  $\gamma$  is Euler-Marscheroni constant. The travel surplus of the agent is defined as the sum of the travel surplus for all their tours.

In addition to the trips performed by agents, truck trips are also simulated. For these trips, mode choice is irrelevant and the departure times are exogenously determined, making route choice the only decision. The primary reason for including truck trips in the simulation is to account for the road congestion they generate, and the impact it has on car trips.

### 3.3.2 Supply Side

Car and truck trips take place on a road network, represented as a directed graph of nodes (intersections) and edges (road segments).

Road congestion occurs in two ways. First, road-level *bottlenecks* limit vehicle entry and exit flows based on a predefined road-level capacity. If two vehicles arrive within a time interval shorter than allowed by the capacity, the second vehicle is delayed at the bottleneck.

If a vehicle arrives and there is already a queue, it must wait at the end of the queue. Capacity is measured in passenger car equivalent (PCE), where we assume, following the *Highway Capacity Manual* (2016), that cars are equal to 1 PCE and trucks are equal to 2 PCE.

Second, the number of vehicles on a road segment is limited by the segment's total length (actual length multiplied by the number of lanes). If a vehicle reaches a full road, it must wait on its current road segment until space becomes available. This phenomenon is known as *spillback*. The length a vehicle occupies while on a road segment is assumed to be 8 m for cars and 16 m for trucks.

Additionally, speed is limited to 90 km/h for trucks in the simulations.

For public transit, bicycle, and walking, travel times are assumed to be constant, independent of departure time or traffic congestion. Details on how travel times are computed for these modes can be found in Section 3.4.3.

## 3.4 Data Input

This section outlines the data collection and preprocessing steps necessary to generate the inputs for the simulation. The preprocessing tasks are divided into three categories: (i) generating the supply side (road and public-transit networks), (ii) generating the demand side (synthetic population), and (iii) additional mode-specific processing. The study area for the simulation is Île-de-France, a 12 011 km<sup>2</sup> administrative region centered on Paris.

### 3.4.1 Supply Side

We use OpenStreetMap (OSM) data to generate the road network (used by cars and trucks) and the walking network (used by pedestrian and bicycles). The data is obtained from Geofabrik for the Île-de-France region, using a snapshot of January 1st, 2024. Various preprocessing steps are undertaken to refine the data. For instance, bus-only roads are

excluded from the road network, and bidirectional roads are represented as two directed edges. The final output consists in two directed graphs representing the road and walking networks.

For the road network, edge-level data includes length, speed limit, number of lanes and additional attributes used for calibration, such as indicators for urban area and the presence of traffic lights. The imported road network includes 294 706 nodes and 610 629 edges, for a total of 72 962 kilometers. A chunk of the road network is shown in Figure 3.2, illustrating the various road types.

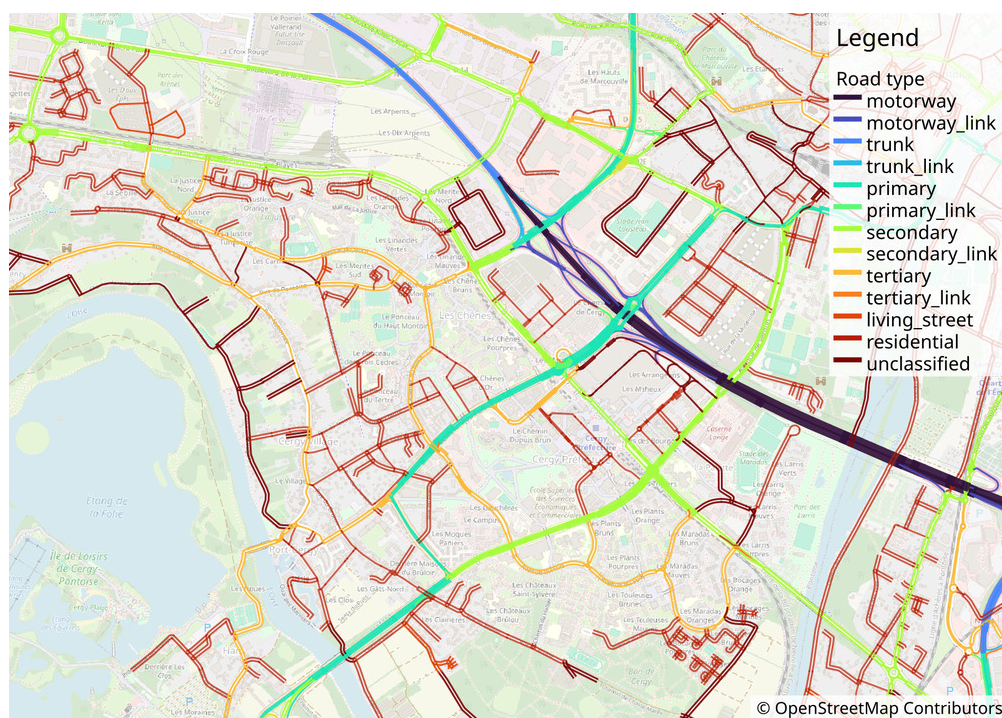


Figure 3.2: Road segments imported from OpenStreetMap, near Cergy

The walking network stores only the length of the edges. It includes 1 003 920 nodes and 2 704 016 edges, for a total of 189 803 kilometers.

Public-transit trip data is provided by *Île-de-France Mobilités* and based on the General Transit Feed Specification (GTFS), which includes schedules for all public-transit lines in the region. Additionally, OSM data is used for the walking legs of the trips.

### 3.4.2 Demand Side

The demand side of the simulation is based on a synthetic population, which consists of households and individuals along with their respective trips (or activities) for a typical working day. A synthetic population is a simulated representation of a real population, designed to replicate its characteristics. The synthetic population of Île-de-France is generated using the methodology described in Hörl and Balac (2021).

This methodology relies on various open datasets, including census data, the regional travel survey and datasets on buildings and firms. The travel survey used is the *Enquête Globale Transport*, a regional survey conducted in 2010 specific to Île-de-France. Throughout this chapter, the survey is referred to as the *regional travel survey* and is used extensively for calibration purpose.

The methodology to generate a synthetic population is extended in three ways.

**Vehicle Characteristics Assignment** Given the number of cars owned by each household, we assign car characteristics (e.g., fuel type and age) from the predicted 2025 vehicle fleet. To do this, we use data from the French Ministry of Ecology, which provides municipality-level fleet characteristics for each year between 2011 and 2022. A basic Markov Chain model is applied to interpolate the 2025 vehicle fleet. These car characteristics are used to compute fuel consumption and pollutant emissions, as well as to identify cars eligible to enter the Low Emission Zone.

**Activity Desired Start Times and Duration** The synthetic population generates a start time and duration for each activity. However, the *desired* start times and duration, required to derive the utility of trips, are not available. We thus develop a methodology to simulate them from the other characteristics available in the synthetic population and from the regional travel survey.<sup>5</sup> Currently, this is done only for *work* activities when they are the

---

<sup>5</sup>The work described here is inspired from an unpublished work with André de Palma and Nathalie Picard.



sole activity in a tour.<sup>6</sup>

The methodology relies on the assumption that, for trips with a time-independent travel time, the *desired* start time and duration coincide with the *actual* start time and duration because the individuals have no reason to choose a different start time or duration.<sup>7</sup> We can thus derive the joint distribution of desired activity start time and duration, from the observed trips with a constant travel time (walking and bicycle trips) in the regional travel survey. However, the probability to travel by walking or cycling can be correlated with variables like job type or job location, which in turn influence the desired activity start time and duration. To control for this endogeneity problem, we derive a distribution for each pair of socio-professional category (employees, intermediate category, upper category and blue-collar workers) and workplace area (Paris, inner ring and outer ring). For example, from these distributions, we observe that blue-collar workers in the outer ring tend to start work two hours earlier on average than upper-category workers in Paris.

To simulate a desired start time and duration for each *work* activity in the synthetic population (when it is the sole activity in a tour), we simply draw randomly values from the distribution corresponding to the individual's socio-professional category and workplace area.

**Mode Choice Prediction** To estimate the transport mode for each trip in the synthetic population, we fit a Multinomial Logit model on the regional travel survey data. The model includes variables such as Euclidean origin-destination distance, departure time and origin /

---

<sup>6</sup>Extending this methodology to all activities and tours would require extensive work and is postponed to future research. For now, the start time and duration provided in the synthetic population are used as the *desired* values.

<sup>7</sup>In the standard alpha-beta-gamma (Arnott et al., 1990), when the travel time is equal to a constant  $\bar{t}$ , the departure time minimizing the generalized cost is

$$\operatorname{argmin}_t \alpha \cdot \bar{t} + \beta[t^* - t - \bar{t}]_+ + \gamma[t + \bar{t} - t^*]_+ = t^* - \bar{t},$$

that is, the departure time such that the arrival time at destination (the activity start time) is equal to the desired arrival time at destination  $t^*$  (the desired activity start time).

destination departments of the trips.<sup>8</sup> For each trip, we draw a mode from the probabilities predicted by the Multinomial Logit model. These mode predictions are used only during the calibration process to select trips for each mode in a manner that is representative of the real population (See Section 3.5). The actual mode choice for all trips is determined dynamically within METROPOLIS2, using the mode choice model described in Section 3.3.

Ultimately, the synthetic population data includes a list of *households*, with their car characteristics and income; a list of *persons*, including their household affiliations and socio-demographic attributes (e.g., age, gender, socio-professional category); and a list of *trips / activities* for each person, including activities' desired start times and duration, as well as their purpose and precise coordinates.

To reduce simulation runtime, the synthetic population is generated to represent 20 % of the total population of Île-de-France. During calibration, road capacities are adjusted to ensure that the road congestion observed with the 20 % synthetic population replicates the congestion of the full population. Aggregate results are subsequently scaled up to represent 100 % of the population, meaning that vehicle-kilometers traveled, for example, are multiplied by five.

The resulting synthetic population consists of 1 095 819 households, 2 451 841 persons and 8 774 929 trips.

The truck trips are generated based on an origin-destination matrix provided by DRIEAT, the regional public authority responsible for transport (DRIEAT Île-de-France, 2021). The trips' origin and destination zones are mapped to actual road intersections in order to spread the departure and arrival points over the zones. A total of 599 669 truck trips are simulated.

---

<sup>8</sup>Note that the Île-de-France region is divided into eight departments.

### 3.4.3 Mode-specific Processing

The simulation considers five transport modes: car driver, car passenger, public transit, bicycle and walking. For car drivers and passengers, travel times are computed dynamically within the simulation based on traffic conditions at the time of departure. For all other modes, travel times are assumed to be constant, independent of departure time or traffic congestion.

**Car driver and Car passenger** To simulate car trips in METROPOLIS2, the road-network origin and destination nodes where the trip starts and ends must be identified. Each origin and destination coordinates are first assigned to the nearest road segment. Then, the origin (or destination) node is defined as the segment endpoint which is closest to the origin (or destination) coordinates. This process is illustrated in Figure 3.3.




Figure 3.3: Process to assign an origin / destination node from activity coordinates  
 Note. The closest road segment from origin is the orange road segment at 25 m. The source node (36 m) is closer than the target node (92 m) and is thus used as the start node of the trip.

The majority of the road segments are residential streets, used only during the first or last

meters of the trips, with congestion unlikely to occur. To optimize the simulator’s runtime, we ignore the congestion on residential streets at the start and end of trips, allowing to reduce the size of the simulated road network by 46 %.

The car driver and car passenger modes are available only to agents whose households own at least one car. Additionally, only agents with a driving license can select the car driver alternative. However, the model does not enforce the constraint that each car in the household can only be used simultaneously by a single agent. Furthermore, any agent can choose the car passenger alternative even if no other agent in the household – or even in the entire simulation – is making the same trip as a car driver. These limitations could be addressed in future research.

Fuel consumption for car trips is calculated based on the fastest route under free-flow conditions, using the EMISENS model (Ho et al., 2014). The value appears as a factor in the trip utility function (3.1). Computing fuel consumption based on the actual route taken would require running the EMISENS model at each iteration of the simulator, considerably increasing computation time. However, an *a priori* analysis reveals a strong correlation (99 %) between fuel consumption under free-flow conditions and actual fuel consumption, justifying the use of free-flow estimates in the simulation.

While both car drivers and passengers are impacted by road congestion, only car drivers contribute to the generation of congestion in the simulator.

**Public Transit** For public-transit, travel time is determined by the least-cost itinerary among all itineraries arriving within 30 minutes of the desired arrival time. Public-transit itineraries are computed using the open-source software OpenTripPlanner. The cost of an itinerary accounts for the in-vehicle time, walking time and waiting time. The public-transit alternative is not available at the tour level if one of the trips is not feasible (e.g., when the origin is too far away from a bus stop or train station). Moreover, the public-transit alternative is omitted from the choice set of the agent if it is faster to walk than to take

public transit for all trips of a tour.

Public-transit pricing is not modeled in this chapter and thus is captured by the public-transit utility constant (see Section 3.5.4).

The current model does not account for crowding in public transit, meaning that agents do not face additional costs when boarding crowded vehicles, and that vehicles can exceed their capacity. This limitation may introduce a bias towards public transit in the simulation, as crowding is a major issue in the Paris area (e.g., Haywood and Koning 2015). However, this bias is likely limited, as the results show a small increase of public-transit use, especially on the most congestion lines (e.g., only a 1.2% rise in passenger-kilometers on the main transit line, RER A). In a future version of METROPOLIS2, public-transit vehicles could be modeled explicitly which would allow for more detailed considerations of waiting time, number of transfers, in-vehicle congestion and reliability.

**Bicycle and Walking** For both bicycle and walking trips, travel time is assumed to be proportional to the distance of the shortest path on the walking network. Similarly to the car modes, the origin and destination coordinates are mapped to the nearest segment of the walking network to compute the shortest path. Average speeds are assumed to be 10 km/h for bicycles and 4 km/h for walking. These two modes are always included in the choice set of the agents (bicycle ownership is not explicitly modeled).

## 3.5 Calibrating the Model

Calibration is a critical yet often overlooked step in running a simulation. It involves adjusting the simulation parameters to closely replicate observed data, such as mode shares or travel times, in the baseline scenario. Given that each simulation run can take several hours and that hundred of parameters may need to be calibrated, this step presents significant challenges that cannot be easily addressed using standard optimization methods.

To overcome these challenges, we designed a four-step calibration methodology, outlined

in Table 3.1. Each step focuses on calibrating a specific subset of parameters while matching some observed data. The first step involves a highly simplified simulation (no congestion, no departure-time choice, no mode choice), where road penalties are calibrated to match observed free-flow travel times. As the process progresses, the simulation’s complexity increases incrementally (first enabling congestion, then departure-time choice, and finally mode choice) and additional parameters need to be calibrated such as road capacities, schedule-delay penalties and values of time. Importantly, the parameters calibrated in earlier steps do not require further adjustment in subsequent steps, as the added complexity does not affect the previously established match between simulated and observed data. This approach ensures that all steps remain computationally feasible.

Table 3.1: Summary of the calibration process

Step	Calibrated parameters	Target values	Target source	Methodology
1	Road constant penalties and free-flow speed	Free-flow travel times	TomTom API	Lasso regression
2	Road capacities	Time-dependent congested travel times	TomTom API	OLS regression
3	Schedule-delay penalties by purpose	Distribution of departure times by cluster	EGT	Bayesian Optimization and Gaussian Process
4	Mode-specific utility parameters by socio-demographic characteristics	Mode shares by cluster	EGT	Random Forest regression

Note. “EGT” is the regional travel survey (*Enquête Globale Transport*).

### 3.5.1 Free-Flow Travel Times Calibration

The first step in the calibration process adjusts road penalties and free-flow speeds to match simulated free-flow travel times to observed data. Free-flow travel times are defined as the travel times on an empty road network, where no delay is caused by congestion. These times do not simply correspond to vehicles traveling at the speed limit, due to factors such as waiting at traffic lights and acceleration / deceleration at intersections or in curves.

To calibrate free-flow travel times, we use data from the TomTom API, which provides both congested and free-flow travel times for any origin-destination (OD) pair. We assume

that the free-flow travel time for each road segment  $e$  is given by

$$t_e = \delta_e^{\text{add}} + \delta_e^{\text{mul}} \cdot t_e^{\text{f}},$$

where  $\delta_e^{\text{add}}$  is an additive penalty,  $\delta_e^{\text{mul}}$  is a multiplicative penalty, and  $t_e^{\text{f}}$  is the “travel time at speed limit”, computed by dividing the segment’s length by its speed limit.

Instead of estimating these penalties for each road segment individually, which could lead to overfitting, we estimate them for  $J$  categories of road segments, based on their characteristics (e.g., speed limit, presence of traffic lights). Each road segment belongs to one or more categories. The additive and multiplicative penalty for a road segment are then calculated as the sum of category-specific penalties:

$$\delta_e^{\text{add}} = \sum_j x_{e,j} \cdot \delta_j^{\text{add}} \quad \text{and} \quad \delta_e^{\text{mul}} = \sum_j x_{e,j} \cdot \delta_j^{\text{mul}}, \quad (3.5)$$

where  $\delta_j^{\text{add}}$  is the additive penalty specific to category  $j$ ,  $\delta_j^{\text{mul}}$  is the multiplicative penalty specific to category  $j$ , and  $x_{e,j} = 1$  if segment  $e$  belongs to category  $j$ .

We estimate the category-specific parameters  $\delta_j^{\text{add}}$  and  $\delta_j^{\text{mul}}$  using a Lasso regression (Tibshirani, 1996), through the following process. First, we retrieve the free-flow travel time,  $y_n$ , and the fastest path,  $p_n$ , for a set of OD pairs,  $\{1, \dots, n, \dots, N\}$ , from the TomTom API (with  $N \approx 200\,000$ ). Then, a map-matching algorithm is used to match each path,  $p_n$ , to the corresponding set of road segments,  $E_n$ , in our OpenStreetMap-based road network. Finally, we run the following Lasso regression:

$$\underset{\delta_j^{\text{add}}, \delta_j^{\text{mul}}}{\operatorname{argmin}} \frac{1}{N} \sum_n \left( y_n - \sum_j \delta_j^{\text{add}} \cdot c_{n,j} - \sum_j \delta_j^{\text{mul}} \cdot t_{n,j}^{\text{f}} \right)^2 + \lambda \sum_j (|\delta_j^{\text{add}}| + |\delta_j^{\text{mul}}|) \quad (3.6)$$

where  $\lambda$  is the Lasso penalty,  $c_{n,j} = \sum_{e \in E_n} x_{e,j}$  is the number of road segments in category  $j$  along path  $p_n$ , and  $t_{n,j}^{\text{f}} = \sum_{e \in E_n} x_{e,j} \cdot t_e^{\text{f}}$  is the total “travel time at speed limit” for road segments in category  $j$  along path  $p_n$ .

The Lasso model mitigates the risk of over-fitting that may occur with a large number of categories, by penalizing the size of the coefficients. The penalty  $\lambda$  is selected via cross-validation, a method where the data is split into subsets: the model is trained on some subsets and tested on others. This process ensures that  $\lambda$  is chosen to balance accuracy on unseen data with simplicity in the model.

Using the estimated parameters  $\delta_j^{\text{add}}$  and  $\delta_j^{\text{mul}}$ , we compute the penalties specific to each edge,  $\delta_e^{\text{add}}$  and  $\delta_e^{\text{mul}}$  from Equation (3.5). Observe that the free-flow travel time for path  $n$  in METROPOLIS2 will thus be equal to

$$\hat{y}_n = \sum_{e \in E_n} t_e = \sum_{e \in E_n} \sum_j x_{e,j} \cdot \delta_j^{\text{add}} + \sum_{e \in E_n} \sum_j x_{e,j} \cdot \delta_j^{\text{mul}} \cdot t_e^f = \sum_j \delta_j^{\text{add}} \cdot c_{n,j} + \sum_j \delta_j^{\text{mul}} \cdot t_{n,j}^f,$$

so that the Lasso regression (3.6) amounts to minimizing the mean squared distance between the free-flow travel times from TomTom,  $y_n$ , and the ones simulated in METROPOLIS2,  $\hat{y}_n$  (with a penalty for the size of the coefficients).

The calibration results show, for example, that an additive penalty of 9.06 seconds applies to road segments with traffic lights in urban areas (7.86 seconds in rural areas). After calibration, the root-mean squared error (RMSE) between TomTom free-flow travel times and those computed in METROPOLIS2 is reduced to 101 seconds, down from 283 seconds for the “uncalibrated” values, defined as using simply the “travel time at speed limit” with no penalty. Figure 3.4 shows a comparison of the distribution of OD-level free-flow travel times between TomTom and METROPOLIS2 (calibrated and uncalibrated).

This calibration step is computationally efficient since it does not require running a full simulation.

### 3.5.2 Congested Travel Times Calibration

The second step in the calibration methodology focuses on adjusting road capacities at bottlenecks to match observed congested travel times. Road capacity measures the number






Figure 3.4: Density of OD-level free-flow travel times (TomTom, calibrated, uncalibrated) of vehicles, in passenger car equivalents (PCE), that can pass through the entry and exit bottlenecks of a road segment per hour, per lane.

To perform this calibration, we run simulations using METROPOLIS2 to compute congested travel times under varying road capacity values. These simulations use fixed departure times and modes (to be calibrated in the next steps). The departure times are set to the ones generated in the synthetic population (they are generated to be representative of the departure time distribution observed in the regional travel survey) and the modes are set to the predicted modes of the trips (see Section 3.4.2). Only the trips by car drivers need to be simulated as the other trips do not influence congestion. This selection ensures that the simulated trips accurately represent the trips observed during an average working day, in terms of departure times, origins and destinations. Truck trips are included to account for their contribution to road congestion.

The simulated congestion levels are compared to observed data from the TomTom API, as in the previous step, but now focusing on congested travel times. For each origin-destination (OD) pair  $n$ , we retrieve the time spent in congestion,  $y_n$  (the difference between congested and free-flow time), and the fastest path,  $p_n$ . We gather data for approximately 200 000 OD

pairs, with randomly selected departure times to observe the varying congestion throughout a typical working day. The day for the requests is set to a random working day in November.

Similar to the free-flow calibration step, we use a map-matching algorithm to match TomTom paths to the corresponding road segments in the OpenStreetMap-based network. Road segments are classified into  $J$  categories based on attributes such as road type, presence of traffic signals, and roundabouts, with each segment assigned to a single category.

The goal is to calibrate  $J$  capacity parameters— one for each category – so that the time spent in congestion from TomTom,  $y_n$ , aligns with the simulated time, denoted by  $\hat{y}_n$ . The congestion time can be decomposed by category as:

$$\hat{y}_n = \sum_j x_{n,j},$$

where  $x_{n,j}$  represents the time spent in congestion on segments in category  $j$  for OD pair  $n$ .

The calibration is performed iteratively. Starting from initial capacity values, we run a METROPOLIS2 simulation using only the selected car and truck trips, with route choice enabled but without departure-time or mode choice. From the simulation, we extract the congested times,  $x_{n,j}$ , for each category  $j$  and each OD pair  $n$ . Using Ordinary Least Squares, the observed congested times,  $y_n$ , are then regressed on the simulated times:

$$\operatorname{argmin}_{\delta_j} \sum_n \left( y_n - \sum_j \delta_j \cdot x_{n,j} \right)^2,$$

where the parameters  $\delta_j$  indicate whether the congestion on segments in category  $j$  is under- or over-estimated. If  $\delta_j < 1$ , the congestion for that category is over-estimated, meaning the capacity should be increased. Conversely, if  $\delta_j > 1$ , the capacity should be reduced.

Since congestion at bottlenecks is roughly inversely proportional to capacity, multiplying the simulated congestion time by  $\delta_j$  can be achieved by multiplying the bottleneck capacity by  $1/\delta$ . However, the relationship becomes more complex with multiple bottlenecks, due to agents adapting their routes and due to spillback effects. After adjusting the capacity

for all segments in each category, the iterative process continues, until the  $\delta_j$  parameters converge to values close to 1, indicating that the simulated and observed congestion times are well-matched

Table 3.2 shows the calibrated road capacities for each road type, distinguishing between segments with and without traffic signals or roundabouts. The results align with expectations, showing that major roads (motorways, trunks, primary roads) have higher capacities compared to minor roads (tertiary roads, residential streets, living streets) and that capacity is lower for roads with traffic signals or roundabouts.

Table 3.2: Calibrated capacities by road type (in PCE per hour per lane)

Road type	Base value	With traffic signals	With roundabout
motorway	1700	N/A	680
motorway_link	1350	945	N/A
trunk	1850	N/A	740
trunk_link	1350	945	N/A
primary	2150	1505	860
secondary	1700	1190	680
tertiary	1400	980	560
residential	1300	910	520
unclassified	1000	700	400
living_street	900	630	360

It can seem surprising that the model predicts the capacity to be higher for primary roads (2150 PCE per hour per lane) than for motorways (1700 PCE per hour per lane). The primary roads usually represent long urban avenues with regular intersections, while motorways are controlled-access roads designed to reduce congestion. Assuming capacity to be proportional to the number of lanes might be inaccurate and could explain the surprising result given that primary roads are typically single-lane roads, while motorways usually have multiple lanes.<sup>9</sup> This could be improved in a future version by introducing a parameter

<sup>9</sup>Assuming that all observed primary roads have 1 lane and all observed motorways have 2 lanes, the model basically predicts that the capacity of a 1-lane primary road is 2150, while the total capacity of a 2-lane motorway is  $2 \times 1700 = 3400$ . If in reality doubling the lane effectively amounts to increasing the total capacity by 50%, then the total capacity of a 2-lane primary road would be  $1.5 \times 2150 = 3225$ , which is smaller than for motorways. Or, equivalently, the capacity of a 1-lane motorway would be  $3400/1.5 \approx 2267$ .

to account for the number of lanes, ensuring that doubling the number of lanes does not necessarily double the road’s capacity.

Figure 3.5 compares the congested travel times between Tomtom and METROPOLIS2, for the 200 000 OD pairs used in the calibration. The figure shows that the observed travel times are generally well replicated, although METROPOLIS2 tends to overestimate travel times for some OD pairs with travel time exceeding thirty minutes.




Figure 3.5: TomTom vs calibrated congested travel times at the OD-level

### 3.5.3 Departure-Time Distribution Calibration

The third step of the calibration methodology focuses on calibrating the departure-time distribution for car driver trips. The observed departure-time distribution is derived from the regional travel survey data.

The goal is not only to match the global departure-time distribution but also the distributions specific to certain trip purposes or origin/destination areas. To achieve this, we cluster the survey trips into 18 groups, each with approximately 2000 trips, based on the preceding and following activity purpose, origin and destination departments and Euclidean

distance. These clusters are created using the  $k$ -means algorithm (Macqueen, 1967) and all have between 1100 and 3500 trips. Figure 3.6 illustrates the departure-time distribution for the 18 clusters, showing significant variability across clusters.

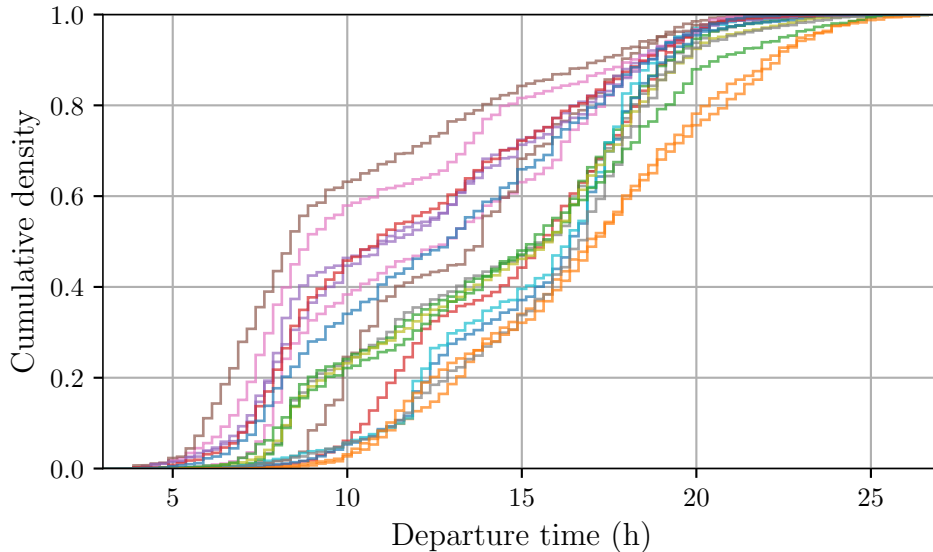


Figure 3.6: Cumulative distribution of departure time, for the 18 clusters

We aim to closely match the observed departure-time distribution in each cluster by solving the following optimization problem:

$$\operatorname{argmin}_{\theta} \sum_m w_m \cdot d_m(\theta),$$

where  $\theta$  is the vector of parameters,  $w_m$  is the number of trips in cluster  $m$  and  $d_m(\theta)$  is the Cramér-von Mises distance between the observed and simulated distribution for cluster  $m$ , defined as

$$d_m(\theta) = \int_{t^0}^{t^1} [F_m(t; \theta) - F_m^*(t)]^2 d F_m^*(t),$$

where  $F_m^*(t)$  is the cumulative distribution function (CDF) of the observed departure-time distribution for cluster  $m$  and  $F_m(t; \theta)$  is the CDF of the simulated distribution for cluster  $m$ , given a vector of parameters  $\theta$ . The parameters included in  $\theta$  are the early and late

schedule-delay penalties for each purpose, relative to the value of time (which is normalized for this step).

To compute  $F_m(t; \theta)$  for any cluster  $m$  and any parameter vector  $\theta$ , we use METROPOLIS2's departure-time choice model. Since road congestion was calibrated in the previous step, the road-network conditions can be treated as exogenous, allowing us to avoid rerunning full simulations for any vector of parameters  $\theta$  to evaluate. This reduces computation time significantly, enabling the evaluation of  $F_m(t; \theta)$  in just a few seconds.

Despite this efficiency, the objective function remains expensive to evaluate and gradient information cannot be obtained, which makes standard optimization algorithms like Newton-Raphson or Powell's methods unfeasible. Instead, we employ Bayesian Optimization with Gaussian Processes, an approach that is adapted to optimize a complicated function that is costly to evaluate by approximating it using a Gaussian process. The calibration follows these steps:

1. Evaluate the objective function for 10 random vectors of parameters using METROPOLIS2.
2. Approximate the objective function using a Gaussian process based on all previous evaluations.
3. Select the next parameter vector by minimizing a function over the Gaussian process prior, balancing exploration (evaluating regions of high uncertainty) and exploitation (focusing on regions where the objective function is predicted to be low).
4. Evaluate the objective function for the new parameter vector using METROPOLIS2.
5. Repeat from Step 2, stopping after 300 evaluations. The calibrated parameter vector is the vector with the smallest objective function value among the 300 vectors evaluated.

Table 3.3 shows the calibrated schedule-delay parameters.

Note that the usual assumption that  $\beta < \gamma$  is not satisfied for all purposes. This might come from the assumption that activity duration is exogenous, which implies that an early arrival at destination cannot be disentangled from an early departure from origin for the

Table 3.3: Calibrated parameter values for departure-time choice

Purpose	$\beta/\alpha$	$\gamma/\alpha$
Work	0.61	1.16
Education	0.00	0.64
Shop	0.85	0.63
Leisure	0.25	0.03
Other	0.66	1.08

next trip, so  $\beta$  and  $\gamma$  measure schedule-delay for both the start and the end of the activities.

Figure 3.7 illustrates the overall match between the observed and simulated departure-time distributions after calibration, while Figure 3.8 shows a comparison for the first 4 clusters. Despite some discrepancies, the overall shapes of the distributions are well replicated.




Figure 3.7: Comparison between the overall departure-time distribution in the travel survey and in METROPOLIS2 after calibration

### 3.5.4 Mode Shares Calibration

In the previous calibrations steps, mode choices were held constant. In this step, we focus on calibrating the parameters governing mode choice – such as utility constants and values




Figure 3.8: Comparison between the observed and simulated departure-time distribution for 4 clusters

Note. Cluster 1 (top left) corresponds mostly to home-to-work trips from the outer ring to either Paris or the inner ring. Cluster 2 (top right) corresponds mostly to trips with *other* activity purpose at origin, with both origin and destination in the *Seine-et-Marne*, *Yvelines* or *Val d'Oise* department. Cluster 3 (bottom left) corresponds mostly to trips with *home* activity purpose at origin, with both origin and destination in the *Seine-et-Marne* department. Cluster 4 (bottom right) corresponds mostly to trips with *home* activity purpose at origin, with both origin and destination in the *Seine-Saint-Denis* department.

of time – to closely replicate the observed mode shares from real-world data.

The survey trips are clustered into 20 groups, of around 5000 observations, following the same methodology described in the previous section, but this time including all trips regardless of mode. The goal is to replicate the observed mode shares within each cluster. The objective function of this calibration step is the Jensen-Shannon distance between the observed mode frequencies and the simulated mode probabilities, averaged over clusters and weighted by cluster size. For cluster  $m$ , the Jensen-Shannon distance,  $d_m(\mathbf{p}_m^*, \mathbf{p}_m)$ , between observed mode probabilities,  $\mathbf{p}_m^* = \{p_{m,j}^*\}_{j \in \mathcal{J}}$ , and simulated mode probabilities,



$\mathbf{p}_m = \{p_{m,j}\}_{j \in \mathcal{J}}$ , is the square root of the Jensen-Shannon divergence:

$$d_m(\mathbf{p}_m^*, \mathbf{p}_m) = \sqrt{\frac{1}{2} \sum_{j \in \mathcal{J}} p_{m,j}^* \cdot \ln(p_{m,j}^*/p_{m,j}) + p_{m,j} \cdot \ln(p_{m,j}/p_{m,j}^*)},$$

with  $\mathcal{J}$  the set of modes. The Jensen-Shannon distance is adapted to measure the similarity between two probability distributions. The objective function is thus

$$\sum_m w_m \cdot d_m(\mathbf{p}_m^*, \mathbf{p}_m),$$

where  $w_m$  is the size of cluster  $m$ .

The simulated mode probabilities are computed using METROPOLIS2 with a Multinomial Logit model (see Section 3.3). As in the previous step, this computation is less computationally intensive than running a full simulation since congestion levels are treated as exogenous. However, calculating mode choice for millions of trips remains computationally demanding, making standard optimization methods impractical. To address this, we adapt the optimization process used in the previous calibration step, with modifications to account for the discrete nature of mode choice (as opposed to the continuous departure-time choice).

The calibration process proceeds as follows:

1. Evaluate the mode probabilities using METROPOLIS2 for 20 random vectors of parameters.
2. Approximate the mode probabilities as a function of the parameters using random forest regression based on the previous evaluations.
3. Use the L-BFGS-B optimization algorithm to find the parameter vector of parameters minimizing the Jensen-Shannon distance between the observed probabilities and those predicted by the random forest model.
4. Evaluate the mode probabilities with METROPOLIS2 for that parameter vector.
5. Repeat from Step 2, stopping after 500 evaluations of the objective function. The

calibrated parameter vector is the vector with the smallest distance among the 500 vectors evaluated.

Random forest regressions are employed to estimate mode probabilities, as they conform to the constraint that probabilities must sum to one while effectively handling non-linear effects.

To account for population heterogeneity, we allow the parameters to vary based on gender (male or female) and socio-professional class (blue-collar workers, employees, intermediate workers, upper-category workers, retirees and students). This results in twelve subpopulations whose parameters can be calibrated independently. Each subpopulation requires nine parameters to be calibrated: one utility constant per mode (except walking, where it is normalized to zero), one value of time for each mode (except car driver, where it is normalized to one) and a fuel factor representing how consuming an additional liter of fuel affects utility.

Tables 3.4, 3.5 and 3.6 present the calibrated values for each subpopulation after 500 evaluations of the objective function.

The mode-specific constants are shown in Table 3.4. Except for walking (normalized to zero), all constants are negative, indicating that, as the trip distance approaches zero, walking becomes the preferred mode. This is intuitive, as other modes involve extra efforts (e.g., getting the car or the bicycle out of the garage, checking public-transit schedules). Note that there was no constraint in the model enforcing the constants to be negative.

All constants range between 487 and 4770 in absolute value. Given that the car value of time is normalized to 1 utility unit per second, the constants are equivalent to between 8 and 80 minutes of car travel time. For example, for male employees, the utility loss of taking the car instead of walking is equivalent to the utility loss of spending 1192 seconds, or about 20 minutes, in a car.

Recall that the car alternative is only available to agents whose households own at least one vehicle, meaning that the fixed costs of car ownership (e.g., purchase price, insurance) are not included in the utility constant. In contrast, the bicycle mode is available to all

Table 3.4: Calibrated mode utility constants

	Car driver	Car passenger	Public transit	Bicycle	Walking
Male / blue-collar	-1094	-2673	-3462	-2217	0
Female / blue-collar	-1051	-1328	-2379	-4081	0
Male / employee	-1192	-2408	-2116	-1390	0
Female / employee	-534	-540	-1241	-1722	0
Male / intermediate cat.	-1242	-1711	-2801	-1263	0
Female / intermediate cat.	-1707	-1801	-2909	-2389	0
Male / upper cat.	-1591	-2506	-1190	-1544	0
Female / upper cat.	-1654	-3470	-2177	-1491	0
Male / retiree	-1134	-1058	-1250	-1777	0
Female / retiree	-487	-734	-1089	-4770	0
Male / student	-2184	-2038	-2146	-2902	0
Female / student	-2653	-2192	-2210	-2164	0

agents (as we do not model bicycle ownership), meaning its utility constant may indirectly reflect the cost of bicycle ownership. Public transit is also available to all agents (except for tours with trips that cannot be completed by public transit), and the utility constant for this mode may indirectly account for ticket costs or transit subscriptions. All this can explain why the car driver constant is the second-largest (least negative) in nearly all categories, except for male and female students, male and female upper-category workers and male retirees. The car driver constant being smaller for students can reflect their lower likelihood of using a car, even if their household owns one, as parents may prioritize the car for their own use (remember that the constraint that the cars cannot be used simultaneously by different household members is not enforced in the model).

Interestingly, the bicycle utility constant is larger (less negative) for men than for women across all socio-professional classes, except upper-category workers and students, which may reflect gender-based preferences for cycling (see e.g., Mitra and Nash 2019).

Table 3.5 reports the mode-specific values of time for each category. The lowest values of time is always associated to public-transit (except for male upper-category workers), with values ranging between 0.11 and 1.01 (where the car driver values are normalized to 1). This suggests that agents prefer spending time in public transit over other modes, likely due to

the possibility of multitasking during transit (e.g., reading or working). However, this could be refined by accounting for factors such as transfer inconveniences or in-vehicle congestion in future versions of the model.

Table 3.5: Calibrated mode-specific values of time

	Car driver	Car passenger	Public transit	Bicycle	Walking
Male / blue-collar	1.00	1.67	0.14	1.38	1.77
Female / blue-collar	1.00	1.39	0.11	1.16	1.50
Male / employee	1.00	1.27	0.40	1.42	1.25
Female / employee	1.00	1.95	0.47	1.72	0.86
Male / intermediate cat.	1.00	1.78	0.28	1.86	1.45
Female / intermediate cat.	1.00	1.64	0.25	0.96	1.47
Male / upper cat.	1.00	1.76	1.01	1.17	1.15
Female / upper cat.	1.00	1.03	0.47	1.51	1.38
Male / retiree	1.00	1.91	0.88	0.92	1.58
Female / retiree	1.00	1.05	0.38	1.28	0.87
Male / student	1.00	1.42	0.38	0.49	1.69
Female / student	1.00	1.84	0.58	1.22	1.91

Table 3.6 shows the fuel factor parameter for each category, representing the utility penalty per additional liter of fuel consumed in the car driver mode. Given that the car value of time is normalized to 1 utility unit per second, and assuming a fuel price of 1.8€/L, the estimated fuel factor parameters allow to compute implied car values of time expressed in €/h. These values range from 6.66€/h, for female blue-collar workers, to 36.00€/h, for male students. One needs to be careful when interpreting the values since the correlation to socio-demographic characteristics is not considered when drawing the car characteristics for the synthetic population, which in turn determine the fuel consumption. Additionally, the values obtained for male and female students might be less reliable due to the low share of trips as car driver in both categories (about 3%).

Figure 3.9 shows a comparison of observed and simulated mode shares over clusters. The average Jensen-Shannon distance between mode shares over clusters, when aggregating all genders and socio-professional classes, is 8.90% after calibration.

Figure 3.10 shows the observed and simulated mode shares, aggregated over all clusters

Table 3.6: Calibrated values for the fuel factor and implicit car value of time

	Fuel factor	Car value of time (€/h)
Male / blue-collar	973	7.11
Female / blue-collar	622	6.66
Male / employee	461	10.42
Female / employee	965	14.06
Male / intermediate cat.	742	8.73
Female / intermediate cat.	415	15.61
Male / upper cat.	923	7.02
Female / upper cat.	852	7.61
Male / retiree	540	12.00
Female / retiree	893	7.26
Male / student	180	36.00
Female / student	719	9.01




Figure 3.9: Comparison between observed and simulated mode shares over all clusters  
 Note. The Pearson correlation coefficient between observed and simulated mode shares is 96.50%.

and categories.




Figure 3.10: Comparison between global observed and simulated mode shares

## 3.6 Low Emission Zone Policy

The methodology outlined so far allows us to build and calibrate a baseline simulation for the Île-de-France region, replicating a typical working day under current conditions. This baseline can then be compared to counterfactual simulations to evaluate the impact of specific policies or exogenous shocks. In this chapter, we evaluate the Low Emission Zone (LEZ) policy being implemented in the Greater Paris area.

The Paris LEZ restricts entry for the most polluting vehicles in an area covering Paris and 76 surrounding municipalities. As shown on Figure 3.11, the LEZ is bordered by the A86 highway, providing an alternative route around the zone. The LEZ covers an area of 367 km<sup>2</sup>, which accounts for only 3.04 % of the Île-de-France region. However, approximately 40 % of the population resides within the LEZ, and 49 % of all trips either originate or terminate inside the LEZ. Additionally, roads within the LEZ represent 9 % of the total road network of Île-de-France.

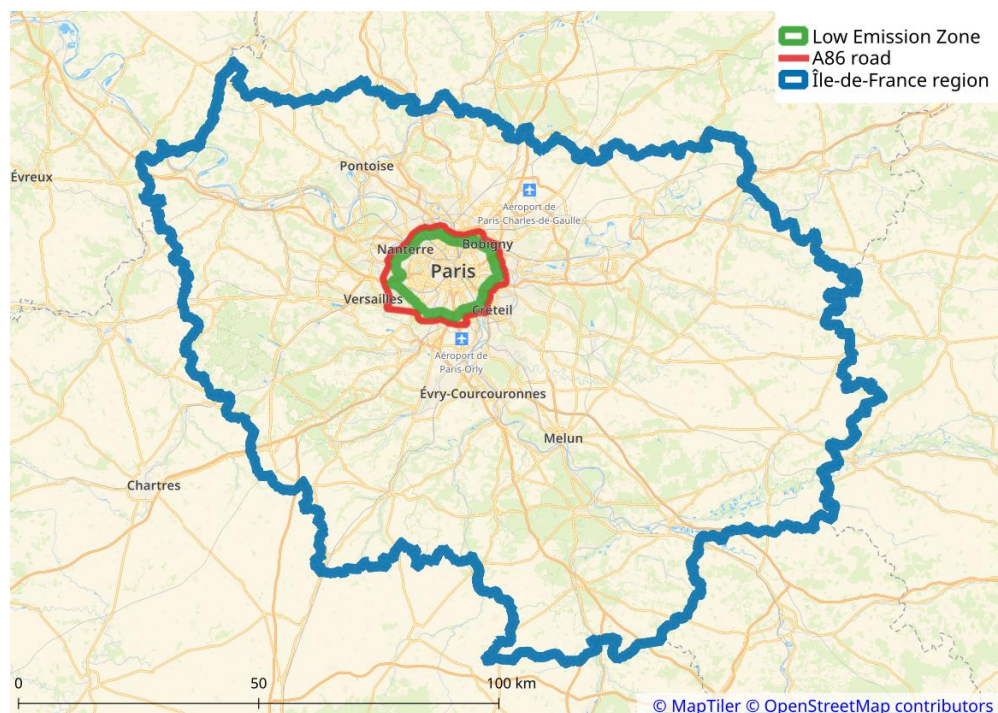


Figure 3.11: Area of the Low Emission Zone within Île-de-France

### 3.6.1 Crit’Air Classification System

Entry restrictions for the LEZ are based on the *Crit’Air* classification system, which categorizes vehicles into six groups depending on engine type (electric, hydrogen, petrol or diesel) and European emission standards. Figure 3.12 shows the evolution of Crit’Air shares within the Île-de-France vehicle fleet between 2011 and 2025, with the last three years being interpolated.

As of June 2021, vehicles in the lowest two Crit’Air categories (4 and 5) as well as unclassified vehicles are banned from entering the LEZ.<sup>10</sup> These vehicles represent 5.4 % of the 2025 Île-de-France vehicle fleet (interpolated). In January 2025, the ban will be extended to Crit’Air 3 vehicles, raising the total share of banned vehicles to 20.8 %. Table 3.7 summarizes the Crit’Air categories, their fleet shares and the dates of the corresponding bans.

<sup>10</sup>We decided to define the baseline situation as the situation with no LEZ, even though a better alternative might have been the June 2021 LEZ situation. This choice is justified by the following reasons: (i) it is simpler to have no LEZ constraint during the calibration process; (ii) the regional travel survey was carried out before the LEZ was implemented; and (iii) it allows interpreting the results as the impact of the LEZ, not the impact of extending the LEZ to Crit’Air 3 vehicles.




Figure 3.12: Evolution of Crit'Air category shares in the Île-de-France vehicle fleet

Note. Values after 2022 are interpolated using a basic Markov Chain model based on the preceding three years.

Source: French Ministry of Ecology.

Table 3.7: Crit'Air categories

Category	Vehicles	Share	Cum. share	Ban date
Unclassified	N/A	1.4 %	1.4 %	July 2019
Crit'Air 5	Diesel cars registered before 2000	0.6 %	2.0 %	July 2019
Crit'Air 4	Diesel cars registered between 2001 and 2005	3.4 %	5.4 %	June 2021
Crit'Air 3	Petrol cars registered before 2005 or diesel cars registered between 2006 and 2010	15.3 %	20.8 %	January 2025
Crit'Air 2	Petrol cars registered between 2006 and 2010 or diesel cars registered after 2011	31.2 %	51.9 %	N/A
Crit'Air 1	Rechargeable gas and hybrid vehicles or petrol cars registered after 2011	45.3 %	97.2 %	N/A
Crit'Air E	Electric and hydrogen	2.8 %	100.0 %	N/A

Note. Shares and cumulative shares are computed from the interpolated 2025 Île-de-France vehicle fleet, using data from the French Ministry of Ecology.



Figure 3.13 shows the spatial distribution of banned vehicle shares for the January 2025 policy. The highest shares of banned vehicles are found in municipalities outside the LEZ, especially in the north of Paris, while the lowest shares are in municipalities in the west and southwest of Paris – inside and outside the LEZ. Interestingly, these shares are negatively correlated with the municipality’s income level, with a Pearson correlation coefficient of  $-29.7\%$ .




Figure 3.13: Share of cars with Crit’Air 3 or worse at the municipality level for the 2025 vehicle fleet

### 3.6.2 Simulating the LEZ

To simulate the LEZ, we use METROPOLIS2’s road restriction feature, defining two vehicle types: authorized vehicles (Crit’Air 2 or better), which have unrestricted access to all roads, and banned vehicles (Crit’Air 3 or worse), which are prohibited from accessing any road within the LEZ. Agents are assigned to these vehicle types based on the car characteristics

generated in the synthetic population (see Section 3.4.2).<sup>11</sup> Owners of banned vehicles may be required to reroute around the LEZ to reach their destination.

For trips that either start or end inside the LEZ, owners of banned vehicles lose access to the car driver or car passenger modes, as the trip becomes infeasible. In total, we identified 91 270 agents (0.88 % of the population) who no longer have access to a car, with at least one trip which cannot be completed by public transit and requires more than 1 h of walking (or over 24 min by bicycle). Given that it is not practically feasible for these “trapped” agents to avoid car use, we assume they can still travel within the LEZ by car, despite the restriction. This assumption accounts for possible exceptions, such as exemptions for disabled individuals, classic cars, or instances where some drivers may disregard the LEZ regulations.

### 3.6.3 Air Quality Impact

Improving air quality is the primary goal of the LEZ. To evaluate its effectiveness, we assess road-traffic emissions and the related health impacts using the methodology developed by Le Frioux et al. (2024) for METROPOLIS2. This methodology computes emissions of nitrogen oxides ( $\text{NO}_x$ ) and particulate matter ( $\text{PM}_{2.5}$ ) due to road traffic, and estimates population exposure to these pollutants. The process includes five steps. Road-traffic *emissions* are calculated based on vehicle characteristics and simulated speeds at the road level, using the EMISENS model (Ho et al., 2014). These emissions are *dispersed* in the atmosphere using a Gaussian plume dispersion model, which generates average daily pollutant concentrations on a grid of resolution  $500 \text{ m} \times 500 \text{ m}$ , assuming a west-to-east wind with speed 10 km/h. *Population density* is computed from METROPOLIS2’s output, showing the location of each agent over the simulation day. Then, *exposure* can be calculated by multiplying pollutant concentrations with population density on the grid. Finally, the *health costs* of pollution exposure are obtained by first computing the relative risks due to pollution

---

<sup>11</sup>Since vehicle ownership is defined at the household level, we assign a random vehicle from the household fleet to each agent.

exposure (the increase in mortality that can be attributed to pollution), then multiplying the result with the mortality rate (1%), the average number of years lost when dying from air pollution (10.4 years), and the price of a year of life lost (106 985 €). See Le Frioux et al. (2024) for additional details.

This methodology allows us to compute exposure at the agent level, which is critical for evaluating spatial and economic disparities resulting from the LEZ. Additionally, CO<sub>2</sub> emissions are computed from the EMISENS model (Step 1).

### 3.6.4 Expected Results

The primary expected outcome of the LEZ is a reduction in air pollutants within the zone. Additional benefits include reducing CO<sub>2</sub> emissions, road congestion and noise pollution.<sup>12</sup> The most direct expected effect of the LEZ is that the most polluting vehicles will no longer circulate within the LEZ, leading to a reduction in both air pollutants and road congestion within the area. Owners of banned vehicles who used to drive in the LEZ have several potential responses: (i) buying an authorized vehicle, (ii) switching to another transport mode (e.g., public transit), or (iii) rerouting around the LEZ, provided that both their origin and destination are outside the zone. The first option, buying a new vehicle, is typically not feasible in the short term.

Second order effects may also arise. For instance, the number of vehicles driving in the LEZ could increase if owners of authorized vehicles change their routes or modes due to decreased congestion (rebound effect). Additionally, air pollutant emissions and road congestion may increase in areas surrounding the LEZ if many drivers choose to reroute around the zone. Verifying these predictions and measuring the magnitude of the effects require running METROPOLIS2 simulations and comparing the results with and without the LEZ.

---

<sup>12</sup>The analysis of the impact of LEZs on noise pollution, albeit important, is outside of the scope of this chapter.

### 3.6.5 Limits of the Methodology

Our evaluation focuses on the short-term impacts of the LEZ, where agents can adjust their mode, departure time, and route, but other variables such as car ownership, activity plans, and residential locations are assumed to remain fixed. In the medium term, agents could respond to the LEZ by purchasing authorized vehicles, changing shopping or other activity locations, or moving closer to public transit, among other adjustments.

Additionally, the study assumes that agents have not anticipated the LEZ policy. This is reflected in the assumption that the vehicle fleet follows the same evolution trend between 2020–2022 and 2023–2025. To account for such medium-term adjustments, more complex models would be required, including car ownership models, activity-based models, and land-use models, all of which fall outside the scope of this chapter.

Another potential adjustment not modeled in this study is the use of park-and-ride strategies, where agents drive their unauthorized vehicles to the LEZ boundaries, park near a train station, and complete their trip using public transit. Such behaviors are unlikely to represent a significant response to the policy, as they represent only 1.4% of the trips in the region according to the travel survey (0.8% when excluding trips by car passengers). In a related case study, Yin et al. (2024) find that park-and-ride trips would not increase substantially when implementing a driving restriction zone in Paris. Nevertheless, park-and-ride trips could be incorporated into the model in future research, as long as the locations of the park-and-ride facilities chosen by agents are well defined (e.g., selecting the facility closest to home, as in Yin et al. 2024).

Despite these limitations, the results provide valuable insights into the short-term effects of the LEZ, representing upper-bound estimates of the policy’s impact if agents do not anticipate it.

Furthermore, the LEZ is modeled as an absolute restriction, with banned vehicles prohibited from entering at any time within the LEZ, while, in reality, the LEZ operates only between 8 a.m. and 8 p.m., on weekdays. There are also several exceptions to the LEZ.

Exemptions exist for disabled individuals, classic cars, and certain professional vehicles. Additionally, enforcement is currently not fully automated, and some individuals may evade compliance. Owners of Crit'Air 3 vehicles are also allowed 12 exceptions per year, granting them temporary access to the LEZ for up to 24 hours. All these cases are not accounted for in the simulations.

## 3.7 Results

We run two simulations of METROPOLIS2, one for the baseline scenario and another for the LEZ scenario, each running for 100 iterations. These iterations are necessary for road congestion and travel decisions to adjust to each other, allowing the simulator to approximate an equilibrium.

### 3.7.1 Convergence

Before interpreting the results of the simulations, we ensure that the convergence of the simulations is satisfactory. Table 3.8 presents statistics related to the convergence of the two simulations.

Both simulations exhibit strong stability in travel decisions, with around only 0.01% of tours experiencing a mode change compared to the previous iteration, with shifts in the selected departure time averaging under 3 seconds, and with road trips taking roads that were not taken during the previous iteration for around 100 meters on average. This suggests that travel decisions and thus road congestion remains stable across iterations.

The results also show that agents misanticipate the travel times they face by about 45 seconds on average, or approximately 3.5% of their travel time, indicating a reasonably accurate forecast of road congestion when making travel decisions. For further details on the convergence behavior of METROPOLIS2, see Chapter 2.

Table 3.8: Statistics on the convergence of the simulations

	Baseline	LEZ
Running time <sup>a</sup>	47:59:52	40:13:05
Share of mode shifts (tour level) <sup>b</sup>	0.01 %	0.01 %
Mean absolute departure-time shift (tour level) <sup>c</sup>	2.6 s	1.7 s
Mean route length shift (road trip level) <sup>d</sup>	110 m	86 m
Mean absolute error of travel-time anticipation (road trip level) <sup>e</sup>	43.8 s	44.3 s
Mean relative error of travel-time anticipation (road trip level) <sup>f</sup>	3.53 %	3.48 %

<sup>a</sup> Simulations are run on a machine with a 8-core Intel Xeon CPU (3.20GHz) and with 128GiB of RAM.

<sup>b</sup> Share of tours with a mode change compared to the previous iteration.

<sup>c</sup> Mean absolute difference in selected departure times compared to the previous iteration, for tours where the selected mode remains the same.

<sup>d</sup> Mean length of the part of the route that was not taken during the previous iteration, over all road trips.

<sup>e</sup> Mean absolute difference between the expected and actual travel time, over all road trips.

<sup>f</sup> Mean of the absolute between the expected and actual travel time, divided by the expected travel time, over all road trips.

### 3.7.2 Global results

Before comparing the results between the baseline and LEZ simulations, we first aim to understand how many agents are directly impacted by the LEZ by examining the characteristics of their tours. Table 3.9 summarizes the share of tours based on car ownership status (authorized car, banned car, or no car) and specific tour characteristics, such as whether a trip occurs inside the LEZ or if public transit is available.

Table 3.9: Tour characteristics by car ownership

	Authorized car	Banned car	No car
Overall	58 %	15.9 %	26.1 %
Agent has driving license (yes / no)	72.4 % / 27.6 %	71.5 % / 28.5 %	49.5 % / 50.5 %
Any trip start / end inside LEZ (yes / no)	48.3 % / 51.7 %	43.1 % / 56.9 %	73.4 % / 26.6 %
Selected mode in baseline scenario (car / other)	48.1 % / 51.9 %	48.9 % / 51.1 %	0.0 % / 100.0 %
Public-transit access (yes / no)	87.6 % / 12.4 %	87.5 % / 12.5 %	93.4 % / 6.6 %

Although 20.8 % of the vehicle fleet is classified as banned under the LEZ, the share of tours directly impacted by the policy is significantly smaller. Specifically, 26.1 % of all tours

are conducted by agents without a car, leaving 15.9% of tours conducted by agents who own a banned vehicle. Nevertheless, not all these 15.9% of tours are directly affected by the LEZ. For example, only 71.5% of them are performed by agents with a driving license, meaning the remaining 28.5% of tours are performed by agents who do not have the option of driving. Furthermore, 43.1% of the tours conducted by banned car owners have a trip that actually starts or ends inside the LEZ. For the remaining 56.9% of tours, the banned car can still be used because the tour does not require entering the LEZ, although detours may be required. In the baseline scenario, fewer than half of the tours conducted by banned car owners are made using their car (48.9%), while the rest (51.1%) are done using alternative modes. This suggests that, in the LEZ scenario, even if agents could still use their banned car, they would often choose not to. Additionally, public transit access is relatively high for tours involving banned cars, with only 12.5% of them having no feasible public transit option for at least part of the trip.

A comparison of the aggregate results for the two simulations is presented in Table 3.10. The travel surplus of agents, representing the utility they expected to derive from their trips, is computed using the logsum formula, described in Equation (3.4). This travel surplus is expressed in utility units and subsequently converted into euros using the fuel factor parameters estimated during calibration. This conversion depends heavily on the mode choice calibration so results should be interpreted with caution (see Section 3.5.4). Fortunately, the sign of the travel surplus variation remains meaningful regardless of the unit: an increase in travel surplus indicates that the agent is better off under the LEZ scenario, and vice versa. Additionally, agents with the same socio-demographic characteristics share the same fuel factor, enabling direct comparison of their travel surplus.

The results show that, on average, the agent travel surplus decreases by 0.13€ when the LEZ is implemented, while the average daily travel time increases by 1 minute and 55 seconds. These outcomes reflect the added constraint of the LEZ, which forces some agents to reroute or switch to modes that provide lower utility.

Table 3.10: Measures of effectiveness for the simulated average day

	Baseline	LEZ	Variation	Observed value
<i>Global output (agent level)</i>				
Average travel surplus	-28.81 €	-28.94 €	-0.13 €	
Average daily travel time	01:09:14	01:11:09	+115 s	01:36:13 (2010) <sup>a</sup>
<i>Mode shares (tour-level)</i>				
Car driver share	31.1 %	29.3 %	-1.8 p.p.	30.7 % (2010) <sup>a</sup>
Car passenger share	5.5 %	5.4 %	-0.1 p.p.	8.8 % (2010) <sup>a</sup>
Public transit share	18.4 %	19.7 %	+1.3 p.p.	24.6 % (2010) <sup>a</sup>
Bicycle share	1.1 %	1.4 %	+0.3 p.p.	1.9 % (2010) <sup>a</sup>
Walking share	43.9 %	44.2 %	+0.3 p.p.	33.9 % (2010) <sup>a</sup>
<i>Mode shares (weighted by Euclidean distance)</i>				
Car driver share	53.3 %	51.2 %	-2.1 p.p.	47.2 % (2010) <sup>a</sup>
Car passenger share	7.8 %	7.7 %	-0.1 p.p.	7.5 % (2010) <sup>a</sup>
Public transit share	33.1 %	34.9 %	+1.8 p.p.	41.6 % (2010) <sup>a</sup>
Bicycle share	1.0 %	1.2 %	+0.2 p.p.	0.8 % (2010) <sup>a</sup>
Walking share	4.8 %	4.9 %	+0.1 p.p.	3.0 % (2010) <sup>a</sup>
<i>Road-traffic output (excluding truck trips)</i>				
Travel time (10 <sup>3</sup> hours)	3502	3307	-5.6 %	5153 (2010) <sup>a</sup>
Time lost to congestion (10 <sup>3</sup> hours)	379	348	-8.2 %	
Vehicle-kilometers (10 <sup>6</sup> km)	126.28	121.40	-3.9 %	
Passenger-kilometers (10 <sup>6</sup> km)	144.59	139.49	-3.5 %	

Note. All results are for an average working day.

<sup>a</sup> Source: Regional travel survey *Enquête Global Transport*.

The LEZ policy leads to a notable reduction in car usage, with the share of tours performed by car (as either driver or passenger) dropping by 1.9 p.p.. Almost two-thirds of this shift is absorbed by public transit (+1.3 p.p.), while the remaining portion is distributed between bicycle and walking tours (+0.3 p.p. each). This trend holds when mode shares are weighted by Euclidean distance.

The decrease in car usage translates into a significant reduction in road usage: total travel time spent on roads by car drivers and passengers decreases by 5.6 %, time lost to congestion drops by 8.2 %, and vehicle-kilometers fall by 3.9 %.

While the aggregate mode share analysis shows a shift away from car use, it is important to investigate the possibility of a rebound effect, where owners of authorized vehicles



switch back to using cars due to reduced congestion. To assess this, Table 3.11 presents the transition matrix of modes between the baseline and LEZ scenarios.

Table 3.11: Transition matrix of modes from baseline to LEZ

	Car driver	Car passenger	Public transit	Bicycle	Walking	<b>Total</b>
Car driver	29.1 %	0.0 %	1.4 %	0.2 %	0.3 %	31.1 %
Car passenger	0.0 %	5.3 %	0.1 %	0.0 %	0.0 %	5.5 %
Public transit	0.2 %	0.0 %	18.1 %	0.0 %	0.0 %	18.4 %
Bicycle	0.0 %	0.0 %	0.0 %	1.1 %	0.0 %	1.1 %
Walking	0.0 %	0.0 %	0.0 %	0.0 %	43.9 %	43.9 %
<b>Total</b>	29.3 %	5.4 %	19.7 %	1.4 %	44.2 %	100.0 %

Reading example: 0.2% of tours have mode *public transit* in the baseline scenario and mode *car driver* in the LEZ scenario.

Mode shifts from car to non-car options represent 2.2% of all tours. The most significant shifts are from car driver to public transit (1.4 p.p.). These changes are likely due to owners of banned vehicles being forced to switch to alternative modes.

Conversely, the rebound effect manifests in 0.3% of all tours switching from non-car to car options, with public transit to car driver (0.2 p.p.) representing the most important shifts. These shifts, involving roughly 48 000 tours, likely reflect the behavior of authorized vehicle owners who prefer to drive when congestion is reduced due to the LEZ.

To better understand how the LEZ impacts road congestion, we visualize the variation in daily average travel times at the road-segment level in Figure 3.14. The results indicate that congestion mostly decreases on roads near the LEZ boundaries.

On the *Boulevard Périphérique*, a heavily congested ring road surrounding Paris, the travel time at 8 a.m. decreases from 59 min to 55 min for a full loop in the clockwise direction, and from 62 min to 56 min in the counter-clockwise direction.<sup>13</sup>

In terms of public transit, Figure 3.15 shows how the daily flow of passengers varies across public transit segments when the LEZ is implemented. A public-transit segment is defined

<sup>13</sup>In October 2024, the speed limit on the *Boulevard Périphérique* decreased from 70 km/h to 50 km/h, a change that was not considered in our simulations. Evaluating the individual and joint impact of the two policies (LEZ and speed limit change) could be the topic of future research.

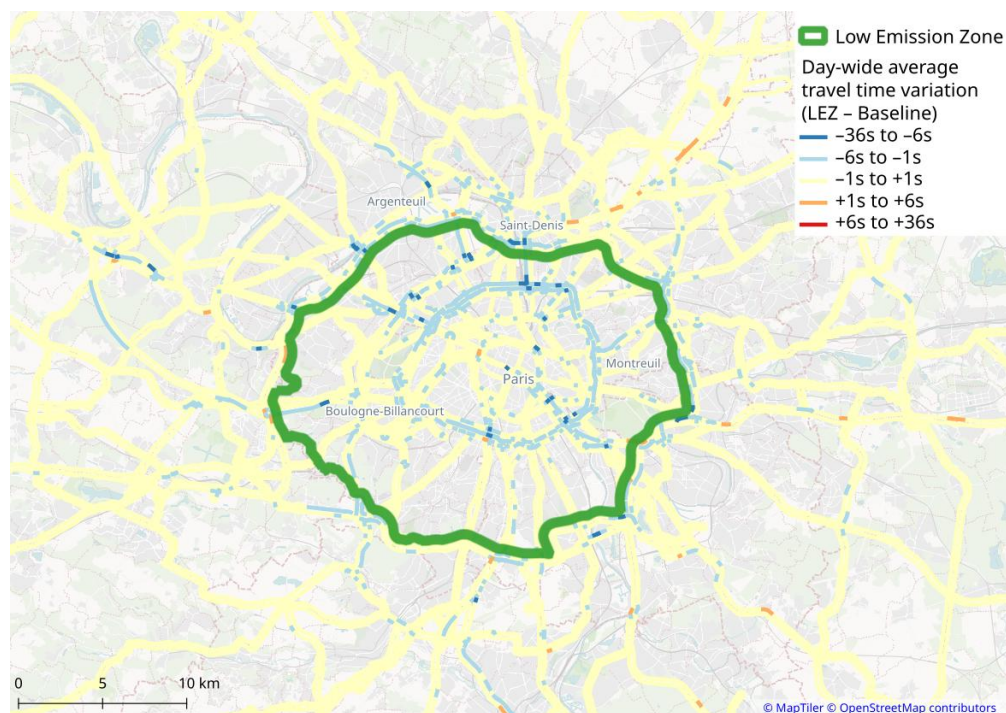


Figure 3.14: Variation in daily average travel times at the road-segment level

Note. Only major roads are represented.

Reading example: A value of  $-5s$  indicates that the average travel time over the day on the road segment is decreased by 5 seconds in the LEZ scenario compared to the baseline scenario.

as a pair of stops directly connected by some public-transit line. Passenger flow increases on almost all segments, especially in the segments within the LEZ, where passenger flows were initially the largest. This is consistent with the predicted increase in public transit mode share (from 18.4% to 19.7% in the LEZ scenario).

The two primary mass-rapid transit lines in Île-de-France, RER A and RER B, show noticeable changes in passenger flows. RER A, which runs along the west-east axis across Paris, sees a 3.2% increase in total daily passenger-kilometers, equivalent to an additional 372 thousand passenger-kilometers. RER B, operating along the north-south axis, records an ever larger relative increase of 5.4%, or 311 thousand additional passenger-kilometers.

The two public-transit lines with the largest relative increases in passenger-kilometers are tramways T7 and T10. Tramway T7 experiences a 28.2% increase, adding 14 thousand passenger-kilometers. Similarly, tramway T10 see a 21.8% rise, equivalent to an additional 13 thousand passenger-kilometers. Both lines run along the southern edge of the LEZ.

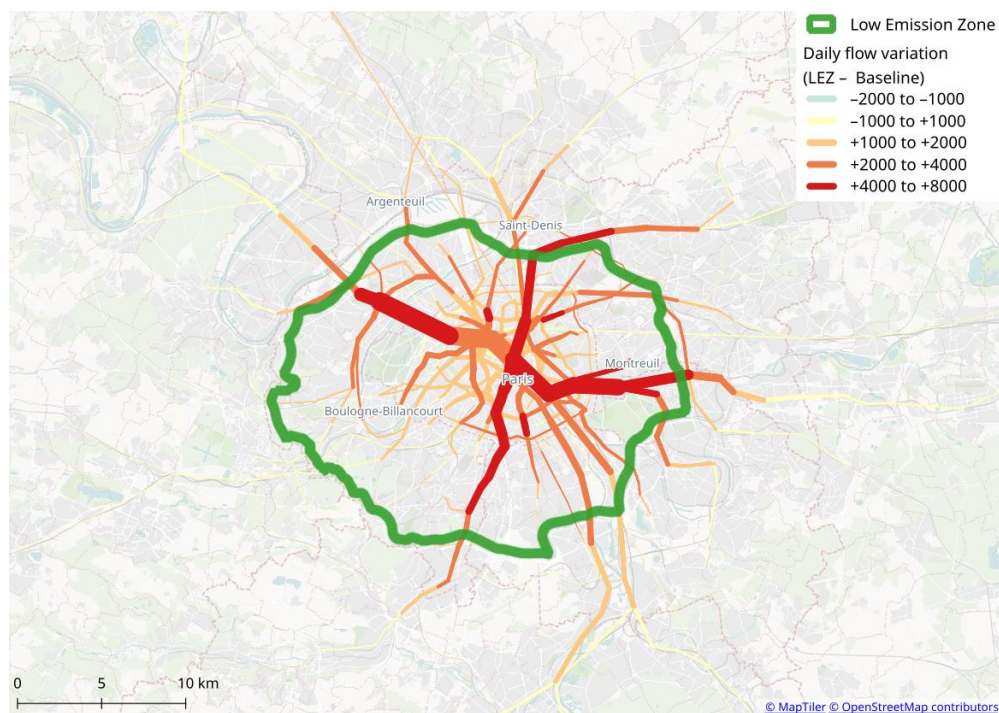


Figure 3.15: Variation of daily flows on public-transit segments

Note. The segment width is proportional to the baseline flow. Bus segments are excluded.

Reading example: A value of +2000 indicates that there are 2000 more daily passengers on that segment in the LEZ scenario compared to the baseline.

### 3.7.3 Pollution Results

Table 3.12 summarizes the results on pollutant emissions and their impact on health. While vehicle-kilometers decrease by 3.9% in the LEZ scenario, pollutant emissions show larger reductions:  $-4.5\%$  for  $\text{CO}_2$ ,  $-9.2\%$  for  $\text{NO}_x$ ,  $-7.6\%$  for  $\text{PM}_{2.5}$ . This is primarily because the reduction in vehicle-kilometers is driven by the most polluting vehicles (see below).

The health costs associated with pollution from  $\text{NO}_2$  and  $\text{PM}_{2.5}$  show even greater reductions, decreasing by  $11.6\%$  from 12.993 million euros to 11.484 million euros per day. The fact that the health costs reduction exceed the drop in emissions can be explained by air quality improving most in densely populated areas, a finding confirmed by the maps presented below. These results highlight the importance of using a detailed pollution model, like METRO-TRACE, which accounts for fleet composition, pollutant dispersion, and population exposure.

Table 3.12: Pollution-related results

	Baseline	LEZ	Variation	Observed value
CO <sub>2</sub> emissions (tonnes)	21 925	20 947	−4.5 %	21 079 (2021) <sup>ab</sup>
NO <sub>x</sub> emissions (tonnes)	33.70	30.61	−9.2 %	46.84 (2021) <sup>ac</sup>
PM <sub>2.5</sub> emissions (tonnes)	2.91	2.69	−7.6 %	2.96 (2021) <sup>ad</sup>
Premature deaths from NO <sub>2</sub>	5.46	4.92	−9.9 %	7.83 (2019) <sup>e</sup>
Premature deaths from PM <sub>2.5</sub>	6.22	5.41	−13.0 %	20.73 (2019) <sup>e</sup>
Total health surplus (10 <sup>6</sup> €)	−12.993	−11.484	−11.6 %	

Note. All results are for an average working day.

<sup>a</sup> Source: [https://data-airparif-asso.opendata.arcgis.com/datasets/73bba8b50bae442697e89b89a70191b7\\_0/explore](https://data-airparif-asso.opendata.arcgis.com/datasets/73bba8b50bae442697e89b89a70191b7_0/explore) and <https://www.citepa.org/fr/secten/>.

<sup>b</sup> In 2021, 11 500 kt of CO<sub>2</sub> were emitted by road transport in Île-de-France (Air’Parif), with private vehicles contributing 55 % (Citepa, national level). The yearly value is divided by 300 to approximate daily emissions.

<sup>c</sup> In 2021, 27 130 t of NO<sub>x</sub> were emitted by road transport in Île-de-France (Air’Parif) with private vehicles contributing 52 % (Citepa, national level). The yearly value is divided by 300 to approximate daily emissions.

<sup>d</sup> In 2021, 1400 t of PM<sub>2.5</sub> were emitted by road transport in Île-de-France (Air’Parif) with private vehicles contributing 63 % (Citepa, national level). The yearly value is divided by 300 to approximate daily emissions.

<sup>e</sup> Source: [https://www.airparif.fr/sites/default/files/document\\_publication/Rapport-Enquete-Mortalite.pdf](https://www.airparif.fr/sites/default/files/document_publication/Rapport-Enquete-Mortalite.pdf). These figures include all pollution sources, not just road transport.

Figure 3.16 confirms that the reduction in vehicle-kilometers is driven by the most polluting vehicles (Crit’Air 3 or worse), which see a decrease of around  $6.24 \times 10^6$  km (−22.5 %). In contrast, vehicle-kilometers for cleaner Crit’Air E, 1 and 2 vehicles increase by around  $1.36 \times 10^6$  km (1.4 %). This can explain why emissions drop more sharply than total vehicle-kilometers.

Next, we analyze the geographic distribution of pollution. Figure 3.17 shows PM<sub>2.5</sub> emissions for both scenarios, along with the difference between them. Emissions in both scenarios are concentrated on major highways and densely populated areas (Paris and surroundings). In the LEZ scenario, emissions decrease by 7.6 %, especially inside the LEZ and on certain highways outside it. However, some roads outside the LEZ see slightly larger emissions, likely due to rerouting from some vehicles. The maps of NO<sub>x</sub> emissions, not shown here, present very similar patterns as PM<sub>2.5</sub> emissions.

Figure 3.18 illustrates PM<sub>2.5</sub> concentrations in the atmosphere, with the highest pollution

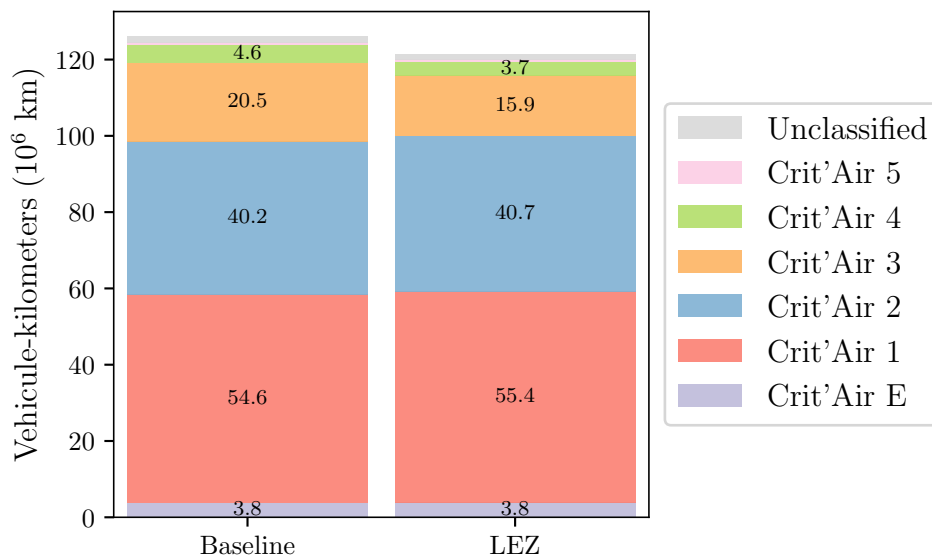


Figure 3.16: Vehicle-kilometers by Crit'Air category

levels observed in Paris and areas extending up to 30 km eastward. This pattern aligns with the assumption of an eastward wind, which is the dominant wind direction in the region.


In the LEZ scenario, air quality improves almost everywhere, especially in highly polluted areas inside and east of the LEZ. While a different wind direction would shift the areas benefiting most from improved air quality, it is unlikely to alter the relatively even distribution of health benefits across the population. Wind speed, however, could have a more significant impact, as it determines how long pollutants remain in the atmosphere. The assumption of a 10 km/h wind speed aligns well with observed data on premature deaths (see Table 3.12), justifying its use in this analysis.

The maps of  $\text{NO}_x$  concentrations, not shown here, present very similar patterns as  $\text{PM}_{2.5}$  concentrations.

Figure 3.19 shows the day-average population density (measured in person-days). It accounts for the time spent at home, the time spent doing activities, and the time spent en-route (only for car trips).<sup>14</sup> As expected, population density is highest within the LEZ,


<sup>14</sup>Individuals who do not travel are assumed to spend the entire day at home. Those who travel by public transit, bicycle or walking are “teleported” from their origin location to their destination location in the middle of the trip. The same assumption is used in Poulhès and Proulhac (2021).





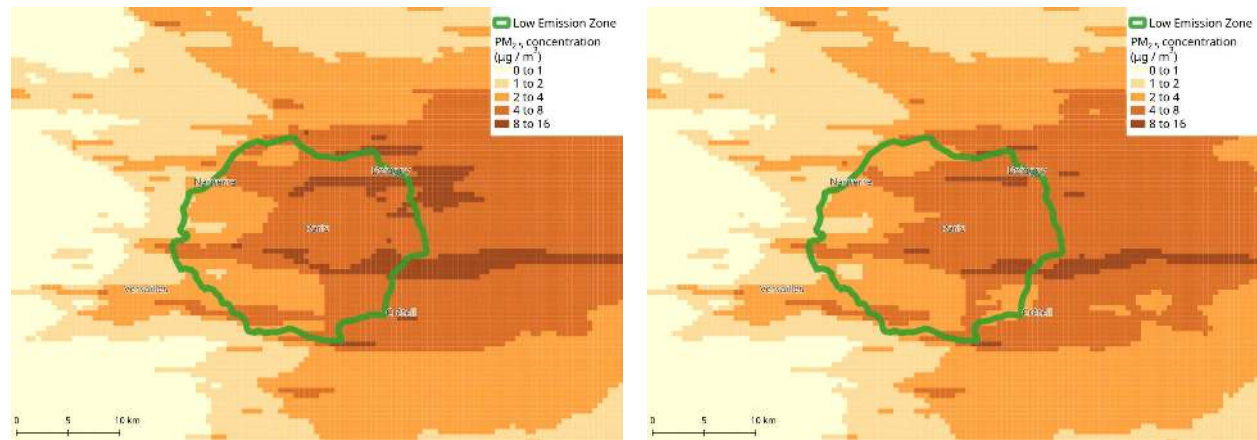
(a) Baseline scenario

(b) LEZ scenario



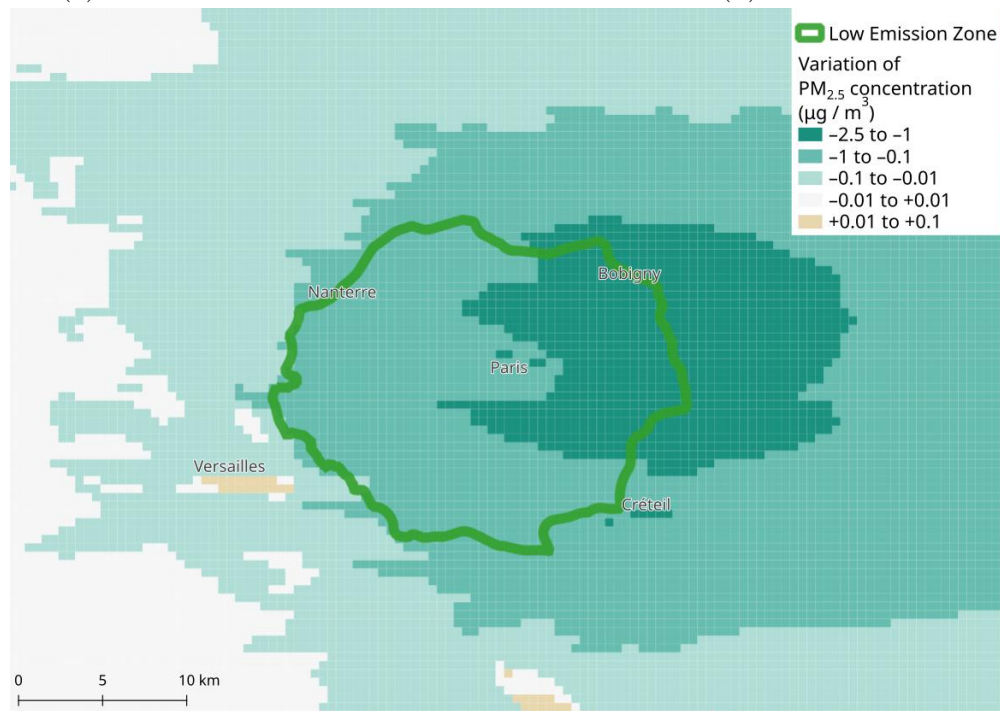
(c) Change in emissions (LEZ vs baseline)

Figure 3.17: Daily PM<sub>2.5</sub> emissions from car trips on a 500 m grid



(a) Baseline scenario

(b) LEZ scenario



(c) Change in concentration (LEZ vs baseline)

Figure 3.18: PM<sub>2.5</sub> concentrations on a 500 m grid

underscoring the importance of improving air quality in this area. The population density patterns remain almost identical between the scenarios, due to activity patterns being fixed.

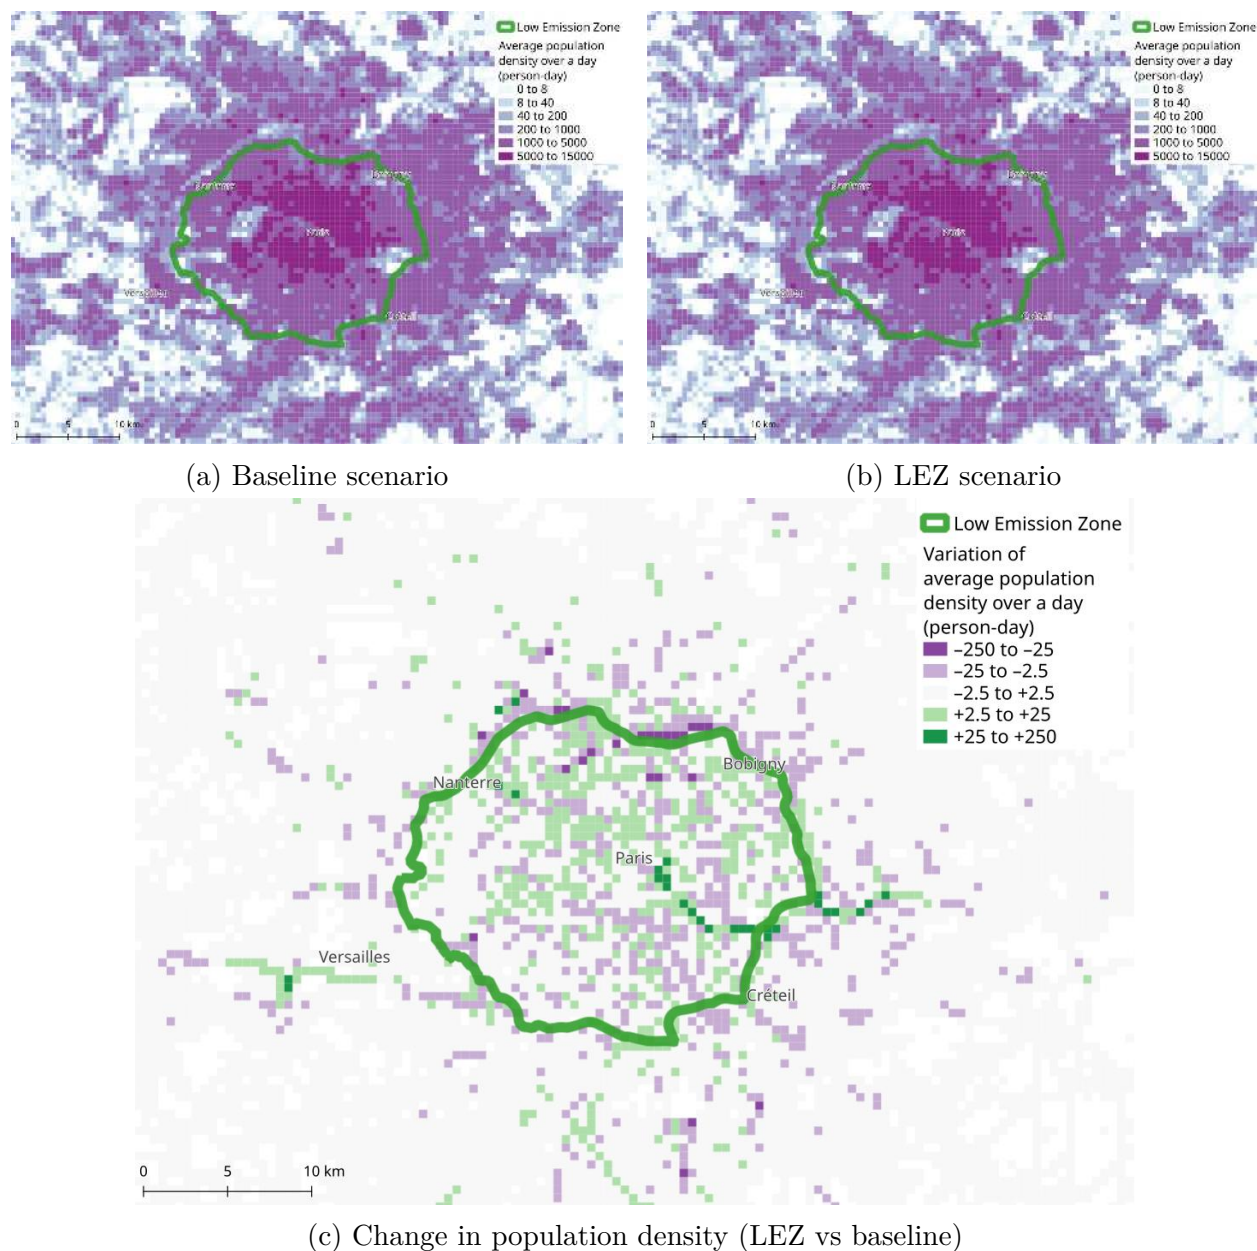


Figure 3.19: Day-average population density on a 500 m grid

Finally, Figure 3.20 maps the daily exposure cost to  $\text{NO}_2$  and  $\text{PM}_{2.5}$ , computed by combining pollutant concentrations with population density. The highest exposure costs are observed in Paris and nearby areas, where both pollution and population density are high. Exposure costs are also large in the parts of the LEZ further away from Paris and in the



east of the LEZ. The LEZ scenario shows a significant reduction in exposure costs in these areas.

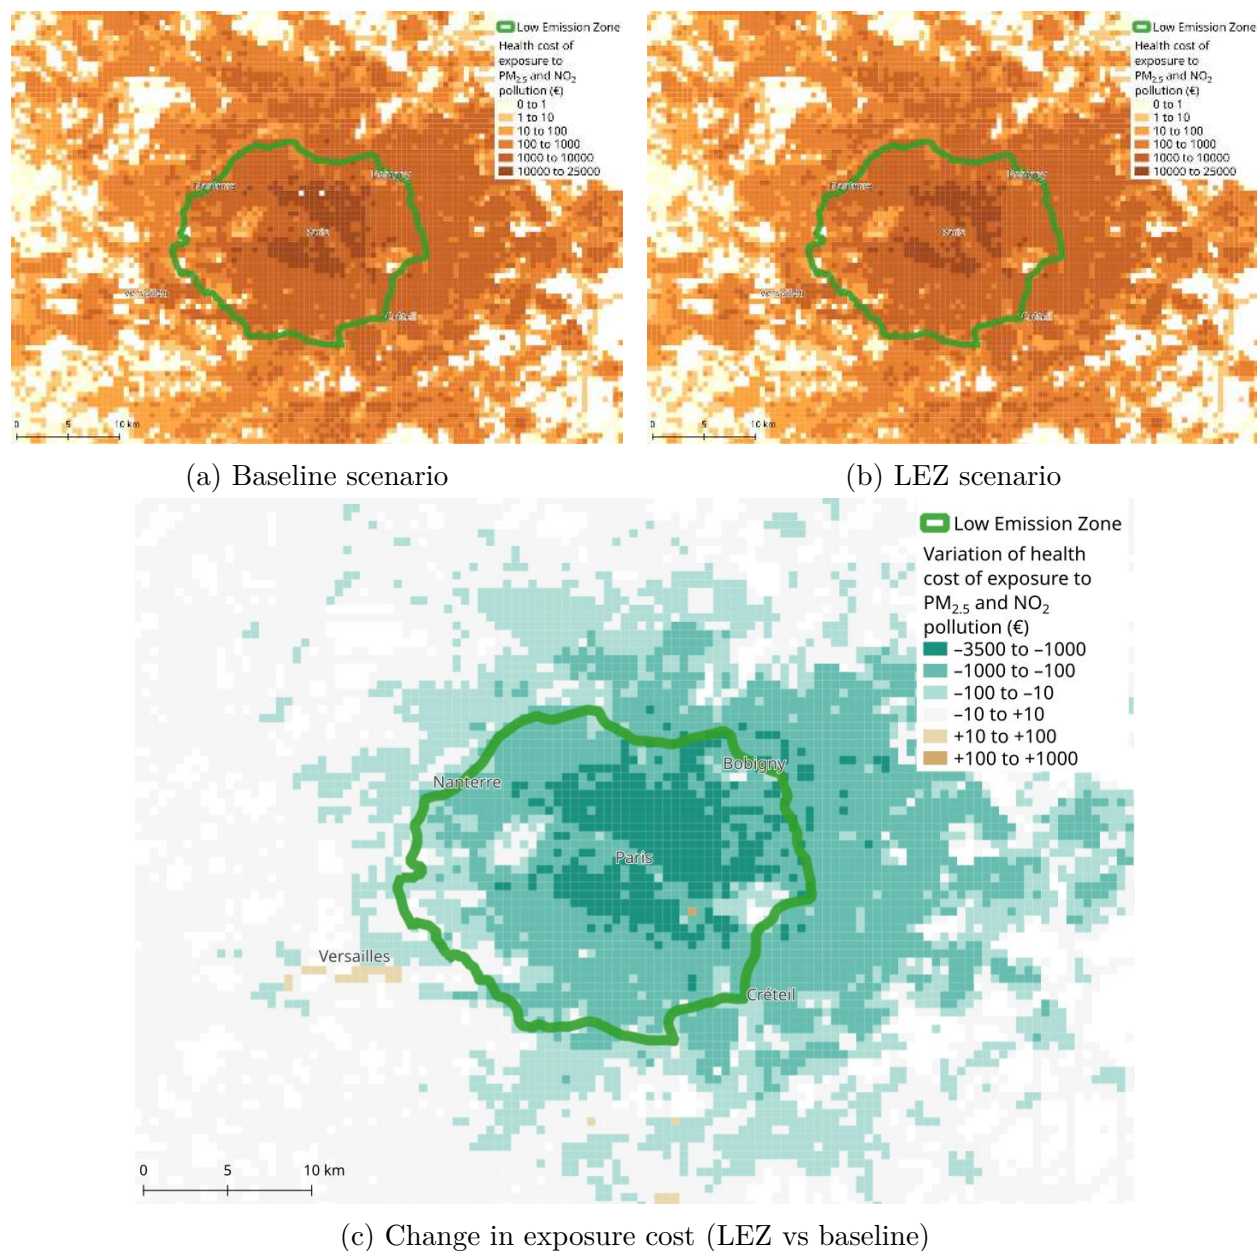


Figure 3.20: Daily exposure cost to  $\text{NO}_2$  and  $\text{PM}_{2.5}$  pollutants on a 500 m grid

### 3.7.4 Spatial and Economic Inequalities

A basic benefit-cost analysis of the LEZ policy is provided in Table 3.13, indicating that the total daily individual and social benefits outweigh its total daily costs.

Table 3.13: Basic benefit-cost analysis of the LEZ policy

Decrease in travel surplus	−1342
Reduction in health costs from pollution	1509
Reduction of CO <sub>2</sub> emissions <sup>a</sup>	196
Total	363

Note. All values are in thousand of euros per day.

<sup>a</sup> Assuming a social cost of CO<sub>2</sub> of 200 €/t.

The social acceptability of the policy might however be undermined by two factors. First, while individuals directly feel the costs – primarily due to reduced freedom in their travel options – the benefits are less tangible. The improvements in air quality, for example, are not easily observable and may not be immediately linked to the policy. Second, there are great disparities in the change in individual travel surplus, as depicted in Figure 3.21a. The total travel surplus decreases by approximately 1.342 million euros, with around 2.11% of individuals experiencing a reduction of more than 1 € in their travel surplus, accounting for a total loss of 1.972 million euros. This loss is partly offset by a surplus increase for approximately 17.69% of the population. In contrast, changes in individual exposure to pollution are more evenly distributed, as shown in Figure 3.21b, with 98% of the population experiencing a change in health exposure within the range of 0.00 € to +0.35 €.




Figure 3.21: Cumulative distribution function of the travel and health surplus (agent-level)

In this section, we analyze the spatial and socio-demographic characteristics of the win-

ners and losers from the LEZ policy, focusing on the individual travel surplus, which shows the most significant inequalities and is directly perceived by the population. We define *LEZ winners* as individuals whose travel surplus increases by more than 1 € with the LEZ. They represent 1.2% of the traveling population. Similarly, we define *LEZ losers* as individuals whose travel surplus decreases by more than 1 €. We also include in this group the 91 270 “trapped” agents assumed to travel with banned cars inside the LEZ because they have no feasible alternatives (see Section 3.6.2). These LEZ losers represent 3.3% of the population. The remaining agents (95.5% of the population) are classified as *LEZ neutrals*. They are not significantly affected by the policy.

Table 3.14 presents the total and mean travel and health surplus for the LEZ winners, losers, and neutrals. LEZ winners and losers receive comparable benefits from health improvements due to the policy (+0.13 € per LEZ winner, +0.15 € per LEZ loser). However, for the LEZ losers, this benefit is insignificant compared to the loss in travel surplus of 7.64 € on average.

Table 3.14: Travel and health surplus by population group

	LEZ Winner	LEZ Loser	LEZ Neutral
Count	122 170	344 800	9 908 530
Share of population traveling	1.2 %	3.3 %	95.5 %
Travel surplus variation (total, €) <sup>a</sup>	+202 000	-1 937 000	+376 000
Travel surplus variation (mean, €) <sup>a</sup>	+1.65	-7.64	+0.04
Health surplus variation (total, €)	+16 000	+53 000	+1 218 000
Health surplus variation (mean, €)	+0.13	+0.15	+0.12

<sup>a</sup> Excluding the “trapped” agents.

Table 3.15 compares the characteristics of agents and their trips across the three groups: winners, losers, and neutrals. LEZ winners are exclusively car owners, with a significant majority needing to travel inside the LEZ (91.7%) and owning an authorized vehicle (98.6%). Their trips are generally longer, averaging 55.1 km, and only a minority of them live inside the LEZ (27.0%) or have access to public transit (42.5%). This group is likely to benefit significantly from the reduced congestion produced by the LEZ. Additionally, some LEZ

winners own banned cars but do not need to travel inside the LEZ, likely benefiting from reduce congestion outside the LEZ.

Table 3.15: Characteristics of the agents and their trips by population group

	LEZ Winner	LEZ Loser	LEZ Neutral
Living in LEZ	27.0 %	53.3 %	40.4 %
At least one trip inside LEZ	91.7 %	99.5 %	61.2 %
Car owners	100.0 %	100.0 %	71.9 %
Banned car owners	1.4 %	99.9 %	18.0 %
Mean Euclidean distance for all trips	55.1 km	30.8 km	15.8 km
Access to public-transit for all trips <sup>a</sup>	42.5 %	65.9 %	84.3 %

<sup>a</sup> Share of agents for which all trips can be done by public transit (or short walking trips).

LEZ losers, in contrast, are almost exclusively composed of banned car owners (99.9%) whose trips are within the LEZ (99.5%).<sup>15</sup> Compared to LEZ winners, their trips tend to be shorter (30.8 km) and they have better access to public transit (65.9%), but not as much as LEZ neutrals. Many in this group might experience a significant loss because they lose access to car-based modes and do not have suitable alternative.

A common belief is that LEZs disproportionately harm individuals living far from city centers who lack access to public transit. However, Table 3.15 paints a different picture: LEZ losers are more concentrated inside the LEZ than outside, and they generally have better access to public transit than LEZ winners (though not as good as the LEZ neutrals).

Table 3.16 shows the mode shares for LEZ winners and losers in both the baseline and LEZ scenarios. LEZ winners are predominantly car drivers or passengers in both scenarios (83.5 % in the baseline, 85.5 % in the LEZ scenario). This group likely includes people who travel by car in the baseline scenario despite facing heavy congestion but benefit from reduced traffic with the LEZ. A notable proportion of LEZ winners are also pedestrians who might also benefit from the reduced road congestion for car alternatives (remember that the travel surplus does not only depend on the utility of the selected mode).

<sup>15</sup>The 0.01 % of LEZ losers owning authorized vehicles are agents negatively impacted by the adjustments in road congestion outside the LEZ.

Table 3.16: Mode shares in the baseline and LEZ scenario for the LEZ winners and losers

	LEZ Winner Baseline	LEZ Winner LEZ	LEZ Loser Baseline	LEZ Loser LEZ
Car driver	72.4 %	74.1 %	72.1 %	21.5 %
Car passenger	11.1 %	11.4 %	6.6 %	3.0 %
Public transit	4.1 %	2.2 %	2.6 %	42.2 %
Bicycle	0.0 %	0.0 %	0.4 %	7.4 %
Walking	12.4 %	12.3 %	18.3 %	26.0 %

LEZ losers on the other hand, are also mostly car users in the baseline scenario (78.7%), but many lose access to their preferred mode of transportation in the LEZ scenario. Only 24.5 % of LEZ losers continue traveling by car – mostly the “trapped” agents – with many switching to public transit (+39.6 p.p.), walking (+7.7 p.p.), or cycling (+7.0 p.p.).

Figures 3.22 and 3.23 illustrate the spatial distribution of LEZ winners and losers by municipality. LEZ losers are predominantly located within the LEZ, particularly around the zone’s boundaries. In contrast, LEZ winners are more dispersed, in municipalities outside the LEZ. This suggests that the LEZ may favor individuals living outside the LEZ more than those within.

Analyzing the impact of the LEZ by socio-demographic characteristics, such as gender and socio-professional class, is unfortunately not feasible at the agent level because household vehicles were generated without considering correlations with these attributes (see Section 3.4.2). However, the vehicle fleets were generated based on the characteristics of municipalities, meaning that the shares of LEZ winners and losers computed at the municipality level are consistent. This allows us to analyze how these shares correlate with socio-demographic characteristics aggregated at the municipality level.

We focus on how the LEZ affects municipalities based on the average monthly disposable income over households. Figures 3.24 and 3.25 illustrate the share of LEZ winners and losers as a function of the mean income of the municipalities. Both figures also present the equation from an Ordinary Least Squares (OLS) regression of the share of LEZ winners or losers on



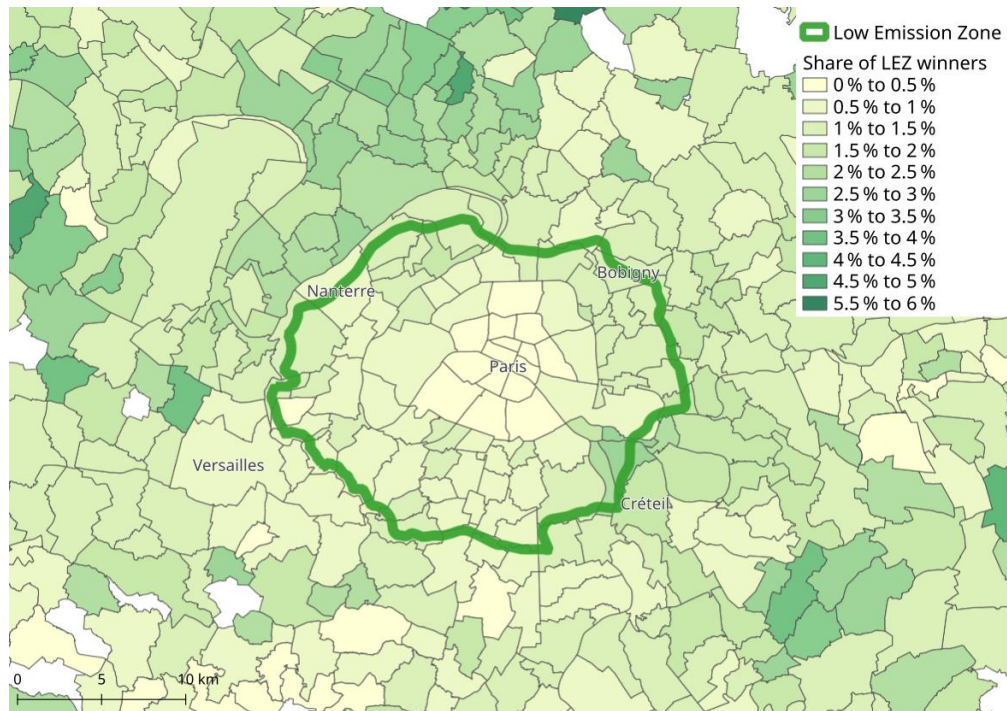


Figure 3.22: Share of LEZ winners by municipality

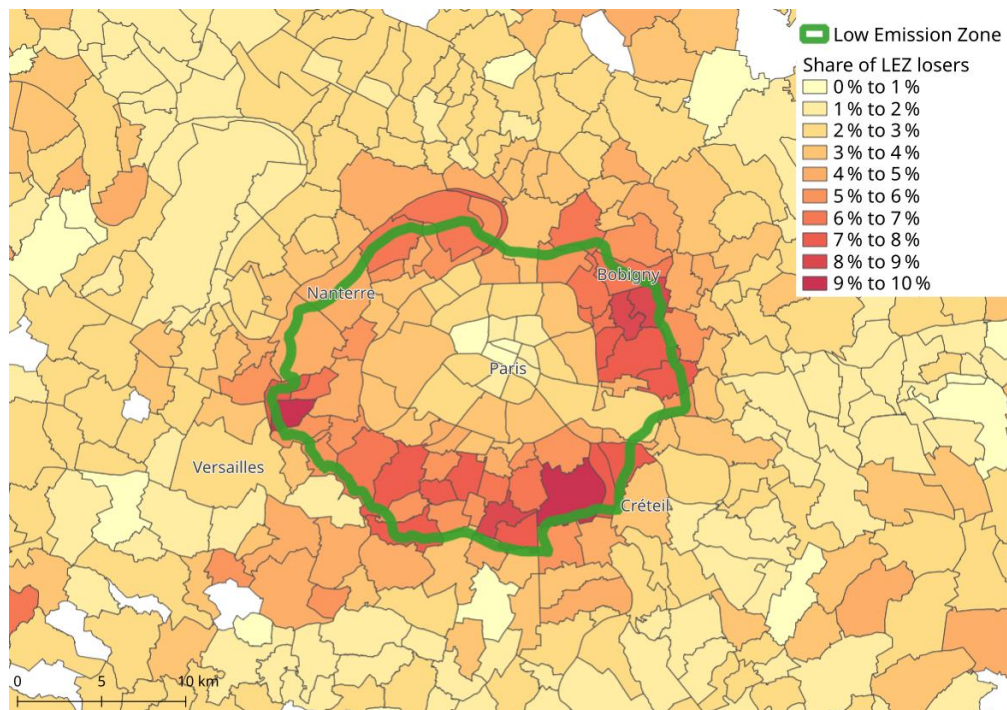


Figure 3.23: Share of LEZ losers by municipality

a constant and the logarithm of the municipality's mean income.

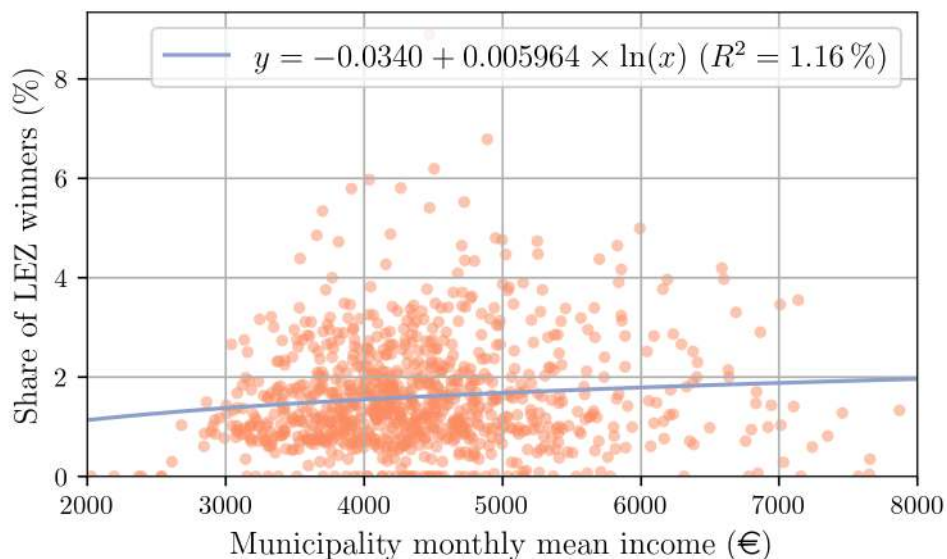


Figure 3.24: Share of LEZ winners as a function of the municipality mean income

Although the income coefficients have the expected sign (richer municipalities are expected to have a larger share of LEZ winners and a smaller share of LEZ losers), both OLS regressions show a low explanatory power, with poor  $R^2$  values (1.16 % for the share of LEZ winners, and 0.07 % for the share of LEZ losers). These results suggest that LEZ winners and losers are distributed relatively evenly across low- and high-income municipalities, indicating that the LEZ does not significantly affect economic inequalities between municipalities.

## 3.8 Conclusion

This chapter presents a comprehensive evaluation of the Low Emission Zone (LEZ) policy being implemented in the Greater Paris area, using a dynamic agent-based transport simulator (METROPOLIS2) to assess its short-term impacts on travel behavior, road congestion, and pollution emissions.

The aggregate results show that the benefits of the LEZ policy (improved air quality, reduced CO<sub>2</sub> emissions) outweigh its costs (decrease of travel surplus). Pollution reduction is




Figure 3.25: Share of LEZ losers as a function of the municipality mean income

permitted by the decrease of car trips, with the mode share of car trips decreasing by 1.9 p.p., leading to a 3.9% decrease in total vehicle-kilometers traveled. The emission reductions far exceed the decrease in car trips, with CO<sub>2</sub> emissions falling by 4.5%, nitrogen oxides (NO<sub>x</sub>) by 9.2%, and particulate matter (PM<sub>2.5</sub>) by 7.6%. These reductions highlight the LEZ's ability to target the most polluting vehicles effectively, supporting its primary goal of improving air quality.

Importantly, the analysis of premature deaths due to pollution exposure demonstrates the significant public health benefits of the LEZ. The number of premature deaths from exposure to NO<sub>2</sub> and PM<sub>2.5</sub> is expected to decline by 9.9% and 13.0%, respectively, with the greatest air quality improvements occurring within the LEZ, where population density is highest.

However, the distributional analysis reveals mixed results in terms of equity. While the health benefits of reduced pollution are distributed relatively evenly across the population, the impact on travel surplus shows greater disparities. Approximately 3.3% of the population experiences a loss of more than 1€ in daily travel surplus, primarily owners of banned vehicles residing inside the LEZ. On the other hand, 1.2% of the population, particularly



authorized vehicle owners who benefit from reduced congestion, experience a daily surplus gain of more than 1 €. This variation in travel surplus suggests that while the policy has clear health benefits, it may exacerbate inequalities. Some compensations for the most penalized individuals might be helpful to improve public acceptability.<sup>16</sup>

Several limitations of this analysis should be noted. First, the calibration of mode choices and values of time – essential for converting travel surplus into euros – relied on strong assumptions and thus may not fully reflect actual preferences. As a result, the expression of travel surplus in euros is somewhat ambitious and should be interpreted with caution. A more accurate estimation of values of time across different socio-demographic groups would improve the precision of these welfare calculations.

Second, while we exploit data on the vehicle fleet at the municipality-level, we did not consider the correlation between vehicle ownership and the socio-demographic characteristics of households within municipalities. This omission could influence the distributional impacts of the LEZ since, within a municipality, lower-income households are more likely to own older, banned vehicles. Incorporating this correlation into the simulation model would offer a more precise understanding of how the LEZ affects different segments of the population.

Moreover, the focus on short-term impacts means that variables such as car ownership, destination choice, and residential location remain fixed. In reality, individuals may adapt to the LEZ over time by purchasing authorized vehicles, relocating closer to public transport, or altering their activity patterns. Additionally, the LEZ is modeled as a strict restriction, whereas in practice, exceptions and partial enforcement may lead to different outcomes. Incorporating such long-term adjustments and policy nuances would require more sophisticated models, such as a car-ownership model able to disentangle the natural renewal of the fleet from the changes in behavior caused by the LEZ.

Despite these limitations, our findings provide valuable insights into the potential effectiveness of the LEZ in reducing emissions and improving air quality in urban areas. By

---

<sup>16</sup>Note that already car owners can benefit from a subsidy of up to 3000 € when replacing their older vehicle for a newer one, if they live or work inside the LEZ.

leveraging a detailed transport simulation, this study highlights the complex interactions between individual travel decisions and transportation policies, offering a robust framework for evaluating similar policies. Future research should extend this analysis to medium- and long-term impacts, as well as investigate the role of complementary policies, such as public transport improvements or changes in speed limits, in mitigating the social inequalities observed in this study.

# General Conclusion

## Contributions

The three chapters of this thesis demonstrate how transport simulators can be effectively applied to evaluate a variety of transport policies, such as ride-sharing and Low Emission Zones. The thesis provides a comprehensive framework for understanding the potential impacts of these policies, not only at a global level but also on individual agents.

A central contribution of this thesis is the development of METROPOLIS2, a simulator that enhances existing models by integrating finer behavioral representations and supporting a wider range of policy applications. The simulator combines discrete-choice models with dynamic congestion modeling and state-of-the-art routing algorithms to simulate the equilibrium that emerges from interactions between agents on a shared transport network. These features make METROPOLIS2 a valuable tool to analyze complex urban transportation interactions at a city or regional level.

Through applications to ride-sharing and Low Emission Zones, this thesis also proposes methodologies, intended to be easily replicable to any French metropolitan area, for calibrating simulations and assessing policy impacts. These applications illustrate how METROPOLIS2 can be used to quantify the expected benefits and limitations of policies, providing decision-makers with valuable insights as they develop transportation policies.

## Methodology Improvements

Despite these advancements, certain limitations remain.

Currently, in METROPOLIS2, public-transit itineraries are computed externally using OpenTripPlanner, and public-transit trips are treated similarly to walking trips, with constant travel times. While this enables the inclusion of public transit in mode choice and allows estimation of passenger flows at each transit stop and line, it does not account for in-vehicle congestion or reliability in the agents' decision-making process, nor does it handle the intrinsic time-dependency of public transit, such as the influence of line frequency.

During this thesis, the endogenization of departure-time and itinerary choice for public-transit trips within the simulator was conceptualized and partially modeled, but development was postponed due to the substantial amount of work involved. Integrating itinerary choice would allow for a more refined simulation by incorporating agent-level preferences to penalize itineraries with many transfers or long walking times for example. Additionally, it would enable capacity constraints on public-transit vehicles and allow agents to reroute in response to overcrowded or disrupted lines.

With the proposed methodology, the simulator would pre-compute a set of "attractive" itineraries for each trip, representing the available choice set. These "attractive" itineraries could be defined as all non-Pareto dominated options connecting an origin and destination, considering variables like walking time, waiting time, bus time, and the number of transfers.

At each iteration of the simulation, agents could then choose the itinerary that maximizes their utility within this choice set, with utility potentially depending on the expected number of passengers inside the public-transit vehicles. The expected passenger numbers could be learned over iterations, similarly to how expected road congestion is updated.

This approach is computationally efficient as it merely requires to compare itineraries, without having to compute least-cost paths at each iteration. A similar method is used for both road and public-transit trips in SimMobility (Adnan et al., 2016), and for public-transit trips only in BusMezzo (Cats et al., 2011).

One major challenge to integrate public transit within the simulator is the departure-time choice. Unlike road trips, where departure times are continuous, public-transit timeta-

bles create discrete departure options, complicating the decision-making process. Moreover, while route choice for road trips typically occurs after the departure time is selected, in public-transit, departure-time and route choice are more closely intertwined and are better understood as a joint decision, adding another layer of complexity to the simulation.

Once the simulator includes departure-time and itinerary choice for public-transit trips, further refinements could include incorporating reliability factors in public transit and modeling the interaction between buses and road congestion.

These enhancements would make METROPOLIS2 especially well-suited to assess the impacts of new public-transit infrastructure, such as the *Grand Paris Express*, by allowing for detailed considerations of in-vehicle congestion, capacity constraints, and variations in service frequency. In the Low Emission Zone evaluation presented in Chapter 3, incorporating in-vehicle congestion into the analysis could refine the predicted policy impact, as the projected increase in public-transit trips (with mode share rising from 19% to 20%) may lead to the saturation of certain transit lines.

Another potential improvement involves incorporating household-level decisions. Currently, agents belonging to the same household make decisions independently in METROPOLIS2, which means that, in Chapter 3, multiple household members might have been using the same car simultaneously for different trips. Future versions could implement joint decision-making, where, for example, spouses decide who gets to use the car or whether to carpool together, as in Picard et al. (2018). Expanding this to households with more than two members may pose additional challenges.

### **Long-Term Decisions**

The thesis focuses primarily on short-term decision-making, such as route choice, mode choice, and departure times, while holding longer-term variables like vehicle ownership, activity patterns, and residential locations constant. METROPOLIS2 could thus benefit from integration with other models, expanding the scope of analysis. One promising integration,

already initiated, is with EUrbanSim, a land-use model applied to Île-de-France, which would enable the creation of a Land-Use Transport Interaction (LUTI) model. This interaction introduces two key feedback loops: traffic congestion influences household residential location choices, and in turn, these location choices affect congestion levels.

Other potential integrations, mentioned in Chapter 3, could further broaden the scope of the simulator. For example, incorporating a car-ownership model could enable agents to make decisions about purchasing new vehicles or selling old ones, an important adjustment variable for policies such as Low Emission Zones.

Integrating an activity-based model would allow agents to adjust the duration and location of their activities based on anticipated road congestion. Given the increasing prevalence of telework, such an activity-based model could account for weekly-activity patterns, instead of the more traditional daily-activity patterns. Simulating a full week, rather than a single average day, would provide a more realistic representation of contemporary travel behaviors.

Looking ahead, future challenges in transportation will call for further advancements in simulators such as METROPOLIS2. By integrating more detailed behavioral models and accounting for new trends, such as electric bicycles or telework, these simulators can continue to evolve, providing policymakers robust tools to design sustainable, efficient, and equitable transport policies for the future.

# Bibliography

- Adnan, M., Pereira, F. C., Azevedo, C. M. L., Basak, K., Lovric, M., Raveau, S., Zhu, Y., Ferreira, J., Zegras, C., and Ben-Akiva, M. (2016). Simmobility: A multi-scale integrated agent-based simulation platform. In *95th Annual Meeting of the Transportation Research Board Forthcoming in Transportation Research Record*, volume 2. The National Academies of Sciences, Engineering, and Medicine Washington, DC.
- Agatz, N. A. H., Erera, A. L., Savelsbergh, M. W. P., and Wang, X. (2011). Dynamic ride-sharing: A simulation study in metro Atlanta. *Transportation Research Part B: Methodological*, 45(9):1450–1464.
- Agence de la transition écologique (2021). ADEME - Bilans GES. <https://www.bilans-ges.ademe.fr/fr/accueil>.
- Alisoltani, N., Leclercq, L., and Zargayouna, M. (2021). Can dynamic ride-sharing reduce traffic congestion? *Transportation research part B: methodological*, 145:212–246.
- Arnott, R., de Palma, A., and Lindsey, R. (1990). Economics of a bottleneck. *Journal of Urban Economics*, 27(1):111–130.
- Arnott, R., de Palma, A., 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.
- Bahat, O. and Bekhor, S. (2016). Incorporating Ridesharing in the Static Traffic Assignment Model. *Networks and Spatial Economics*, 16(4):1125–1149.
- Barceló, J. and Casas, J. (2005). Dynamic network simulation with aimsun. In *Simulation approaches in transportation analysis: Recent advances and challenges*, pages 57–98. Springer.
- Basu, R., Araldo, A., Akkinepally, A. P., Nahmias Biran, B. H., Basak, K., Seshadri, R., Deshmukh, N., Kumar, N., Azevedo, C. L., and Ben-Akiva, M. (2018). Automated Mobility-on-Demand vs. Mass Transit: A Multi-Modal Activity-Driven Agent-Based Simulation Approach. *Transportation Research Record: Journal of the Transportation Research Board*, 2672(8):608–618.
- Batz, G. V., Geisberger, R., Sanders, P., and Vetter, C. (2013). Minimum time-dependent travel times with contraction hierarchies. *ACM Journal of Experimental Algorithmics*, 18.
- Ben-Akiva, M. and Bierlaire, M. (1999). Discrete choice methods and their applications to short term travel decisions. In *Handbook of transportation science*, pages 5–33. Springer.
- Ben-Akiva, M., Litinas, N., and Tsunokawa, K. (1985). Continuous spatial choice: The continuous logit model and distributions of trips and urban densities. *Transportation Research Part A: General*, 19(2):119–154.
- Ben-Akiva, M. and Watanatada, T. (1981). Application of a continuous spatial choice logit

- model. *Structural analysis of discrete data with econometric applications*, pages 320–343.
- Björklund, G. and Swärdh, J.-E. (2017). Estimating policy values for in-vehicle comfort and crowding reduction in local public transport. *Transportation Research Part A: Policy and Practice*, 106:453–472.
- Börjesson, M., Bastian, A., and Eliasson, J. (2021). The economics of low emission zones. *Transportation Research Part A: Policy and Practice*, 153:99–114.
- Bou Sleiman, L. (2023). Displacing congestion: Evidence from paris. Technical Report 2302, CEPREMAP.
- Bruck, B. P., Incerti, V., Iori, M., and Vignoli, M. (2017). Minimizing CO2 emissions in a practical daily carpooling problem. *Computers & Operations Research*, 81:40–50.
- Cats, O., Koutsopoulos, H. N., Burghout, W., and Toledo, T. (2011). Effect of real-time transit information on dynamic path choice of passengers. *Transportation Research Record*, 2217(1):46–54.
- Chamberlain, R. C., Fecht, D., Davies, B., and Laverty, A. A. (2023). Health effects of low emission and congestion charging zones: A systematic review. *The Lancet Public Health*, 8(7):e559–e574.
- Chan, N. D. and Shaheen, S. A. (2012). Ridesharing in North America: Past, Present, and Future. *Transport Reviews*, 32(1):93–112.
- Charypar, D. and Nagel, K. (2005). Generating complete all-day activity plans with genetic algorithms. *Transportation*, 32(4):369–397.
- Cici, B., Markopoulou, A., Frias-Martinez, E., and Laoutaris, N. (2014). Assessing the potential of ride-sharing using mobile and social data: A tale of four cities. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '14*, pages 201–211, New York, NY, USA.
- Cornut, B. (2017). *Le Peak Car En Ile-de-France : Etude de l'évolution de La Place de l'automobile et de Ses Déterminants Chez Les Franciliens Depuis Les Années 1970*. Theses, Université Paris-Est.
- Coulombel, N., Boutueil, V., Liu, L., Vigiuié, V., and Yin, B. (2019). Substantial rebound effects in urban ridesharing: Simulating travel decisions in Paris, France. *Transportation Research Part D: Transport and Environment*, 71:110–126.
- De Borger, B. and Proost, S. (2013). Traffic externalities in cities: The economics of speed bumps, low emission zones and city bypasses. *Journal of Urban Economics*, 76:53–70.
- de Palma, A., Ben-Akiva, M., Lefèvre, C., and Litinas, N. (1983). Stochastic Equilibrium Model of Peak Period Traffic Congestion. *Transportation Science*, 17(4):430–453.
- de Palma, A., Javaudin, L., Stokkink, P., and Tarpin-Pitre, L. (2022a). Ride-sharing with inflexible drivers in the Paris metropolitan area. *Transportation*, pages 1–24.
- de Palma, A., Kilani, M., and Lindsey, R. (2005). Congestion pricing on a road network: A study using the dynamic equilibrium simulator METROPOLIS. *Transportation Research Part A: Policy and Practice*, 39(7-9):588–611.
- de Palma, A. and Lindsey, R. (2002). Comparison of morning and evening commutes in the vickrey bottleneck model. *Transportation Research Record*, 1807(1):26–33.
- de Palma, A. and Marchal, F. (1999). Analysis of Travel Cost Components Using Large-Scale, Dynamic Traffic Models. *Transportation Research Record: Journal of the Transportation Research Board*, 1676(1):177–183.
- de Palma, A. and Marchal, F. (2002). Implementation of a Dynamic Traffic Simulator to



- the Paris Area. In *Traffic And Transportation Studies (2002)*, pages 1216–1223, Guilin, China. American Society of Civil Engineers.
- de Palma, A., Marchal, F., and Nesterov, Y. (1997). METROPOLIS: Modular System for Dynamic Traffic Simulation. *Transportation Research Record: Journal of the Transportation Research Board*, 1607(1):178–184.
- de Palma, A. and Nesterov, Y. (2006). Park and ride for the day period and morning-evening commute. In *Mathematical and Computational Models for Congestion Charging*, pages 143–157. Springer.
- de Palma, A., Stokkink, P., and Geroliminis, N. (2022b). Influence of dynamic congestion with scheduling preferences on carpooling matching with heterogeneous users. *Transportation Research Part B: Methodological*, 155:479–498.
- Degraeuwe, B., Thunis, P., Clappier, A., Weiss, M., Lefebvre, W., Janssen, S., and Vranckx, S. (2017). Impact of passenger car NOX emissions on urban NO2 pollution – Scenario analysis for 8 European cities. *Atmospheric Environment*, 171:330–337.
- Delhomme, P. and Gheorghiu, A. (2016). Comparing French carpoolers and non-carpoolers: Which factors contribute the most to carpooling? *Transportation Research Part D: Transport and Environment*, 42:1–15.
- Di Febbraro, A., Gattorna, E., and Sacco, N. (2013). Optimization of Dynamic Ridesharing Systems. *Transportation Research Record*, 2359(1):44–50.
- Diao, M., Kong, H., and Zhao, J. (2021). Impacts of transportation network companies on urban mobility. *Nature Sustainability*, pages 1–7.
- Donais, F. M., Abi-Zeid, I., Waygood, E. O., and Lavoie, R. (2019). A review of cost–benefit analysis and multicriteria decision analysis from the perspective of sustainable transport in project evaluation. *EURO Journal on Decision Processes*, 7(3-4):327–358.
- DRIEAT Île-de-France (2021). Modélisation des déplacements en IDF avec MODUS 3.1. Technical report, DRIEAT Île-de-France.
- Durrmeyer, I. and Martinez, N. (2022). The welfare consequences of urban traffic regulations. Technical report, TSE Working Paper.
- Enquête Globale Transport (2010). Enquête Globale Transport 2010. Technical report, Enquête Globale Transport, Paris.
- Fattah, M. A., Morshed, S. R., and Kafy, A.-A. (2022). Insights into the socio-economic impacts of traffic congestion in the port and industrial areas of chittagong city, bangladesh. *Transportation Engineering*, 9:100122.
- Flötteröd, G. (2009). Cadyts a free calibration tool for dynamic traffic simulations. In *9th Swiss Transport Research Conference*.
- Flötteröd, G. (2016). *MATSim as a Monte-Carlo Engine*, pages 327–336. Ubiquity Press.
- Flötteröd, G., Chen, Y., and Nagel, K. (2012). Behavioral Calibration and Analysis of a Large-Scale Travel Microsimulation. *Networks and Spatial Economics*, 12(4):481–502.
- Friedrich, M., Hartl, M., and Magg, C. (2018). A modeling approach for matching ridesharing trips within macroscopic travel demand models. *Transportation*, 45(6):1639–1653.
- Galland, S., Knapen, L., Yasar, A.-U.-H., Gaud, N., Janssens, D., Lamotte, O., Koukam, A., and Wets, G. (2014). Multi-agent simulation of individual mobility behavior in carpooling. *Transportation Research Part C: Emerging Technologies*, 45:83–98.
- Geisberger, R. and Sanders, P. (2010). Engineering Time-Dependent Many-to-Many Short-

- est Paths Computation. In Erlebach, T. and Lübbecke, M., editors, *10th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS'10)*, volume 14 of *Open Access Series in Informatics (OASICs)*, pages 74–87, Dagstuhl, Germany. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- Gheorghiu, A. and Delhomme, P. (2018). For which types of trips do French drivers carpool? Motivations underlying carpooling for different types of trips. *Transportation Research Part A: Policy and Practice*, 113:460–475.
- Guo, R.-Y., Yang, H., and Huang, H.-J. (2018). Are We Really Solving the Dynamic Traffic Equilibrium Problem with a Departure Time Choice? *Transportation Science*, 52(3):603–620.
- 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.
- Hensher, D. A. and Rose, J. M. (2007). Development of commuter and non-commuter mode choice models for the assessment of new public transport infrastructure projects: A case study. *Transportation Research Part A: Policy and Practice*, 41(5):428–443.
- Herbawi, W. and Weber, M. (2012). The ridematching problem with time windows in dynamic ridesharing: A model and a genetic algorithm. In *2012 IEEE Congress on Evolutionary Computation*, pages 1–8.
- Ho, B. Q., Clappier, A., and Blond, N. (2014). Fast and Optimized Methodology to Generate Road Traffic Emission Inventories and Their Uncertainties. *CLEAN – Soil, Air, Water*, 42(10):1344–1350.
- Ho, C. and Mulley, C. (2013). Tour-based mode choice of joint household travel patterns on weekend and weekday. *Transportation*, 40:789–811.
- Holman, C., Harrison, R., and Querol, X. (2015). Review of the efficacy of low emission zones to improve urban air quality in European cities. *Atmospheric Environment*, 111:161–169.
- Hörl, S. and Balac, M. (2021). Synthetic population and travel demand for Paris and Île-de-France based on open and publicly available data. *Transportation Research Part C: Emerging Technologies*, 130:103291.
- Hörl, S., Balac, M., and Axhausen, K. W. (2018). A first look at bridging discrete choice modeling and agent-based microsimulation in MATSim. *Procedia Computer Science*, 130:900–907.
- Horni, A., Nagel, K., and Axhausen, K. W. (2011). High-resolution destination choice in agent-based demand models. working paper or preprint.
- Horni, A., Nagel, K., and Axhausen, K. W., editors (2016). *The Multi-Agent Transport Simulation MATSim*. Ubiquity Press, London.
- Host, S., Honoré, C., Joly, F., Saunal, A., Le Tertre, A., and Medina, S. (2020). Implementation of various hypothetical low emission zone scenarios in Greater Paris: Assessment of fine-scale reduction in exposure and expected health benefits. *Environmental Research*, 185:109405.
- Huang, J., Cui, Y., Zhang, L., Tong, W., Shi, Y., and Liu, Z. (2022). An overview of agent-based models for transport simulation and analysis. *Journal of Advanced Transportation*, 2022(1):1252534.
- Hörl, S., Balac, M., and Axhausen, K. W. (2019). Pairing discrete mode choice models and agent-based transport simulation with matsim. In *2019 TRB Annual Meeting Online*,

- pages 19–02409, Washington, DC. Transportation Research Board.
- Île-de-France Mobilités (2019). Enquête Globale Transport H2020 Île-de-France Mobilités-OMNIL-DRIEA / Partial results 2018. Technical report, Île-de-France Mobilités.
- Iryo, T. (2013). Properties of dynamic user equilibrium solution: Existence, uniqueness, stability, and robust solution methodology. *Transportmetrica B: Transport Dynamics*, 1(1):52–67.
- Jong, G. D., Fox, J., Daly, A., Pieters, M., and Smit, R. (2004). Comparison of car ownership models. *Transport Reviews*, 24(4):379–408.
- Kickhöfer, B. and Nagel, K. (2016). *Microeconomic Interpretation of MATSim for Benefit-Cost Analysis*, pages 353–364. Ubiquity Press.
- Kleiner, A., Nebel, B., and Ziparo, V. (2011). A Mechanism for Dynamic Ride Sharing based on Parallel Auctions. In *22th International Joint Conference on Artificial Intelligence*, pages 266–272.
- Kong, H., Zhang, X., and Zhao, J. (2020). How does ridesourcing substitute for public transit? A geospatial perspective in Chengdu, China. *Journal of Transport Geography*, 86:102769.
- Kumar, P., Hama, S., Nogueira, T., Abbass, R. A., Brand, V. S., Andrade, M. d. F., Asfaw, A., Aziz, K. H., Cao, S.-J., El-Gendy, A., Islam, S., Jeba, F., Khare, M., Mamuya, S. H., Martinez, J., Meng, M.-R., Morawska, L., Muula, A. S., Shiva Nagendra, S. M., Ngowi, A. V., Omer, K., Olaya, Y., Osano, P., and Salam, A. (2021). In-car particulate matter exposure across ten global cities. *Science of The Total Environment*, 750:141395.
- Kumar, P. and Khani, A. (2020). An algorithm for integrating peer-to-peer ridesharing and schedule-based transit system for first mile/last mile access. *arXiv:2007.07488 [cs, math]*.
- Le Frioux, R., de Palma, A., and Blond, N. (2024). Speeding towards cleaner air: An evaluation of maximum speed restrictions in île-de-france during high pollution days. Technical Report 2024-10, THEMA (THéorie Economique, Modélisation et Applications), CY Cergy Paris Université.
- Lemp, J. D., Kockelman, K. M., and Damien, P. (2010). The continuous cross-nested logit model: Formulation and application for departure time choice. *Transportation Research Part B: Methodological*, 44(5):646–661.
- Li, M., Di, X., Liu, H. X., and Huang, H.-J. (2020a). A restricted path-based ridesharing user equilibrium. *Journal of Intelligent Transportation Systems*, 24(4):383–403.
- Li, Z., Hong, Y., and Zhang, Z. (2016). An Empirical Analysis of On-Demand Ride Sharing and Traffic Congestion. SSRN Scholarly Paper ID 2843301, Social Science Research Network, Rochester, NY.
- Li, Z.-C., Huang, H.-J., and Yang, H. (2020b). Fifty years of the bottleneck model: A bibliometric review and future research directions. *Transportation Research Part B: Methodological*, 139:311–342.
- Lian, Z. and Van Ryzin, G. (2021). Optimal growth in two-sided markets. *Management Science*, 67(11):6862–6879.
- Liu, X., Titheridge, H., Yan, X., Wang, R., Tan, W., Chen, D., and Zhang, J. (2020). A passenger-to-driver matching model for commuter carpooling: Case study and sensitivity analysis. *Transportation Research Part C: Emerging Technologies*, 117:102702.
- Lo, J. and Morseman, S. (2018). The Perfect uberPOOL: A Case Study on Trade-Offs.

- Ethnographic Praxis in Industry Conference Proceedings*, 2018(1):195–223.
- Lopez, P. A., Wiessner, E., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., Lucken, L., Rummel, J., and Wagner, P. (2018). Microscopic Traffic Simulation using SUMO. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 2575–2582, Maui, HI. IEEE.
- Lu, W., Liu, L., Wang, F., Zhou, X., and Hu, G. (2020). Two-phase optimization model for ride-sharing with transfers in short-notice evacuations. *Transportation Research Part C: Emerging Technologies*, 114:272–296.
- Lu, Y., Basak, K., Carrion, C., Loganathan, H., Adnan, M., Pereira, F. C., Saber, V. H., and Ben-Akiva, M. (2015). SimMobility mid-term simulator: A state of the art integrated agent based demand and supply model. In *94th Annual Meeting of the Transportation Research Board*, pages 15–3937, Washington, DC. Transportation Research Board.
- Ma, R., Yao, L., Song, L., and Jin, M. (2019a). A novel algorithm for peer-to-peer ridesharing match problem. *Neural Computing and Applications*, 31(1):247–258.
- Ma, T.-Y., Rasulkhani, S., Chow, J. Y. J., and Klein, S. (2019b). A dynamic ridesharing dispatch and idle vehicle repositioning strategy with integrated transit transfers. *Transportation Research Part E: Logistics and Transportation Review*, 128:417–442.
- Macqueen, J. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and Probability/University of California Press*.
- Madow, W. G. and Madow, L. H. (1944). On the Theory of Systematic Sampling, I. *The Annals of Mathematical Statistics*, 15(1):1–24.
- Malichová, E., Pourhashem, G., Kováčiková, T., and Hudák, M. (2020). Users’ Perception of Value of Travel Time and Value of Ridesharing Impacts on Europeans’ Ridesharing Participation Intention: A Case Study Based on MoTiV European-Wide Mobility and Behavioral Pattern Dataset. *Sustainability*, 12(10):4118.
- Masoud, N., Nam, D., Yu, J., and Jayakrishnan, R. (2017). Promoting Peer-to-Peer Ridesharing Services as Transit System Feeders. *Transportation Research Record*, 2650(1):74–83.
- McNally, M. G. (2007). The Four-Step Model. In Hensher, D. A. and Button, K. J., editors, *Handbook of Transport Modelling*, volume 1, pages 35–53. Emerald Group Publishing Limited.
- Miller, E. J., Roorda, M. J., and Carrasco, J. A. (2005). A tour-based model of travel mode choice. *Transportation*, 32:399–422.
- Mitra, R. and Nash, S. (2019). Can the built environment explain gender gap in cycling? An exploration of university students’ travel behavior in Toronto, Canada. *International Journal of Sustainable Transportation*, 13(2):138–147.
- Nagel, K. (1996). Particle hopping models and traffic flow theory. *Physical Review E*, 53(5):4655–4672.
- Nagel, K. and Flötteröd, G. (2016). *Agent-Based Traffic Assignment*, pages 315–326. Ubiquity Press.
- Neoh, J. G., Chipulu, M., and Marshall, A. (2017). What encourages people to carpool? An evaluation of factors with meta-analysis. *Transportation*, 44(2):423–447.
- Nguyen, J., Powers, S. T., Urquhart, N., Farrenkopf, T., and Guckert, M. (2021). An overview of agent-based traffic simulators. *Transportation Research Interdisciplinary Per-*

- spectives*, 12:100486.
- Oh, S., Seshadri, R., Azevedo, C. L., Kumar, N., Basak, K., and Ben-Akiva, M. (2020). Assessing the impacts of automated mobility-on-demand through agent-based simulation: A study of Singapore. *Transportation Research Part A: Policy and Practice*, 138:367–388.
- Otsubo, H. and Rapoport, A. (2008). Vickrey’s model of traffic congestion discretized. *Transportation Research Part B: Methodological*, 42(10):873–889.
- Patriksson, M. (2015). *The Traffic Assignment Problem: Models and Methods*. Dover Publications, Mineola, New York.
- Picard, N., Dantan, S., and de Palma, A. (2018). Mobility decisions within couples. *Theory and Decision*, 84(2):149–180.
- Pinto, G. A., Vieira, K. C., Carvalho, E. G., and Sugano, J. Y. (2019). Applying the lazy user theory to understand the motivations for choosing carpooling over public transport. *Sustainable Production and Consumption*, 20:243–252.
- Player, L., Prosser, A. M., Thorman, D., Tirion, A. S., Whitmarsh, L., Kurz, T., and Shah, P. (2023). Quantifying the importance of socio-demographic, travel-related, and psychological predictors of public acceptability of low emission zones. *Journal of Environmental Psychology*, 88:101974.
- Poulhès, A. and Proulhac, L. (2021). The Paris Region low emission zone, a benefit shared with residents outside the zone. *Transportation Research Part D: Transport and Environment*, 98:102977.
- Qian, Z. and Zhang, H. (2011). Modeling multi-modal morning commute in a one-to-one corridor network. *Transportation Research Part C: Emerging Technologies*, 19(2):254–269.
- Rasouli, S. and Timmermans, H. (2014). Activity-based models of travel demand: Promises, progress and prospects. *International Journal of Urban Sciences*, 18(1):31–60.
- Reck, D. J. and Axhausen, K. W. (2020). Subsidized ridesourcing for the first/last mile: how valuable for whom? *European Journal of Transport and Infrastructure Research*, 20(4):59–77.
- Saifuzzaman, M., de Palma, A., and Motamedi, K. (2012). Calibration of METROPOLIS for Ile-de-France. working paper or preprint.
- Saifuzzaman, M., Engelson, L., Kristoffersson, I., and de Palma, A. (2016). Stockholm congestion charging: An assessment with METROPOLIS and SILVESTER. *Transportation Planning and Technology*, 39(7):653–674.
- Schaller, B. (2021). Can sharing a ride make for less traffic? Evidence from Uber and Lyft and implications for cities. *Transport Policy*, 102:1–10.
- Shaheen, S. and Cohen, A. (2019). Shared ride services in North America: Definitions, impacts, and the future of pooling. *Transport Reviews*, 39(4):427–442.
- Shaheen, S. A., Chan, N. D., and Gaynor, T. (2016). Casual carpooling in the San Francisco Bay Area: Understanding user characteristics, behaviors, and motivations. *Transport Policy*, 51:165–173.
- Small, K. A. (1982). The scheduling of consumer activities: Work trips. *The American Economic Review*, 72(3):467–479.
- Standing, C., Standing, S., and Biermann, S. (2019). The implications of the sharing economy for transport. *Transport Reviews*, 39(2):226–242.
- Tarriño-Ortiz, J., Soria-Lara, J. A., Gómez, J., and Vassallo, J. M. (2021). Public Acceptability of Low Emission Zones: The Case of “Madrid Central”. *Sustainability*, 13(6):3251.

- Teubner, T. and Flath, C. M. (2015). The Economics of Multi-Hop Ride Sharing. *Business & Information Systems Engineering*, 57(5):311–324.
- Tibshirani, R. (1996). Regression Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288.
- Train, K. and McFadden, D. (1978). The goods/leisure tradeoff and disaggregate work trip mode choice models. *Transportation research*, 12(5):349–353.
- Vickrey, W. S. (1969). Congestion Theory and Transport Investment. *The American Economic Review*, 59(2):251–260.
- Vosough, S., de Palma, A., and Lindsey, R. (2022). Pricing vehicle emissions and congestion externalities using a dynamic traffic network simulator. *Transportation Research Part A: Policy and Practice*, 161:1–24.
- Waddell, P. (2002). UrbanSim: Modeling Urban Development for Land Use, Transportation, and Environmental Planning. *Journal of the American Planning Association*, 68(3):297–314.
- Waddell, P., Bhat, C., Eluru, N., Wang, L., and Pendyala, R. M. (2007). Modeling interdependence in household residence and workplace choices. *Transportation Research Record*, 2003(1):84–92.
- Wardman, M. (2001). A review of British evidence on time and service quality valuations. *Transportation Research Part E: Logistics and Transportation Review*, 37(2):107–128.
- Xu, H., Pang, J.-S., Ordóñez, F., and Dessouky, M. (2015). Complementarity models for traffic equilibrium with ridesharing. *Transportation Research Part B: Methodological*, 81:161–182.
- Yan, S. and Chen, C.-Y. (2011). An optimization model and a solution algorithm for the many-to-many car pooling problem. *Annals of Operations Research*, 191(1):37–71.
- Yin, B., Diallo, A. O., Seregina, T., Coulombel, N., and Liu, L. (2024). Evaluation of Low-Traffic Neighborhoods and Scale Effects: The Paris Case Study. *Transportation Research Record: Journal of the Transportation Research Board*, 2678(1):88–101.
- Yu, X., van den Berg, V. A. C., and Verhoef, E. T. (2019). Carpooling with heterogeneous users in the bottleneck model. *Transportation Research Part B: Methodological*, 127:178–200.
- Zhao, Y. and Sadek, A. W. (2012). Large-scale agent-based traffic micro-simulation: Experiences with model refinement, calibration, validation and application. *Procedia Computer Science*, 10:815–820.
- Ziemke, D., Kaddoura, I., and Nagel, K. (2019). The MATSim Open Berlin Scenario: A multimodal agent-based transport simulation scenario based on synthetic demand modeling and open data. *Procedia Computer Science*, 151:870–877.

