



UNIVERSITÀ
DI PAVIA

Department of Economics and Management

**Master's Programme in Economics,
Development and Innovation (MEDI)**

***The Role of IoT-Driven Smart Traffic
Systems in Reducing Economic Loss from
Congestion in Milan***

Supervisor:

Prof. Elisabetta Maria Venco

elisabettamaria.venco@unipv.it

Student:

**Mohamed Mamdouh Sayed Baioumy Ibrahim
mohamedmamdouh.ibrahim01@universitadipavia.it**

ID. 533681

Academic Year 2023-2025

Dedication

To my mother and my sister, for every sacrifice made quietly across the distance, and every weight carried so I could carry this one.

To my girlfriend, for walking this journey beside me, step by step.

Acknowledgements

This thesis would not have been possible without the guidance of Professor Elisabetta Maria Venco, whose expertise in urban planning and smart city frameworks proved invaluable to a work that sat at the intersection of economics and urban systems.

What made her contribution especially meaningful was that she chose to support this research despite it falling outside her own department, a gesture of intellectual generosity I will always remember.

Her patience in following my progress, her willingness to engage with each update, and her steady encouragement through every revision shaped this thesis in ways that go far beyond the footnotes. I am deeply grateful.

Abstract

Urban congestion generates substantial economic losses for contemporary cities through time loss, reduced reliability, lower productivity, and environmental externalities. Although Internet of Things (IoT)-driven smart traffic systems are increasingly used to improve traffic management, empirical evidence on their measurable economic effects remains limited, particularly in Europe.

This thesis investigates the effects of an IoT-enabled congestion regulation system in Milan (Italy), focusing on the Area C charging zone. The analysis is based on a 36-month daily administrative dataset derived from ANPR gate-entry records published by the Municipality, recording vehicles accessing the restricted area during active charging periods. The empirical strategy adopts a quasi-experimental approach centered on the October 2023 tariff increase, combining interrupted time-series (ITS) logic, a local event-study design, a longer-run filtered regression, a log-specification robustness check, and a placebo timing test.

The results provide evidence of a significant reduction in vehicle entries following the tariff increase, including a short-run decline of approximately 12.8% in daily entries and a smaller but persistent longer-run effect. The response varied across user categories, with the strongest adjustment concentrated among more flexible payer groups. These effects are interpreted through a partial welfare framework based on value-of-time estimates, allowing an assessment of congestion-relief benefits associated with reduced traffic demand.

The findings contribute to the literature on smart mobility and urban transport economics by showing that Area C operates not simply as a pricing mechanism, but as smart regulatory infrastructure in which pricing, enforcement, and continuous data generation are operationally unified. From a policy perspective, the study indicates that IoT-enabled congestion pricing can serve as an effective demand-management instrument when embedded in a broader planning framework and supported by continuous administrative monitoring.

Keywords: Internet of Things (IoT), Congestion Pricing, Automatic Number Plate Recognition (ANPR), Area C, Milan, Urban Congestion, Value of Time, Welfare Analysis

Titolo Italiano: *Il ruolo dei sistemi di traffico intelligenti basati sull'IoT nella riduzione delle perdite economiche dovute alla congestione a Milano*

Abstract

La congestione urbana genera rilevanti perdite economiche per le città contemporanee, in termini di perdita di tempo, minore affidabilità degli spostamenti, riduzione della produttività ed esternalità ambientali. Sebbene i sistemi intelligenti di gestione del traffico basati sull'Internet of Things (IoT) siano sempre più utilizzati per migliorare la gestione della mobilità, le evidenze empiriche sui loro effetti economici misurabili restano ancora limitate, in particolare in Europa.

Questa tesi analizza gli effetti di un sistema di regolazione della congestione abilitato dall'IoT nella città di Milano (Italia), con specifico riferimento alla zona a pedaggio Area C. L'analisi si basa su un dataset amministrativo giornaliero di 36 mesi, ricavato dai registri di accesso ai varchi ANPR pubblicati dal Comune, che rilevano i veicoli in ingresso nell'area soggetta a restrizione durante i periodi di attivazione della tariffazione. La strategia empirica adotta un approccio quasi-sperimentale incentrato sull'aumento tariffario dell'ottobre 2023, combinando la logica delle interrupted time series (ITS), un disegno event-study locale, una regressione filtrata di più lungo periodo, un controllo di robustezza con specificazione logaritmica e un test placebo.

I risultati mostrano una riduzione significativa degli ingressi veicolari in seguito all'aumento tariffario, con una diminuzione di breve periodo pari a circa il 12,8% degli ingressi giornalieri e un effetto di più lungo periodo più contenuto ma persistente. La risposta varia inoltre tra le

categorie di utenti, con l'aggiustamento più marcato concentrato nei gruppi di pagatori più flessibili. Tali effetti sono interpretati attraverso un quadro di welfare parziale basato sul valore del tempo, che consente di valutare i benefici di riduzione della congestione associati alla diminuzione della domanda di accesso.

I risultati contribuiscono alla letteratura sulla smart mobility e sull'economia dei trasporti urbani mostrando che Area C non opera semplicemente come un meccanismo di tariffazione, ma come un'infrastruttura regolativa intelligente nella quale tariffazione, enforcement e generazione continua di dati sono operativamente integrati. Dal punto di vista delle politiche pubbliche, lo studio indica che la congestion pricing abilitata da sistemi IoT può costituire uno strumento efficace di gestione della domanda, a condizione che sia inserita in un più ampio quadro pianificatorio e supportata da un monitoraggio amministrativo continuo.

Parole chiavi:

Internet of Things (IoT), tariffazione della congestione, riconoscimento automatico delle targhe (ANPR), Area C, Milano, congestione urbana, valore del tempo, analisi del welfare

Table of Contents

Introduction	11
· Background and Context	12
· Research Problem and Questions	14
· Objectives and Scope of the Study	15
· Thesis Structure	16
Chapter 1: Literature Review and Theoretical Framework	17
1.1 Urban Traffic Congestion and Economic Implications	17
1.1.1 Conceptual Understanding	17
1.1.2 Historical and Urban Evolution of Congestion	18
1.1.3 Causes and Contributing Factors	20
1.1.3.1 Urbanization and Population Growth	20
1.1.3.2 Rising Vehicle Ownership and Parking Accessibility	23
1.1.3.3 Infrastructure, Network Constraints, and Short-Run Disruptions	24
1.1.3.4 Modal Imbalance and Evolving Demand Patterns	25
1.1.4 Economic Implications of Congestion	27
1.2 Smart Cities and IoT in Traffic Management	28
1.2.1 The Smart City Paradigm	28
1.2.2 IoT in Urban Traffic Management: Mechanisms and Policy Relevance	29
1.2.3 Milan's IoT and Smart Mobility Ecosystem	30
1.3 Milan Area C as a Case Study	32
1.3.1 Milan's Congestion Problem: Structural Pressures	32
1.3.2 From Ecopass to Area C and Area B: Milan's Demand-Management Policy Path ...	33
1.3.3 Area C as a Smart Mobility Policy Case in Milan	36

1.4 Theoretical Models and Analytical Framework	36
1.4.1 Transport and Congestion Theories	37
1.4.2 Urban Systems and Complexity	38
1.4.3 Smart City Governance Models	39
1.4.4 Planning Perspectives and the Economic Results in Context	40
1.4.5 Conceptual Impact Pathway	43
Chapter 2: Research Methodology	44
2.1 Case Setting, Scope, and Unit of Analysis	46
2.2 Research Design and Identification Strategy	47
2.3 Data Sources and Dataset Construction	49
2.4 Econometric Models	51
2.4.1 Local Event-Study Around the 30 October 2023 Price Increase	51
2.4.2 Longer-Run Filtered Level Model	52
2.4.3 Long-Run Log Specification	53
2.4.4 Placebo Timing Test	53
2.5 Economic Valuation Methodology	54
2.6 Methodological Validity, Limitations, and Delimitations	56
Chapter 3: Data Analysis and Impact Assessment	59
3.1 Area C as the Empirical IoT Traffic System Under Study	59
3.2 Descriptive Traffic Patterns and Policy Timing	61
3.3 Empirical Design	63
3.4 Main Empirical Results	66
3.5 Robustness and Validation	69
3.6 Economic Valuation Framework	71

3.6.1	Scope, Appraisal Basis, and Valuation Logic	71
3.6.2	Valuation Assumptions and Scenario Design	72
3.6.3	Monetized Congestion-Relief Results	74
3.6.4	Interpretation, Limitations, and Extension Toward Fuller Appraisal	75
Chapter 4:	Discussion and Policy Recommendations	76
4.1	Behavioral Findings and Economic Interpretation	76
4.2	Area C as Smart Regulatory Infrastructure	80
4.3	Governance, Data Ethics, and Public Legitimacy	82
4.4	Policy Implications for Milan	83
4.5	Broader Implications for Smart Urban Mobility Policy	84
4.6	Limitations	86
Chapter 5:	Conclusion	87
5.1	Main Findings and Contributions	87
5.2	Directions for Future Research	88
References	90
Appendices	103
-	Full scenario monetization grid for estimated Area C traffic reductions (All scenarios)	103
-	Weekly-binned event-study coefficients for payer response within a ± 35 -day window	104
-	Daily Area C transits after filtering August and the year-end holiday period	105
-	Data and code availability	106

List of Tables

<i>Table 3.1. IoT monitoring structure: average daily transits by areac flag, selected months 2023</i>	60
<i>Table 3.2. Summary statistics of traffic data for Area C: regression estimates</i>	67
<i>Table 3.3. Value-of-time inputs used in the valuation (Italy, short-distance urban)</i>	73
<i>Table 3.4. Monetized congestion-relief benefits under the central VoT scenario</i>	74
<i>Table A-1. Full scenario monetization grid (all VoT and time-saving combinations)</i>	103

List of Figures

<i>Figure 1-1. Map of Milan population density across districts (2014–2024)</i>	21
<i>Figure 1-2. Milan municipality within the metropolitan catchment and primary road structure</i>	23
<i>Figure 1-3. Administrative spatial borders of Area C and Area B in Milan within its 9 zones</i>	35
<i>Figure 1-4. Spatial extent of Milan’s Area C limited traffic zone with camera gate entry points</i>	37
<i>Figure 2-1. Conceptual framework of the research methodology</i>	45
<i>Figure 3-1. Monthly average Area C traffic, December 2021 – November 2024</i>	61
<i>Figure 3-2. Daily Area C transits zoomed around October 2023</i>	62
<i>Figure 3-3. Weekly-binned event study for total transits around 30 October 2023</i>	66
<i>Figure 3-4. Payer composition indexed series around October 2023</i>	68
<i>Figure A-1. Event-study estimates for Area C payer transits around the 30 October 2023 price</i>	104
<i>Figure A-2. Daily Area C transits after filtering August and the year-end holiday period</i>	105

Glossary of Terms and Acronyms

ANPR: Automatic Number Plate Recognition — camera-based technology that reads and records vehicle license plates automatically, used in Area C and Area B gates to identify and classify vehicles in real time.

Area B: Milan's Low Emission Zone (LEZ), launched in 2019: a large-scale electronically monitored perimeter covering approximately 70% of the city area (132 km²) with 185 ANPR gates, designed to exclude the most polluting vehicles from the majority of the city.

Area C: Milan's congestion-charging zone in the historic center (Cerchia dei Bastioni, 8.2 km²), monitored by 43 ANPR-equipped gates, where a daily access charge applies to most motorized vehicles during weekday operating hours.

CBA: Cost-Benefit Analysis — a structured economic appraisal that quantifies and compares all relevant costs and benefits of a policy or project to assess its net social value.

Ecopass: Milan's predecessor to Area C (2008–2012): a pollution charge differentiated by vehicle emission class, applied within the same *Cerchia dei Bastioni* perimeter. Phased out after the share of exempt vehicles reached ~90%, undermining its effectiveness.

Event-study (local): An econometric method that traces the dynamic response of an outcome variable in relative-time bins around a focal intervention date, allowing identification of immediate and post-treatment effects while controlling for pre-existing trends.

ITS: Intelligent Transportation Systems — the integration of information and communication technologies with transport infrastructure to improve safety, efficiency, and sustainability.

IoT: Internet of Things — a network of physical devices embedded with sensors, software, and connectivity to collect and exchange data in real time.

LEZ: Low Emission Zone — a defined geographic area from which vehicles not meeting specified emission standards are restricted, typically enforced through ANPR monitoring.

MaaS: Mobility as a Service — the integration of multiple transport modes into a single user-facing platform supporting trip planning, booking, and payment.

Pigouvian tax/charge: A charge set equal to the marginal external cost imposed by an activity on third parties, designed to internalize the externality and align private and social costs.

Smart Infrastructure: Urban infrastructure enhanced by digital technologies to optimize performance, enable adaptive management, and improve service quality for users.

Smart Mobility: The use of innovative technologies, data analytics, and connected systems to create more efficient, sustainable, and user-centric transportation solutions.

SUMP: Sustainable Urban Mobility Planning — the EU framework for integrated, participatory, evidence-based urban mobility planning that emphasizes multimodality, accessibility, equity, and environmental quality.

TOD: Transit-Oriented Development — an urban planning approach that concentrates density, mixed uses, and pedestrian-oriented design around transit nodes to reduce car dependence.

VoT: Value of Time — a monetary measure of the welfare value of a unit of time saved or lost, used as the central input for congestion-relief valuation.

ZTL: Zona a Traffico Limitato — Limited Traffic Zone; the Italian regulatory designation for a restricted access area, within which vehicle access is controlled and monitored.

Introduction

Urban congestion is one of the most persistent and economically costly features of modern city life. It converts productive time into delay, raises logistics costs, degrades air quality, and imposes a structurally unequal burden on those least able to avoid it. Yet for all its familiarity, congestion remains analytically difficult to address: it is a network-scale externality that cannot be resolved simply by expanding road capacity, and its effective management requires instruments capable of observing, measuring, and responding to mobility behavior in real time. This thesis asks whether, and to what extent, a specific class of such instruments, IoT-driven smart traffic systems, can be empirically demonstrated to reduce the economic loss associated with urban congestion. It does so through a quantitative quasi-experimental evaluation of Milan's Area C congestion-charging zone, exploiting the continuous administrative monitoring data generated by its ANPR gate network to identify and monetize the behavioral response to the October 2023 tariff reform. Milan provides the empirical setting for two distinct reasons. The first is institutional. As Italy's economic capital and the functional core of a polycentric Lombardy region of approximately 9.5 million inhabitants, Milan has sustained an electronically monitored demand-management policy across three successive schemes spanning more than fifteen years (AMAT, 2023). Ecopass (2008), Area C (2012), and Area B (2019) together constitute a layered, ANPR-based regulatory architecture that is unique in the Italian context and notable in the European one (Sevino, 2019). The second reason is empirical. Unlike comparable congestion-charging systems in London, Stockholm, or Singapore, Milan's Area C publishes its transaction-level monitoring records as open data, providing a continuous, high-frequency, and publicly accessible record of behavioral responses to policy change. It is this combination of institutional maturity and data openness that makes Milan one of the most tractable empirical foundations available in the European context for the research question this thesis pursues.

In this thesis, Area C is treated as an IoT-driven smart traffic system in a functional rather than rhetorical sense. Its ANPR gate network does not simply apply a charge; it continuously senses entries, classifies vehicles, automates enforcement, and generates a structured administrative record of access behavior over time. This combination of monitoring, processing, and data production is what makes the system analytically important for this study. Thus, the thesis not only examines whether pricing affects behavior, but also whether a digitally mediated traffic-governance system can make congestion reduction empirically observable and economically interpretable. The quasi-experimental identification strategy in Chapter 2 and the smart-city framing in Chapter 1 are connected by this distinction between an administratively imposed charge and a system that continuously records its own effects.

Background and Context

The economic cost of urban congestion operates across multiple channels. At the individual level, congestion raises both the expected duration of trips and their variance, imposing a reliability cost that often exceeds the cost of average delay alone. When arrival times become unpredictable, travelers must add precautionary buffer time, and firms must absorb scheduling uncertainty in ways that compound the welfare loss beyond what average minutes lost would suggest. At the network level, the relationship between vehicle numbers and travel time is nonlinear: as demand approaches road capacity, small additional volumes trigger disproportionate delays. At the city scale, these individual and network-level effects aggregate into measurable productivity losses, raised logistics costs, higher retail prices, and degraded air quality. Understanding congestion as a nonlinear, network-scale externality rather than a per-vehicle inconvenience is the conceptual foundation on which the economic valuation framework of this thesis is built. Addressing it through demand-management rather than infrastructure expansion has become the dominant policy orientation in European cities

precisely because the induced-demand mechanism shows that additional road capacity generates additional vehicle travel, leaving congestion approximately unchanged in the long run (Duranton & Turner, 2011). Pricing and access controls, by contrast, raise the generalized cost of a trip to the social marginal cost, internalizing the externality and reducing demand toward the socially efficient level. The practical challenge is that effective demand management requires the ability to observe behavior, enforce charges, and assess outcomes in real time. These are exactly the functions that IoT-based monitoring systems are designed to support. Milan's structural position makes it an acute case of these dynamics. Its radial road network, organized around the historic Cerchia dei Bastioni inner ring, funnels metropolitan commuting demand from a catchment of millions through a limited set of access corridors. The city's sustained investment in electronically monitored demand-management reflects both the severity of that structural pressure and a political tradition, consolidated by the 2011 referendum on Area C, of treating access pricing as a democratically legitimate governance instrument rather than a technocratic imposition.

Milan's Area C has been operating since January 2012. It evolved from Ecopass (2008 to 2012), a pollution charge whose design flaw, pricing emissions rather than access, allowed drivers to upgrade vehicles and re-enter freely until the exempt share approached 90%, eroding its effectiveness (Sevino, 2019). Area C corrected this by applying a flat daily charge, irrespective of emission class. Area B, launched in 2019, extended electronically monitored access control to approximately 70% of the city's territory via 185 ANPR gates, serving as a complementary low-emission layer. The October 2023 reforms evaluated in this thesis, a vehicle eligibility tightening on 2 October and a tariff increase from 5 to 7.50 euros on 30 October, are the most recent policy episode within this trajectory (Moulin & Urbano, 2025).

Research Problem and Questions

Despite the widespread adoption of congestion pricing and IoT-based traffic management in European cities, the empirical literature on the effects of tariff reforms within mature, established regimes remains thin. Most quasi-experimental studies evaluate the introduction of congestion pricing into previously unregulated environments, and few exploit the continuous administrative records that digital monitoring systems generate for post-reform behavioral evaluation. This thesis addresses that gap.

Its primary research question is: **to what extent has the October 2023 Area C tariff reform produced a measurable and economically meaningful reduction in congestion-related access demand in Milan, and how can that reduction be translated into a bounded welfare interpretation?**

Three supporting questions structure the analysis.

First, what was the magnitude and temporal profile of the behavioral response, specifically whether the effect was concentrated in the immediate post-reform window or persisted over a longer horizon?

Second, was the response uniform across user categories, or did it concentrate among specific payer groups in ways that reveal differential price sensitivity and distributional consequences?

Third, what do the planning perspectives on accessibility, the 15-minute city model, SUMP governance, and transport-related social equity imply for interpreting the welfare estimates that the empirical analysis produces?

Objectives and Scope of the Study

The thesis has three objectives. The first is empirical: to identify the behavioral response to the October 2023 tariff reform using quasi-experimental methods applied to the Area C open administrative dataset, establishing its magnitude, direction, and persistence. The second is economic: to translate the identified traffic reduction into a monetized partial welfare estimate of congestion-relief benefit using a transparent value-of-time framework grounded in European Commission transport appraisal guidance. The third is analytical: to situate both findings within the theoretical and planning frameworks that give them their full interpretive meaning, including Pigouvian pricing, induced demand, urban systems complexity, the 15-minute city model, SUMP governance, and transport-related social equity.

The study is deliberately bounded in scope. It does not evaluate the full social cost-benefit balance of Area C, estimate fiscal performance, or model citywide traffic equilibrium. The economic valuation is a partial welfare approximation that monetizes one identified benefit channel, congestion relief at the charging perimeter, rather than a general-equilibrium account of all traffic and welfare effects across the Milan road network. This methodological restraint is a strength rather than a limitation: it keeps the empirical claims aligned with the evidentiary base and avoids the spurious precision that broader appraisals often generate when data constraints are understated. The geographic scope is restricted to the Area C monitored perimeter; the temporal scope covers 36 months from December 2021 to November 2024; and the unit of analysis is the daily entry count. These delimitations are described and justified in full in Chapter 2.

Thesis Structure

The thesis is organized across five chapters. Chapter 1 reviews the literature and develops the theoretical framework, covering urban congestion as an economic externality problem, Milan's IoT and smart mobility ecosystem, the evolution of demand-management policy from Ecopass through Area C to Area B, and the planning perspectives, including the 15-minute city, TOD, SUMP, and transport-related social equity, that situate the economic results in a broader governance context. Chapter 2 describes the research methodology: the construction of the daily entry dataset from the Area C open-data platform, the four-layer identification strategy, and the value-of-time valuation framework. Chapter 3 presents the empirical results, including the short-run and longer-run traffic effects, the payer-category decomposition, and the corresponding monetized welfare estimates. Chapter 4 interprets these findings in relation to the IoT governance argument, the planning literature, and the implications for Milan and comparable European cities, and closes with five policy recommendations. Chapter 5 synthesizes the contributions of the thesis, addresses its limitations, and identifies directions for future research.

Chapter 1: Literature Review & Theoretical Framework

1.1 Urban Traffic Congestion & Economic Implications:

1.1.1 Conceptual Understanding

Academics have interpreted the term "congestion" from various perspectives. In the context of traffic, it is broadly defined as the condition in which demand for mobility exceeds the capacity of the road infrastructure, resulting in significantly slow car speeds and increased risk of traffic jams or possibly gridlock (Pi et al., 2021). Initially, it occurs due to population growth and an ongoing increase in urban density, and is subsequently accompanied by the rapid expansion of ridesharing and delivery services, motor-vehicle use, and supporting transportation infrastructure (Afrin & Yodo, 2020; Jayasooriya & Bandara, 2017). In urban settings, congestion is rarely a purely physical phenomenon: it is also an economic problem, because each additional vehicle incurs time costs on other users that are not fully internalized in individual travel decisions (Mosquera, 2024a).

Congestion is mainly divided into two types: recurrent and non-recurrent. Where recurrent congestion refers to regular, predictable traffic congestion that occurs at specific locations and times, such as during daily rush hours, it is mainly caused by steady demand exceeding road capacity. Non-recurrent congestion, on the other hand, is caused by factors such as accidents, roadwork, unstable weather, and other unpredictable events, including operational failures. Those events interrupt the usual pace and may result in significant, albeit brief, delays (Tang & Heinemann, 2018; Lieberthal et al., 2024). Policy and technology responses differ by congestion type: recurrent congestion, associated with regular demand peaks, is more commonly addressed through demand-management instruments and traffic-control optimization, whereas non-recurrent congestion, triggered by disruptions such as incidents or other perturbations, depends more heavily on rapid detection, response, and incident-management systems (Gu et al., 2020).

Urban traffic congestion is not merely a collective inconvenience; it is a mechanism through which productive time is systematically converted into economic waste at the metropolitan scale. When road networks are near or at capacity, the relationship between the number of vehicles and travel time becomes non-linear. This non-linearity is shown in speed-flow curves. As traffic volume approaches a road's capacity, small increases in demand lead to significantly larger decreases in speed and longer delays, a dynamic that intensifies as the network approaches saturation (Zefreh & Török, 2020).

The practical implication is that congestion does not scale linearly with the number of vehicles on the road. It compounds. A road carrying traffic well above its design capacity does not produce a proportional increase in delay; it can produce multiples of that figure, particularly when spillback propagates across adjacent junctions and upstream corridors (Afrin & Yodo, 2020; Tang & Heinemann, 2018). This compounding behavior is not incidental to the analysis that follows. It is central to it. If the relationship between vehicle entries and delay were proportional, a policy-induced reduction of several thousand daily accesses into a congested urban core would yield only a modest welfare gain. Because it is not, the same reduction can generate time savings, and therefore economic value, that are substantially larger than the raw entry figures suggest. Understanding congestion as a nonlinear, network-scale loss function rather than as a per-vehicle inconvenience is the conceptual foundation on which the economic valuation framework of this thesis is built.

1.1.2 Historical and Urban Evolution of Congestion

The history of urban traffic congestion is shaped by some pivotal stages that can be traced to the interrelated events of industrialization, post-war economic expansion, and suburban growth. Initially, industrialization drew significant populations to urban regions due to centralized economic activities within business districts (CBDs), thereby placing considerable pressure on transportation infrastructures (Gu, 2019).

Historically, manufacturing activities have been concentrated, often resulting in densely organized urban environments that allow residents to live near their workplaces (Bueno-Suárez & Coq-Huelva, 2020). In the years following World War II, Milan, like many other European cities, expanded as populations grew and incomes rose. This was accompanied by accelerated suburban development and a surge in car ownership, a phenomenon commonly known as Mass Motorization and Mass Transit, driven by the large production of affordable automobiles and government investment in highways (Paolini, 2009). Unfortunately, improvements in road infrastructure did not keep pace with suburban growth and the increasing number of vehicles, resulting in persistent traffic congestion (Wang & Debbage, 2021; Lu et al., 2021). This driving force led to suburban growth and urban sprawl, facilitated by permissive zoning and post-war highway construction, and characterized by single-use, low-density development that consistently increased reliance on automobiles at the city's edges. Prolonged commute distances and scattered origins and destinations heightened reliance on private vehicles, increased vehicle-kilometers traveled (VKT), and aggravated congestion across metropolitan road networks. While compact urban designs can reduce overall travel demand, they may also concentrate traffic in central areas, leading to localized congestion (Ewing et al., 2018). Subsequently, these urban centers experienced significant development as the focus shifted from traditional industries to the service, technology, and knowledge sectors. This shift has significantly altered the shape of urban infrastructure and the layout of mobility patterns. With the establishment of employment centers and an increase in socioeconomic diversity in residential areas, people generally experienced longer commutes as they traveled farther from their homes to various workplaces. Collectively, these historical dynamics established the structural underpinnings of modern congestion issues, in which congestion is more often managed through demand control and operational strategies rather than just expanding capacity (Qin et al., 2022).

1.1.3 Causes and Contributing Factors

Congestion results from the interaction of structural demand pressures and operational constraints. In empirical terms, this means that observed delays are shaped not only by the number of trips generated but also by where and when they concentrate, the available modal alternatives, and how efficiently the network is operated (Xiao et al., 2024).

For analytical clarity, the determinants can be grouped into demographic and land-use drivers; travel behavior and vehicle ownership; infrastructure and network constraints; and short-run disruptions (Wang & Debbage, 2021). Because the thesis later quantifies the economic loss from congestion, these factors are treated as channels through which time costs and unreliability arise: they either shift the demand curve for road space upward, reduce effective capacity, or increase the variance in travel times (Büchel & Corman, 2020).

1.1.3.1 Urbanization and Population Growth

Urbanization and population growth are both widely recognized as key contributors to traffic congestion, yet their relationship is not straightforward. City density, suburban sprawl, and commuter patterns all mediate how each process translates into actual congestion, and the mechanisms involved can differ considerably. Urbanization, broadly understood as the spatial concentration of population in cities and densely built-up areas, occurs through rural-to-urban migration, natural demographic increase, and the physical or administrative expansion of urban boundaries (Knorr et al., 2018). The scale of this shift in Europe is telling: according to the United Nations' World Urbanization Prospects 2025, 40.4% of the European population resided in cities in 2025, 32.6% in towns, and 27.0% in rural areas, with the urban share projected to reach 42.1% by 2050. Within Southern Europe, a similar dynamic is evident, although the settlement structure in Italy is more town-dominant, with 42.5% of its population residing in towns and 34.5% in cities in 2025. That spatial distribution matters because urban pressure is shaped not only by demographic growth, but by how population is organized

across cities, towns, and peri-urban areas — a configuration that directly conditions land use, commuting patterns, and transport demand (United Nations, 2025). In Milan, this mechanism is important to interpret carefully. Although the resident population has not expanded dramatically in recent years, with the municipality reporting an increase of about 4.2% from 2014 to 2024, congestion pressures remain substantial because mobility demand is generated by residents and non-residents alike.

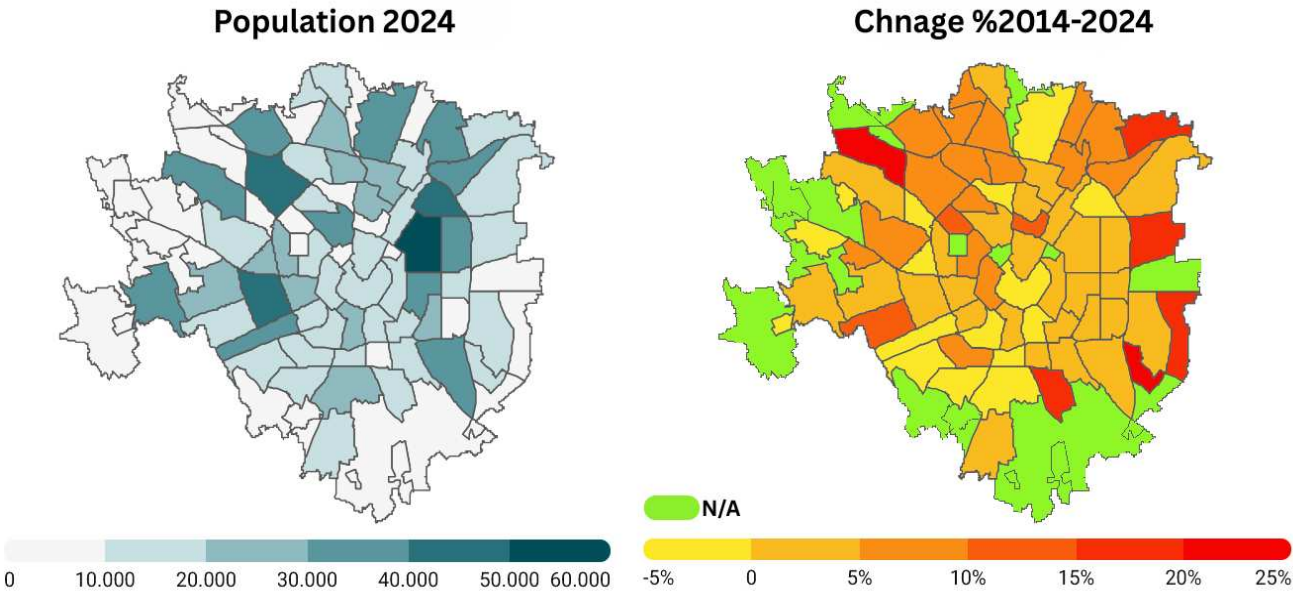


Figure 1-1. Left in Blue is the map of the density of Milan’s Population in 70 of 88 districts in 2024. On the right, in colors ranging from yellow to red, is a map showing how the population changed across those 70 districts over 10 years. Source: (Andamento Della Popolazione 2014 – 2024, n.d.)

Congestion is also shaped by where residents are spatially concentrated across the city’s districts, since population distribution determines where trips originate; however, the resulting congestion pressure depends equally on where employment and services are concentrated, which governs where and when trips converge (Wang & Debbage, 2021).

A key Milan-specific channel is the daily inflow of “city users” (work and study mobility), which increases the effective population present in the city and concentrates trips by time-of-day and location, thereby intensifying peak congestion even under modest resident growth. ISTAT’s daytime population evidence for Milan indicates very large incoming units and a coexistence index substantially above 100%, consistent with a city in which daytime

demand exceeds the resident baseline (Istituto Nazionale di Statistica (ISTAT), 2020).

In a long-run perspective, Milan's resident population today is also lower than its historical peak (early 1970s), which supports the point that current congestion cannot be interpreted as a simple function of resident population size alone (Comune di Milano, 2025).

While city density and traffic congestion do not exhibit a linear correlation, increased densification typically diminishes dependence on automobiles. It also promotes the utilization of public transportation if the enhancement of human density within current urban areas is strategically managed. Compact urban designs tend to reduce travel distances due to facilitating proximity, thereby alleviating congestion and promoting active mobility (Akbaridin et al., 2020).

On the other hand, the patterns of urban sprawl involve outward growth into lower-density, peripheral zones. This process increasingly consumes undeveloped areas as residents move into them. These areas are farther from employment and primary services, thereby increasing mobility needs and placing greater strain on transportation systems as the number of vehicles on the road rises, a trend most evident when infrastructure expansion does not keep pace with these developments (H. Lu et al., 2023).

In Milan, this aligns with the broader metropolitan pattern. Specifically, people living in peripheral districts and in the surrounding metropolitan area commute daily toward the core, increasing pressure on the main radial access routes and the ring-road system. This, in turn, worsens congestion and its unpredictability, particularly for work trips concentrated in peak hours (AMAT, 2022, 2023, 2024).

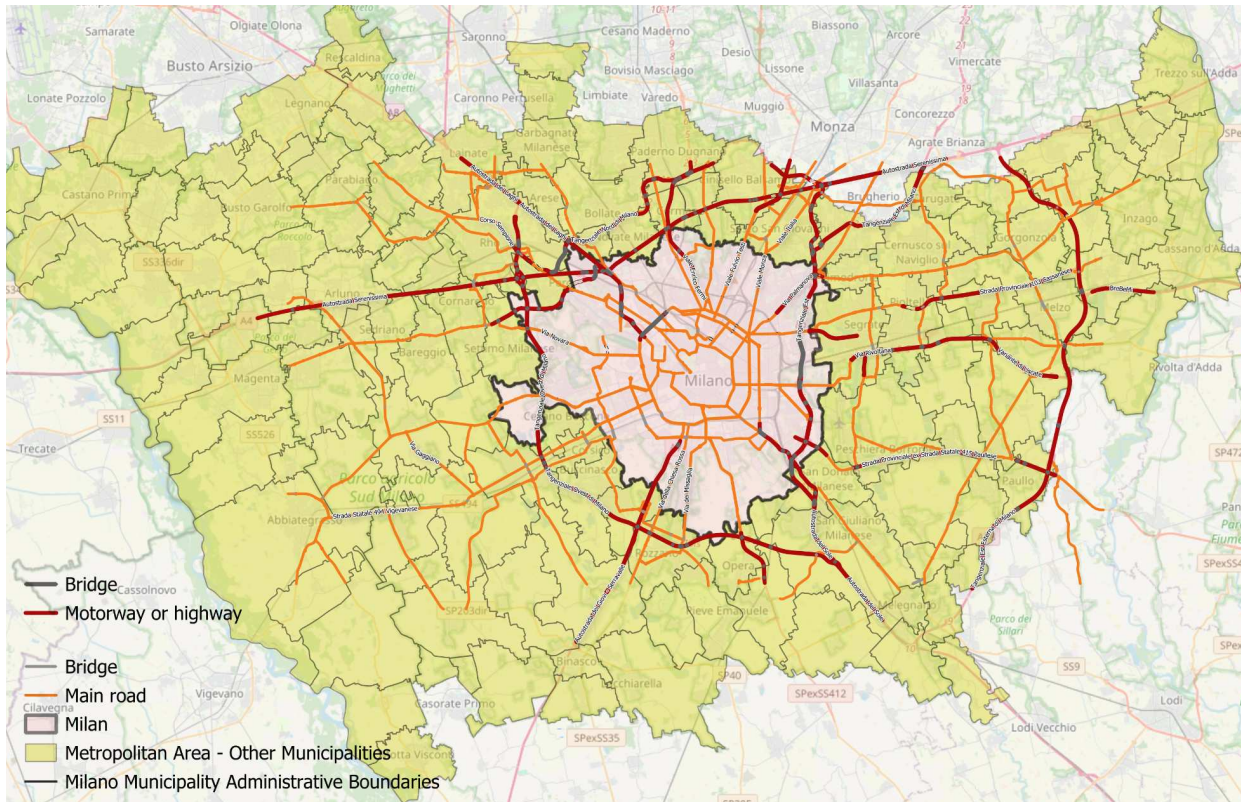


Figure 1-2. Milan municipality within the metropolitan catchment and primary road structure (main corridors/rings & main roads). The map highlights the main access corridors and ring-road structure (in red) and the broader set of main roads (in blue), providing the spatial basis for metropolitan commuting flows. (Source: Author's elaboration in QGIS based on Comune di Milano open Spatial data)

1.1.3.2 Rising Vehicle Ownership and Parking Accessibility

Along with urbanization and economic growth, vehicle ownership has emerged as a central mechanism of car dependence. The latest 2024 national estimate for Italy indicates approximately 700.8 passenger cars per 1,000 inhabitants, one of the highest motorization levels in Europe (Automobile Club d'Italia, 2025). Milan specifically reports an ownership level of approximately 0.49 cars per resident, confirming that the city's motorization rate, while lower than the national average, is still substantial and continues to exert pressure on the urban transport system (Marchetti & Antonelli, 2024).

A survey analysis of Milan shows that vehicle ownership is consistently the strongest predictor of a household's travel mode choice, outperforming travel distance, travel time, household size, and income across all comparative models. As the number of cars owned in a family increases, the likelihood of using public transportation decreases.

On the other hand, more and more people are opting to drive themselves, either alone or alongside public transportation (Liang et al., 2019). The convenience of parking increases the likelihood of vehicle ownership and the intention to purchase a vehicle. Parking density at both residences and workplaces is positively correlated with vehicle usage for commuting. When parking variables are included in the model, their presence clarifies the influence of built environment factors and improves model fit, indicating that parking oversupply masks or mediates the effects of density, transit access, and land-use mixture (Yin et al., 2018). These patterns are further amplified by urban spatial structure. Based on a comparative analysis of 19 European cities, Berrill et al. (2024) find that car use decreases exponentially with higher population density, and that distance to the city center is the most critical predictor of car ownership, car mode share, and longer trip distances. Low-density areas exhibit noticeably higher car ownership and much longer car trips, thus exerting disproportionate pressure on radial corridors. This is especially applicable to Milan, where the spatial heterogeneity between the dense urban core and the neighboring low-density towns is well established. Pucci (2021) shows that the surrounding Milan Urban Region is characterized by dispersed settlement patterns, multidirectional trip flows, and uneven public transport provision. More evidently seen in non-urban landscapes in municipalities where car dependency remains structurally strong. The same study reports that, while Milan itself exhibits a “very low motorization rate” relative to the surrounding region, 76% of work trips in the Milan Urban Region are made by car, clearly indicating a car-dependent dynamic developing beyond the city center.

1.1.3.3 Infrastructure, Network Constraints, and Short-Run Disruptions

Alongside demand-side drivers, congestion is reinforced by infrastructure and network constraints that limit effective capacity even when the physical road supply remains unchanged. In practice, bottlenecks arise from network design features such as lane drops,

constrained junction geometry, high signal density, and limited ability to adjust operations dynamically over time and space. These constraints create recurrent delay because the same locations repeatedly become capacity-limiting points during peak demand, and they also heighten the network's sensitivity to shocks (C. Zhang et al., 2019). Operationally, non-adaptive or weakly coordinated signal control and capacity-constrained intersections are particularly important because urban delay is often concentrated at junctions rather than along links (Chen et al., 2019). Infrastructure limitations also interact with disruption risks: events, construction, and temporary closures reduce overall capacity and can cause delays on connected routes, worsening congestion beyond the initial problem (Kukkapalli & Pulugurtha, 2020).

In dense urban areas, short-run disruptions also arise from the increased intensity of curbside activity. Deliveries, loading/unloading operations, and frequent stopping behavior reduce effective roadway capacity by creating localized friction, particularly near intersections and on narrow streets. In such contexts, even when total traffic volumes are stable, operational interference can increase queue formation and spillback, making congestion more persistent and more sensitive to minor disturbances (Janjevic & Winkenbach, 2020).

Economically, supply-side constraints turn demand spikes into longer delays, and disruptions that cut capacity increase travel-time variability, undermining reliability (Batley et al., 2019). In Milan, the mobility reporting framework shows congestion trends tied to specific times of day and road types, indicating that recurring bottlenecks and operational challenges shape daily delay patterns (AMAT, 2022; AMAT, 2023; AMAT, 2024).

1.1.3.4 Modal Imbalance and Evolving Demand Patterns

Congestion also reflects an imbalance in transportation modes. When private vehicles dominate travel needs and other options are not sufficiently attractive, peak demand concentrates on roadways. Even with public transport, issues such as unreliable service,

difficulties getting to and from stations, and poor integration between transport systems can prevent people from switching from cars. (Papadakis et al., 2024). From a mechanistic perspective, modal imbalance affects congestion by altering the flexibility of road demand: when public transport or active mobility options are perceived as less reliable or less accessible, households may maintain car ownership and continue to use private vehicles even as travel times rise. This is especially relevant in Milan's broader urban region, characterized by spread-out settlements and uneven public transportation coverage outside the city center (Pucci, 2021). Apart from mode choice, congestion is influenced by economic and behavioral patterns that concentrate travel during specific times. Factors like work-hour synchronization, school runs, and trip chaining bundle demand into short time periods, intensifying peak congestion even if overall daily travel stays roughly the same. This clustering of travel times increases peak road use and diminishes the ability of scheduling shifts to ease congestion (Yildirimoglu et al., 2021).

Rising mobility and logistics trends could exacerbate these existing pressures. Ride-hailing and other platform-based services, for instance, could lead to more vehicles on the road because of low-occupancy trips and the need to reposition them. Simultaneously, the expansion of e-commerce is driving up last-mile deliveries and service-related journeys.

While the operational impacts of curbside activity are discussed above, these trends are also relevant on the demand side because they add trips that are often highly time- and location-concentrated, contributing to higher peak-hour traffic volumes and, in dense areas, interacting with curbside friction to exacerbate delays (Janjevic & Winkenbach, 2020; G. S. Nair et al., 2020). Modal imbalance increases time costs because limited substitution keeps demand high during periods of high activity. It further renders less reliability, given that congestion is susceptible to small fluctuations or disruptions in demand. Such patterns are also important for policy because demand management measures are more effective when

viable modal alternatives exist (Romero et al., 2020). That is directly related to the later empirical design around central access regulation.

1.1.4 Economic Implications of Congestion

The congestion mechanisms described above ultimately matter because they convert everyday mobility into an economic loss function. When metropolitan road networks operate near capacity, congestion is not a proportional outcome: minor increases in demand or small operational frictions can trigger queue formation and spillback, producing delays that are disproportionate to the initial shock. Recent work that evaluates congestion through system cost metrics supports this non-linear interpretation, reinforcing that congestion should be treated as a recurring efficiency loss embedded in daily network functioning rather than a rare deviation from normal conditions (Serok et al., 2022).

The most immediate economic channel is time loss, but contemporary evidence suggests that a substantial share of the burden also comes from unreliability. In dense urban settings, congestion increases not only expected travel time but also the risk of late arrival, which leads travelers and firms to respond by adding buffer time, departing earlier, and reducing scheduling precision (Zang et al., 2022). Utilizing real-time congestion pricing data, Bento et al. (2024) demonstrated that the willingness-to-pay to circumvent congested alternatives is driven primarily by a discrete “value of urgency” (lateness penalties), rather than by marginal minutes saved, implying that congestion costs are closely associated with lateness risk and uncertainty, rather than merely average delay. Consistently, recent congestion externality estimates incorporate both expected travel time and reliability and find that reliability materially increases the measured welfare loss of congestion (Herzog, 2024).

As a result, these time and reliability losses lead to higher costs manifested in inefficient vehicle operations and reduced logistics productivity. Stop-go traffic raises fuel use and wear, but in urban economies, the more consequential effect is that congestion reduces the number

of activities that can be completed per hour, whether commuting, service visits, or deliveries. The value of travel time savings in freight transport is well-documented, providing a strong basis for interpreting congestion as a direct cost multiplier for logistics activity (Binsuwadan et al., 2022). In addition to operator costs, congestion can influence market outcomes: urban traffic time is associated with higher retail prices and lower consumer welfare, indicating that congestion costs are reflected not only in drivers' time loss but also in the prices households pay (Yilmazkuday, 2025). Within Milan, the AMAT mobility reports are useful as supporting local context rather than a primary evidence base: they operationalize congestion as additional time relative to free-flow conditions and document its concentration in peak windows, which is consistent with the academic emphasis on time-loss and reliability as the dominant economic channels (AMAT, 2022; AMAT, 2023; AMAT, 2024).

1.2 Smart Cities and IoT in Traffic Management

1.2.1 The Smart City Paradigm

Contemporary scholarship frames the smart city not as a technological end-state but as a governance orientation: an urban strategy in which digital infrastructure is deliberately integrated with institutional capacity and human behavior to generate measurable improvements in service delivery and resource allocation (Yigitcanlar et al., 2018). This distinction matters for the present thesis because it positions IoT not as a self-justifying investment but as a tool whose value depends on whether it produces behavioral and operational change. In the transport domain, that orientation is operationalized through Intelligent Transportation Systems (ITS), which couple information and communication technology with physical infrastructure, enabling the network to respond in real time to actual demand rather than to pre-programmed schedules (Ren et al., 2019). The ITS logic maps directly onto the congestion-economic loss pathway developed in Section 1.1: if congestion

arises because demand repeatedly exceeds capacity without any corrective feedback, then a system capable of sensing that imbalance and adjusting control parameters, or prices, in response is precisely the mechanism through which welfare losses can be reduced (Cheng et al., 2020). The technical foundation of modern ITS is the Internet of Things: an infrastructure of interconnected sensors, cameras, and edge computing devices that convert physical traffic states into continuous data streams (Cordoş et al., 2025; Oladimeji et al., 2023). What distinguishes IoT-enabled systems from earlier generations of traffic technology is not sensing alone, but the combination of real-time sensing with automated decision logic and persistent data archiving. This combination enables two analytically distinct functions. The first is operational control, in which the system adjusts signal timing, activates dynamic access restrictions, or triggers enforcement in response to observed conditions.

The second is policy evaluation, in which the same high-frequency records serve as the empirical basis for assessing whether a regulatory intervention has produced the intended behavioral change (Cheng et al., 2020; Ren et al., 2019). Both functions are present in Area C, and their coexistence makes the Milan case analytically tractable in ways that other European charging systems are not.

1.2.2 IoT in Urban Traffic Management: Mechanisms and Policy Relevance

The deployment of IoT in traffic management is most commonly discussed in terms of its operational outputs: adaptive signal control, real-time incident detection, smart parking guidance, and multimodal monitoring. Each of these mechanisms addresses a specific source of the delay and unreliability costs identified in Section 1.1.4. Adaptive signal systems reduce the recurrent delay concentrated at signalized intersections by realigning phase timing with actual traffic conditions rather than fixed-cycle plans, thereby directly reducing the stop-go friction that raises both average travel time and its variance (Li et al., 2023).

Parking guidance systems reduce search-related cruising, a source of congestion that is

disproportionate to its share of total traffic because it concentrates low-speed, high-frequency circulation in precisely the dense central areas where road capacity is most constrained (Shoup, 2018). Incident detection systems reduce the duration of the capacity drop by decreasing the detection-to-response interval for non-recurrent events, in which Section 1.1.3.3 identifies as a key amplifier of network sensitivity. An analytically important feature of IoT-enabled access-control systems is their data-generation function. ANPR-based systems not only support charging and enforcement but also generate structured, time-stamped records of vehicle movements that can be used for traffic analysis, flow estimation, and policy evaluation (Afandizadeh Zargari et al., 2021; Sánchez-Cambronero et al., 2021; Tang et al., 2022). In the Milan case, the Comune di Milano open-data platform publishes the Area C access dataset described as monthly records of total entries detected every 30 minutes by cameras located at the entry gates to the zone. This continuous administrative record is central to the empirical strategy of this thesis because it makes it possible to construct consistent before-and-after comparisons around policy changes, including the October 2023 tariff reform (Comune di Milano, 2025).

1.2.3 Milan's IoT and Smart Mobility Ecosystem

Milan's IoT traffic management ecosystem is relevant to this thesis not simply as background context but as evidence of institutional capability. A city that has sustained two decades of sensor deployment, data platform development, and iterative policy reform across multiple governance scales is one whose monitoring outputs can be treated as credible and consistently produced, rather than as artifacts of a single project or political cycle. That credibility underpins the empirical use of Area C gate data in Chapters 3 and 4. The foundational layer of this ecosystem is the Sharing Cities project (2016-2021), in which Milan served as an EU Horizon 2020 lighthouse city and deployed an integrated bundle of IoT sensing, data platform, and shared mobility infrastructure in the Porta Romana district. Critically, the model

demonstrated that IoT-enabled mobility governance could be institutionally anchored: its outcomes were subsequently incorporated into Milan's Sustainable Urban Mobility Plan for citywide scaling (Salvia & Morello, 2020). This institutional embedding is significant because it signals that Milan's smart mobility investments are governed by a SUMP framework of the kind described in Section 1.4.3, in which IoT tools are positioned as instruments for measurable mobility goals rather than as standalone technological experiments.

At the metropolitan scale, the Safety21 initiative (2020-2022) extended IoT monitoring and enforcement to provincial roads, deploying sensors explicitly designated for congestion data collection and integrating the outputs through a centralized cloud platform (Città Metropolitana di Milano, 2022; Greco, 2021). In shared mobility governance, AMAT's collaboration with Vianova from 2020 onward created a real-time integration layer over e-scooter, bike-sharing, and car-sharing fleets, operationalizing the kind of multimodal demand monitoring that the MaaS discussion in Section 1.4.3 treats as a precondition for effective mode-shift governance (Buxo, 2021). Most recently, the edge AI camera pilot on Viale Monza (2022-2024) addressed a specific gap in existing monitoring: the systematic under-representation of active mobility flows in ANPR and induction-loop data. By demonstrating that real-time multimodal classification is technically and ethically viable, the pilot lays the groundwork for adaptive signal control systems that incorporate the full demand picture that current Area C monitoring, limited to motorized-vehicle entries, cannot provide (AMAT, 2024b; Diomedede, 2024).

All considered, these initiatives situate the October 2023 tariff reform evaluated in this thesis not as an isolated decision but as one policy episode within a sustained institutional trajectory of technology-enabled demand management, coordinated by AMAT and grounded in a coherent planning framework.

1.3 Milan as A Case Study

1.3.1 Milan's Congestion Problem: Structural Pressures

Milan's congestion is not a temporary or cyclical problem; it is structurally produced by the interaction of metropolitan geography, urban form, and the absence, until the introduction of demand pricing, of any mechanism to internalize the external costs that each additional vehicle imposes on other road users. As Italy's economic capital and the functional core of a polycentric Lombardy region, Milan concentrates a volume and diversity of trip generation that its inherited road network was never designed to accommodate. The radial structure organized around the “Cerchia dei Bastioni” inner ring funnels metropolitan commuting demand from a catchment of approximately 9.5 million inhabitants through a limited set of access corridors, making peak-hour congestion both severe and highly predictable in its spatial distribution (Rotaris et al., 2010). The geography of the Po Valley compounds this problem by limiting atmospheric dispersion and promoting pollutant accumulation, so that congestion-related emissions carry an especially high environmental and public-health cost. In this context, demand-management policies have acquired a dual justification throughout their evolution: economic efficiency on the one hand, and environmental quality and reduced exposure on the other (Crova et al., 2024; Hart et al., 2024). The structural nature of Milan's congestion is also visible in the limits of supply-side responses. Successive metro extensions and surface transit improvements partially absorbed modal demand but could not resolve the core externality problem identified in Section 1.4.1: without pricing, every driver rationally discounts the delay they impose on others, and demand continues to press against the capacity of the radial access network during peak periods regardless of transit availability. This is precisely the mechanism that motivates the Pigouvian logic underpinning Area C and explains why Milan became one of the first Italian cities to shift from supply expansion to demand-side regulation as the primary instrument of congestion management.

1.3.2 From Ecopass to Area C and Area B: Milan's Demand-Management Policy Path

The policy lineage of Milan's traffic regulation instruments is analytically important because it illustrates a sequence of design failures and corrections that clarifies what a durable demand-management architecture requires. Ecopass, introduced in January 2008, was primarily conceived as a pollution charge: it applied a variable daily fee differentiated by Euro emission standard, with cleaner vehicles exempt (Percoco, 2017; Rotaris et al., 2010). The scheme achieved genuine early reductions in entries and congestion, but it contained a structural flaw that the externality framework of Section 1.4.1 makes clear. By pricing emissions rather than access, Ecopass incentivized the wrong behavioral response: drivers upgraded to cleaner vehicles and re-entered the zone, progressively expanding the exempt share from roughly 50% at inception to over 90% by 2010, thereby eroding both revenue and the demand-reduction effect without suppressing the underlying trip (Sevino, 2019). In effect, the policy targeted the vehicle rather than the access decision itself.

The city-wide referendum of June 2011, in which 79% of voters supported extending the scheme and using the revenues to strengthen public transport and sustainable mobility, resolved this design problem and gave the transition a democratic legitimacy that comparable reform attempts in Edinburgh and Manchester failed to achieve (Beria et al., 2017; Sevino, 2019). Area C, launched as a pilot in January 2012 and made permanent in March 2013, applied a flat daily charge to all vehicles irrespective of emission class, closing the exemption pathway that had undermined Ecopass and realigning the instrument with the Pigouvian logic of pricing access rather than emissions (Comune di Milano, 2013, as cited in Moulin & Urbano, 2025). The commitment to reinvest revenues in public transport strengthened the modal alternative available to drivers and therefore reinforced the durability of the pricing mechanism.

Beria, Tosi, and Nuccio (2017) show that the effects of this redesign were not temporary. Across the first four years of Area C, average daily entries during charged hours remained approximately 22 to 23 percent below the 2011 pre-Area C baseline, equivalent to roughly 33,000 fewer vehicles per day, with little sign of erosion over time. The same study reports associated emissions reductions within the zone, including about 18 percent for PM10, 18 percent for nitrogen oxides, and 35 percent for CO₂, alongside a marked shift in vehicle composition: paying vehicles fell from 62 percent of total entries in 2011 to around 45 percent by 2015, while public, authorised, and ecological categories came to represent the majority of access. These figures are directly relevant to the present thesis because they indicate that Area C had already reached a mature and relatively stable demand regime before the October 2023 tariff increase evaluated here.

In February 2019, Milan extended this logic through Area B, a Low Emission Zone covering approximately 132 km², or about 70% of the municipal territory, monitored by 185 ANPR gates operating under the same broader IoT architecture (Sevino, 2019). Whereas Area C manages the volume and cost of vehicle access to the historic core, Area B regulates fleet composition across most of the city by restricting the most polluting vehicle categories on weekdays from 07:30 to 19:30. Residents and registered firms are granted 25 annual derogation days before the restriction fully applies, subject to registration in the system (Direzione Mobilità e Trasporti - Comune di Milano, 2021). The measure is institutionally grounded in Milan's Sustainable Urban Mobility Plan and, as Sevino (2019) notes, drew explicitly on the experience of the Greater London Low Emission Zone.

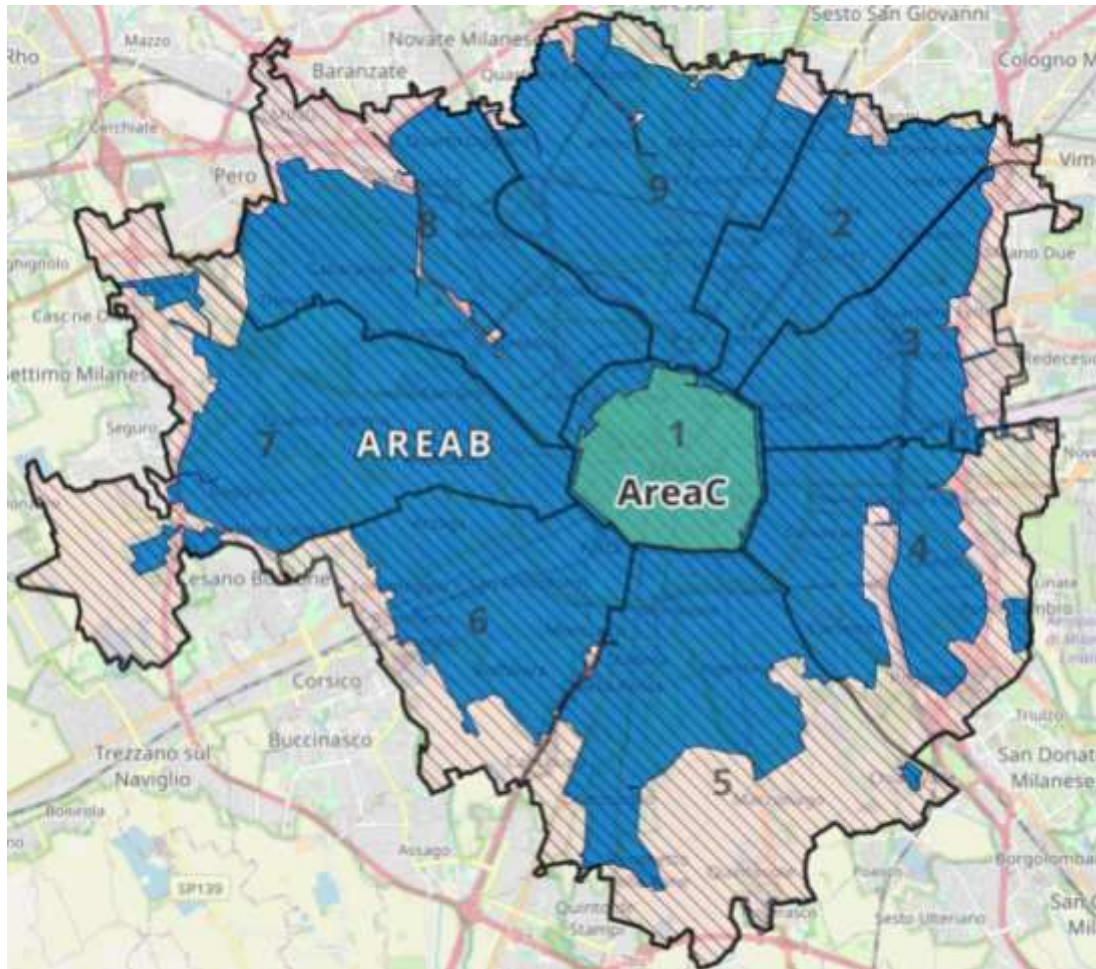


Figure 1-3. The administrative spatial borders of Area C and Area B in Milan, within its 9 zones
 (Source: Author's elaboration in QGIS based on Comune di Milano open Spatial data)

For this thesis, Area B's relevance is twofold: it confirms that electronically monitored access control is an institutionally settled governance technology in Milan rather than a novel experiment, and it affects the fleet composition approaching the Area C perimeter. The congestion-relief estimates in Chapter 3, therefore, operate on a vehicle mix already shaped by upstream environmental regulation. Following the pandemic suspension and reactivation of Area C in August 2023, two reforms in October 2023 provide the identifying policy variation evaluated in Chapters 3 and 4: a tightening of vehicle eligibility on 2 October and a tariff increase from €5 to €7.50 on 30 October (Moulin & Urbano, 2025). Because these changes occurred within the same ANPR-based monitoring infrastructure that has recorded access behavior continuously since 2012, they are amenable to the interrupted time-series and event-study strategy developed in Chapter 2.

1.3.3 Area C as a Smart Mobility Policy Case in Milan

Area C is best understood not as a pricing instrument that happens to use cameras, but as an IoT-enabled regulatory system in which the enforcement mechanism and the data-generation infrastructure are operationally unified. The scheme is enforced through 43 camera-equipped entry portals that use automated license plate recognition, with gate-level records generated as part of the charging and enforcement process (Gibson & Carnovale, 2015; Percoco, 2013).

This architecture instantiates the ITS logic of Section 1.2.1 at the regulatory perimeter of the urban core: distributed sensing, automated classification, rule-based control, and persistent data recording are integrated in a single operational system. As Beria et al. (2017) show, the system's architecture is not limited to plate detection itself: the OCR-generated transit record is cross-checked against exemption lists, payment tokens, and the national vehicle registry, while unmatched or non-compliant entries are transmitted into the fines process. This is analytically important because it clarifies that Area C's digital infrastructure is simultaneously a sensing system, an enforcement system, and a data-production system.

The analytical value of this architecture for the present thesis lies in the existence of high-frequency administrative records derived from gate-based ANPR monitoring. These records support daily aggregation across the gate network and make possible the interrupted time-series and event-study designs described in Chapter 2 (Gibson & Carnovale, 2015).

Compared with the congestion-pricing cases most commonly discussed in the literature, the Milan case offers an unusually tractable basis for before-and-after empirical evaluation because its charging system is institutionally mature, has been repeatedly reformed, and is observed through continuous administrative monitoring (Moulin & Urbano, 2025; Rotaris et al., 2010). Area C, therefore, offers a combination that is unusual in the empirical congestion-pricing literature: a long-run, high-frequency record of behavioral responses to a charging instrument whose design logic is theoretically grounded and institutionally mature.

1.4.1 Transport and Congestion Theories

Congestion as an externality and Pigouvian pricing. Urban road congestion is a textbook negative externality: each additional vehicle imposes delay costs on other users that are not internalized in private driving decisions, so unregulated demand exceeds the socially efficient level (Mosquera, 2024b). A Pigouvian access charge corrects this by pricing the marginal external cost of road use (Anas, 2020; Tikoudis, 2023). Area C is precisely such an instrument, and the October 2023 tariff increase is a deliberate upward revision of that correction. This logic connects all three analytical stages of the thesis: it establishes why the instrument exists, motivates the event-study identification of a demand response, and grounds the value-of-time welfare valuation in Chapter 3.

Downs-Thomson paradox. Where road capacity expands without a matching improvement in transit quality, mode shift back to driving erodes any speed advantage, potentially leaving generalized travel cost unchanged (Otsubo et al., 2025; Pandey & Lehe, 2024). This mechanism explains both why Milan's earlier supply-side expansions failed to resolve central congestion and why Area C's commitment to reinvesting charge revenues in public transport is not incidental but structurally necessary for the pricing instrument to remain effective over time.

Induced demand. The theory suggests that vehicle travel expands in response to added road capacity, implying that infrastructure expansion alone is unlikely to deliver durable congestion relief (Duranton & Turner, 2011; Ossokina et al., 2023). This principle directly informs the choice of case: the thesis evaluates a demand-management intervention rather than an infrastructure investment. The reason is that induced-demand dynamics render the latter strategy unreliable. It also helps interpret the persistence of the post-reform traffic reduction in Chapter 3.

1.4.2 Urban Systems and Complexity

Congestion is an emergent property of a complex network rather than a proportional outcome of a single cause: minor demand increases near capacity can trigger nonlinear queue formation and spillback (Saberli et al., 2020). This justifies the interrupted time-series design in Chapter 2, which treats the Area C entry series as a dynamic system requiring high-frequency longitudinal observation to detect structural breaks. IoT-driven traffic management adds a digital layer to this physical system: sensors, communications, and control logic interact continuously with road users, and the gate entry records produced by that digital layer are the precise, policy-responsive outcome variable used throughout the empirical analysis (Orozco et al., 2023; Zhu et al., 2020).

1.4.3 Smart City Governance Models

SUMP. The EU SUMP framework requires integrated, multi-modal, evidence-based mobility planning with explicit monitoring and evaluation (European Commission, 2019; European Union, 2023). Area C is embedded in Milan's SUMP cycle, which means the behavioral responses observed in Chapter 3 occur within a policy environment of simultaneous transit investment and shared-mobility governance, not in isolation from other demand-side initiatives.

TOD. It explains why congestion in Milan's corridors is spatially structured rather than random: land-use concentration channels metropolitan trip demand toward a constrained set of radial access points (Hrelja et al., 2020; Ibraeva et al., 2020; Liu et al., 2020). This supports the interpretation that a pricing intervention at the 43-gate Area C perimeter, where access is most spatially concentrated, produces a higher per-entry behavioral effect than a more dispersed instrument would.

MaaS. MaaS platforms increase the practical substitutability of car trips by integrating real-time multimodal information and booking (Kriswardhana & Esztergár-Kiss, 2023), but empirical evidence cautions that MaaS can also stimulate car use if bundles are not designed with explicit substitution goals (Alyavina et al., 2024; An & Shen, 2025; Storme et al., 2020). In the Milan context, MaaS availability is a moderating condition: it helps explain why more flexible payer categories exhibit stronger pricing responses in Chapter 3, whereas residents with more constrained mobility show weaker responses.

1.4.4 Planning Perspectives and the Economic Results in Context

The monetized congestion-relief estimates in Chapter 3 require interpretation through both urban planning and welfare-economics lenses. Four perspectives are especially relevant.

The 15-minute city and accessibility. The 15-minute city literature argues that sustainable urban mobility depends not only on regulating vehicle demand, but also on reorganizing urban form so that essential daily needs can be reached within short walking or cycling times, thereby reducing structural car dependence (Moreno et al., 2021; Pozoukidou & Chatziyiannaki, 2021). A related distinction is important here: access refers to the ability to reach the transport system or a destination, whereas accessibility concerns whether the reachable opportunities actually correspond to people's needs (Brussel et al., 2019). Applied to Area C, this means that the policy improves traffic conditions for continuing users, but does not necessarily improve accessibility for those who suppress trips because the charge has become burdensome. The welfare gains estimated in Chapter 3 should therefore be read as real but partial, capturing delay reduction within the existing urban structure rather than the fuller gains that would arise if Milan's urban form reduced car dependence more fundamentally. This is consistent with the broader sustainable-mobility literature, which emphasizes that pricing is most effective when the surrounding spatial environment already supports viable non-car alternatives (Holden et al., 2020).

TOD, park-and-ride, and spatial preconditions. A similar implication emerges from the transit-oriented development literature. TOD reduces car dependence by concentrating mixed uses, density, and accessibility around high-capacity transit nodes, thereby changing the relative attractiveness of driving rather than simply increasing its monetary cost (Thomas et al., 2018). Pricing and land-use strategy are therefore complementary rather than substitutable. Park-and-ride systems reinforce this complementarity by allowing outer-area commuters to transfer from car to transit before entering the congested core, reducing the number of trips that must be priced out at the charging perimeter itself (Li et al., 2021). For Milan, the implication is straightforward: if outer districts and interchange nodes offered broader and more reliable non-car alternatives, the congestion-relief and welfare effects of Area C would likely be larger, more persistent, and more evenly distributed.

SUMP and multi-dimensional evaluation. From a planning perspective, the results should also be situated within the logic of Sustainable Urban Mobility Planning. The European Commission's evaluation of the Urban Mobility Package makes clear that urban mobility policy should be assessed across multiple dimensions, including congestion, accessibility, affordability, modal split, safety, and environmental performance, rather than through a single indicator alone (European Commission, 2021). Recent work on sustainable urban mobility indicators reinforces this multidimensional logic by identifying congestion, affordability, energy efficiency, access to mobility services, and multimodality as core dimensions of policy evaluation (Chatziioannou et al., 2023). In this context, the interrupted time-series evidence developed in this thesis should be read as one analytically rigorous component of a broader evaluation framework: it identifies a behavioral response with unusual precision, but captures only one channel of policy performance. The broader implication is that Milan's future SUMP cycles could use the Area C data infrastructure not only to track traffic responses but also to embed them within a wider appraisal of accessibility, equity, and modal reorganization.

Social exclusion and distributional limits. A final planning perspective concerns transport-related social exclusion. Recent work emphasizes that exclusion arises not simply from fare or price levels, but from the interaction between cost burdens, limited modal alternatives, spatial mismatch, and unequal ability to reach socially valued opportunities (Bantis & Haworth, 2020; Luke, 2024). This perspective is particularly relevant for the payer heterogeneity documented in Chapter 3. The weaker response among resident payers is consistent with the idea that some users face tighter spatial or social constraints and therefore absorb higher charges rather than substituting away from the trip. The aggregate welfare estimates are therefore distributionally incomplete: they identify net delay-reduction benefits, but do not by themselves show who is least able to avoid the burden of the charge. Beria et al. (2017) qualify this point by showing that under the earlier Area C regime, the average monetary burden was concentrated in a relatively small share of users, with only about 4 percent of residents and 3 percent of other car users paying more than €1 per day on average. This does not remove the equity concern, but it clarifies it: the key issue is not whether all users are heavily burdened, but which groups are least able to substitute away from charged access.

Collectively, these perspectives do not challenge the empirical findings; they extend their meaning. The congestion-relief benefits identified in Chapter 3 are econometrically grounded and substantively important, but they remain partial in planning terms. Larger, more durable, and more equitable welfare gains would require the pricing instrument to operate within a complementary planning environment that also addresses proximity, modal alternatives, interchange design, and the unequal distribution of accessibility constraints across users.

1.4.5 Conceptual Impact Pathway

The theoretical strands above are integrated into a four-stage conceptual impact pathway that links the IoT-driven treatment to measurable outcomes and monetized value, and that maps directly onto the empirical chapters.

Stage 1: Treatment. Area C's ANPR gate network provides the IoT-driven smart traffic system. The October 2023 tariff increase is the specific, sharply timed policy shock applied through that system, observable in the administrative entry record.

Stage 2: Mechanism. The tariff raises the generalized cost of entry during active charging hours, prompting trip suppression, mode substitution, rescheduling, and destination changes. The instrument is demand-side: it does not alter road capacity but alters the incentive to consume it, which is the theoretically appropriate lever for an unpriced externality.

Stage 3: Outcome. Daily recorded entries into Area C during active charging periods, aggregated across the observed gate universe, are the primary outcome variable. The interrupted time-series and event-study designs in Chapter 2 identify structural breaks in this series around the October 2023 reform dates, controlling for weekday and seasonal structure.

Stage 4: Valuation. The estimated entry reduction is monetized using the value-of-time framework in Section 2.6: the reduction in daily entries multiplied by the assumed average time saved per avoided transit and by Italy-specific urban values of time. The result is a bounded partial welfare estimate of congestion-relief benefit, explicitly distinguished from a full cost-benefit analysis. The planning perspectives in Section 1.4.4 then situate this estimate within the broader accessibility and equity context that a welfare-economics lens alone does not capture.

Chapter 2: Research Methodology

This chapter sets out the methodological framework used for evaluating the role of an IoT-driven traffic-management system in reducing congestion-related economic losses in Milan. The thesis adopts a quantitative *ex post* policy-evaluation design centered on Area C, Milan's electronically monitored congestion-charging zone. The methodological task of the chapter is to explain how the research question is translated into an empirical strategy capable of identifying whether a clearly timed reform within this digitally monitored system was followed by a measurable change in traffic behavior, and how that change can be interpreted economically through a bounded valuation framework. This empirical strategy follows directly from the conceptual impact pathway developed in Section 1.4.5: the October 2023 tariff reform is treated as the policy intervention, the change in daily Area C entries as the observable behavioral outcome, and the value-of-time framework as the basis for translating that outcome into a bounded welfare interpretation. The chapter is organized around five main components. First, it defines the empirical setting, the scope of the analysis, and the main unit of observation. Second, it describes the data sources, variable construction, and processing workflow used to build the analytical dataset from the municipal files. Third, it presents the identification strategy — including the local event-study design, the longer-run filtered regressions, the placebo timing test, and the log-specification robustness check — alongside the descriptive analytical work that motivates those modeling choices. Fourth, it explains the value-of-time-based methodology used to monetize the traffic effect. Fifth, it discusses the methodological validity, limitations, and delimitations of the study.

The methodological logic is intentionally narrow and cumulative. The thesis does not attempt to estimate every possible effect of Area C, nor does it provide a full fiscal or social appraisal of the scheme. Instead, it focuses on the best-identified behavioral margin available in the data: daily recorded entries into Area C during active charging periods. The analysis then asks

whether these entries changed around the October 2023 reform period, whether any such change persisted beyond the immediate intervention window, and how the estimated reduction can be translated into a bounded economic interpretation. This approach is consistent with quasi-experimental evaluation principles for interventions introduced at clearly defined times and observed through repeated measures over time (Bernal et al., 2017; Miller, 2023). A further methodological principle guides the chapter. Wherever the administrative data support direct estimation, the analysis is explicit and formal. Where direct observation is unavailable, especially in the economic valuation stage, the thesis relies on transparent scenario assumptions rather than spurious precision. This combination of empirical discipline and bounded interpretation is central to the research design adopted here. **Figure 2.1** presents the conceptual framework underlying the research design, tracing the logical progression from the empirical setting through to the economic valuation outcome.



Figure 2-1. Conceptual framework of the research methodology. Source: Author's elaboration.

2.1 Case Setting, Scope, and Unit of Analysis

The empirical setting of the thesis is Area C in Milan. The choice of this case is methodological rather than illustrative. Area C provides a clearly delimited charging zone, a repeated administrative record of observed entries generated through electronic monitoring, and a sharply timed reform period in October 2023. These features make it suitable for a timing-based ex post evaluation using high-frequency observational data. The case is therefore selected because it offers the combination of policy timing, digital monitoring, and repeated outcome measurement required by the research design.

The analytical scope of the study is intentionally restricted to the monitored boundary of Area C rather than to the full Milan road network. This delimitation follows directly from the research question. The thesis examines whether an IoT-enabled traffic-management system can be linked to a measurable reduction in congestion-related economic losses. For that purpose, the most defensible observable outcome is the recorded number of entries into the charging zone while the system is active. A broader attempt to model all traffic conditions across Milan would extend beyond the structure of the available data and weaken the connection between the outcome variable and the monitored system under analysis.

The temporal scope covers thirty-six months, from December 2021 to November 2024. This period provides sufficient pre- and post-reform coverage to examine local dynamics around the October 2023 intervention period and to observe broader traffic patterns across multiple annual cycles. Within this horizon, two policy dates are relevant: the 2 October 2023 rules change and the 30 October 2023 tariff increase. Both dates are incorporated into the descriptive analysis. However, the tariff increase serves as the principal intervention date in the short-run econometric design because it represents the clearest price-based policy shock in the observed entry series.

The main unit of analysis is the daily recorded number of entries into Area C during active charging periods. Although the raw data are recorded at 30-minute intervals, the thesis aggregates them to the daily level at an early stage of data processing. This aggregation is methodologically appropriate because it reduces intraday volatility that is not central to the policy question, aligns with the weekly rhythm of urban mobility, supports the use of day-of-week controls, and yields coefficients that are directly interpretable as changes in daily access demand.

The observed gate universe used in the analysis consists of 40 gate IDs in the municipal count files, which differs from the 43 access points referenced in official Area C documentation. Because the analytical files do not contain a clean gate-type variable and the main empirical objective is to evaluate the observed entry series recorded in the open-data system, the thesis uses the full observed gate universe. This choice keeps the methodology aligned with the actual empirical record available in the data and is discussed further in the context of study limitations in Section 2.7.

2.2 Research Design and Identification Strategy

The thesis adopts a single-case quantitative evaluation design based on interrupted time-series logic and local event-study estimation. It does not use a treated-versus-control framework as its main design. This choice reflects the structure of the available data and the nature of the case. No untreated comparison area with sufficiently similar institutional conditions and equally rich public entry-count data was identified for the relevant period. By contrast, Area C offers a clearly timed reform, repeated high-frequency observations, and an outcome series generated by the monitored system itself. Interrupted time-series designs are well-suited to interventions introduced at clearly defined moments and observed through repeated measurements over time, particularly when randomized implementation is not feasible, and

the intervention affects an entire population or zone simultaneously (Bernal et al., 2017). In this thesis, that logic is applied to a monitored urban charging perimeter. The key empirical question is whether the entry series exhibits a meaningful shift at or around the reform date, once the regular temporal structure is accounted for.

The design also incorporates an event-study perspective. Event-study methods are useful for tracing the dynamic profile of response around a focal intervention, distinguishing immediate post-event movement from pre-existing patterns, and examining whether the timing of the response aligns with the policy shock rather than with unrelated broader fluctuations (Miller, 2023). In this thesis, the event study is deliberately local: it is centered on the tariff increase and estimated within a narrow symmetric window to keep the comparison as close as possible to the intervention of interest.

The identification strategy combines four layers. First, descriptive time-series analysis establishes the broader structure of the traffic series and positions the October 2023 reform dates within that pattern. Second, a local event-study around the 30 October 2023 tariff increase estimates the immediate response. Third, a longer-run filtered regression tests whether part of the response persists beyond the local shock window. Fourth, a placebo timing test and a log-specification robustness check strengthen the interpretation of the results. Each layer is described in detail in Section 2.5.

Before formal models are estimated, the daily entry series is examined using two descriptive time-series figures that inform the subsequent modeling choices. The first presents monthly averages of daily recorded entries over the full study horizon, revealing the seasonal structure of the series. In particular, it shows that August represents a recurring and pronounced trough in Area C traffic — an observation that directly informs the design of the filtered longer-run regressions, where all August dates are excluded to avoid confounding persistent post-shock effects with structurally atypical seasonal conditions.

The second figure zooms in on the period surrounding the October 2023 reform dates, marking both the October Second rules change and the 30 October tariff increase. This clarifies the relative timing and spacing of the two interventions and supports selecting the later date as the principal shock in the local event- study.

The resulting design should be understood as bounded causal inference. The thesis does not claim the certainty of an experiment, nor does it recover a fully observed, untreated counterfactual path for Milan's traffic. Rather, it evaluates whether a sharply timed policy reform coincided with a statistically meaningful and substantively interpretable change in the monitored entry series, whether that change remains visible under alternative specifications, and whether comparable patterns emerge under placebo timing conditions. In the context of this study, this is the appropriate scale of methodological claim.

2.3 Data Sources and Dataset Construction

The primary data source for this analysis is the Comune di Milano open-data platform (*dati.comune.milano.it*), which provides official administrative records of vehicle transits through the Area C congestion charging zone. Specifically, the dataset consists of monthly CSV files covering the period from December 2021 to November 2024, recorded at the 30-minute interval level by cameras installed at the entrance gates. The main raw variables used in the analytical workflow are *dataora*, *id_varco*, *areac*, *classe_areac*, and *numero_transiti*. This three-year window was deliberately constructed to place the two shock dates occurring in October 2023 approximately at the midpoint of the observation period, yielding roughly 23 months of pre-intervention data and 13 months of post-intervention data. This structure is both methodologically principled and empirically adequate.

The interrupted time series literature establishes that a minimum of 8 time points per phase is required for reliable estimation of level and slope parameters (Ewusie et al., 2020; Penfold &

Zhang, 2013), a threshold that this study far exceeds at the daily frequency. When observations are measured daily, even a single year yields over 250 usable data points per phase, providing substantial statistical power to identify a structural break.

The placement of the intervention near the center of the observation window further ensures that pre-intervention trend estimation is not compressed and that the post-intervention trajectory can be observed with sufficient resolution to distinguish an immediate level shift from a sustained slope change.

The variable `areac` is central to the empirical logic of the dataset. The main sample is restricted to observations where `areac == 1`, meaning the analysis focuses exclusively on traffic recorded while the Area C charging regime was active, rather than on all transits captured by the same infrastructure at all times. Since the research question concerns the behavioral response to the monitored charging system, this restriction is necessary to maintain coherence between the institutional setting and the estimated outcome.

The raw records are observed at the 30-minute level, but the thesis aggregates them to daily totals early in the processing stage. For each date, interval-level counts are summed across the observed gate universe to produce the main outcome series, `transits_total`. This aggregation is methodologically justified because the intervention of interest is defined at the daily level, the underlying weekly structure of urban mobility is strong, and the analysis focuses on changes in daily access demand rather than intraday micro-patterns.

In addition to the total daily series, class-level and payer-level daily datasets are constructed.

The variable `classe_areac` permits the identification of analytically relevant payer categories: class 3 corresponds to resident payers, class 4 to service payers, and class 5 to other payers.

A total payer series is then defined as the daily sum of classes 3–5, permitting the analysis of heterogeneity in behavioral response across payer groups rather than limiting the study to aggregate entry counts.

All stages of the empirical analysis, including data preparation, econometric estimation, and figure production, were conducted in R using a fully reproducible, scripted workflow in the RStudio integrated development environment. The principal libraries used were `data.table` and `lubridate`: the former for fast import, grouping, joining, and aggregation of repeated monthly administrative files, and the latter for consistent parsing and manipulation of time-indexed records. Every transformation that materially affects the estimand — including the `areac == 1` filter, the daily aggregation, the August and holiday exclusions, the event-study window, and the log transformation — is explicit in the code, making the empirical path from raw municipal files to final model output fully documented and replicable. Selected Milan mobility materials are additionally used as contextual supporting sources, though their role is descriptive and institutional rather than identificatory. The full analysis script is available upon request.

2.4 Econometric Models

2.4.1 Local Event-Study Around the 30 October 2023 Price Increase

The purpose of this first model is to test whether Area C entries changed sharply and immediately around the 30 October 2023 tariff increase, and to trace the short-run timing of that response as closely as possible. The main short-run specification is a **local event-study** centered on the **30 October 2023** tariff increase. In the R workflow, event time is defined as the number of days relative to the shock date, and these day-level distances are grouped into **7-day relative-time bins** over a symmetric **±35-day window**. The final pre-treatment weekly bin is used as the omitted category, denoted **bin -1**, and **day-of-week fixed effects** are included.

The model is:

$$Y_t = \alpha + \sum_{k \neq -1} \beta_k 1(kbin_t = k) + \sum_d \delta_d DOW_{dt} + \varepsilon_t$$

where Y_t is the daily recorded number of entries into Area C, $1(kbin_t = k)$ is an indicator for the event-time bin k , bin -1 is omitted, DOW_{dt} are day-of-week controls, and ε_t is the disturbance term. This specification is preferred for several reasons.

The local window keeps the comparison close to the intervention and reduces contamination from broader seasonal movement. Weekly binning smooths high-frequency noise while preserving short-run temporal precision. The choice of the final pre-treatment week as the baseline makes the coefficients directly interpretable as deviations from the last observed pre-shock traffic level. Day-of-week fixed effects are necessary because urban mobility follows a clear weekday rhythm that must be separated from any policy-related shifts. The event study also functions as a diagnostic device: pre-treatment coefficients allow inspection of whether the series was already moving sharply before the intervention, while post-treatment coefficients trace the short-run dynamic response — one of the main methodological advantages of event-study models (Miller, 2023).

2.4.2 Longer-Run Filtered Level Model

While the event-study identifies the immediate shock effect, the second model asks whether part of that response remained visible beyond the narrow reform window under more regular traffic conditions. The local event-study estimates the immediate effect of the price shock, but it does not by itself answer whether the response persisted beyond the short-run adjustment window. To address persistence, the thesis estimates a longer-run post-treatment model on the wider daily panel after applying a structural filtering rule. Specifically, all **dates in August** and all dates from **23 December to 6 January** are excluded. A post-treatment indicator is then defined for dates on or after **30 October 2023**, and **day-of-week fixed effects** are

included.

The model is:

$$Y_t = \alpha + \beta Post_t + \sum_d \delta_d DOW_{dt} + \varepsilon_t$$

where $Post_t$ equals one after the tariff increase. This model asks whether the daily entry level remained lower after the reform, once the most structurally atypical traffic periods are removed. The exclusion rule is methodologically important because August and the late-December/early-January holiday interval represent irregular traffic conditions that would make it harder to interpret the persistent effect as a regular post-reform adjustment.

2.4.3 Long-Run Log Specification

The third model serves as a functional-form check, testing whether the longer-run result remains substantively similar when traffic is modeled in approximate percentage rather than absolute terms. To assess whether the persistent result depends materially on the scale of the dependent variable, the thesis estimates a log-transformed version of the same filtered model. In the script, the dependent variable is defined as $\log(1 + transits_total)$, and the model includes the same post-treatment indicator and weekday controls.

$$\ln(1 + Y_t) = \alpha + \beta Post_t + \sum_d \delta_d DOW_{dt} + \varepsilon_t$$

This specification serves as a **functional-form robustness check**. It allows the persistent effect to be interpreted in approximate percentage terms and tests whether the substantive conclusion survives a change in outcome scale. Consistency across the level and log models strengthens the credibility of the finding.

2.4.4 Placebo Timing Test

The final model is a falsification exercise: it asks whether a comparable break appears around a non-policy date, which helps assess whether the main result could simply reflect recurring

seasonal movement rather than the October 2023 reform itself. The thesis also includes a placebo timing test. The script constructs a placebo model object, `m_placebo`, and reports its coefficient `post_fake` alongside the main estimates. This indicates that a parallel specification was applied to an alternative non-policy timing reference.

The placebo model is:

$$Y_t = \alpha + \beta Post_t^{fake} + \sum_d \delta_d DOW_{dt} + \varepsilon_t$$

The placebo does not conclusively prove causal identification. Rather, it evaluates whether a similar shock-like pattern appears when the same logic is applied to a non-policy period.

A weaker placebo effect strengthens confidence that the observed 2023 result is not simply a recurring calendar distortion.

It should be noted that a non-zero placebo coefficient does not invalidate the design; rather, it suggests that the placebo can be understood as a falsification device that only partially, not completely, eliminates recurring calendar effects. This limitation is covered in more detail in Section 4.2. In quasi-experimental research, this is a standard falsification logic rather than a definitive proof device (Bernal et al., 2017; Miller, 2023).

2.5 Economic Valuation Methodology

After estimating the traffic effects, the thesis translates the identified entry reduction into an economic interpretation through a value-of-time (VoT)-based partial welfare valuation of congestion relief. This stage is intentionally narrower than a full cost-benefit analysis. The thesis does not estimate the overall fiscal or welfare performance of Area C, nor does it calculate a net present value or benefit-cost ratio for the scheme as a whole. Instead, it monetizes a specific, empirically grounded benefit channel: the congestion-relief value associated with reduced traffic pressure at the charging perimeter.

A fuller CBA would require treatment-area-specific inputs on travel-time changes, environmental effects, noise, accidents, and operating costs, all data that were not publicly available, and would require, in the majority of them, an access request from parties of interest (administrative authorities, private consultancies, or relevant data providers), especially for the full 2021–2024 study period. A narrower but more robust welfare-based valuation is therefore preferred.

The value-of-time framework is the appropriate welfare concept for this context for a specific reason. Congestion pricing operates by internalizing the delay cost that each additional vehicle imposes on other road users, defining the gap between the private cost of a trip and its social marginal cost. As the Handbook establishes, the average travel cost borne by road users is derived from the product of the value of time and average travel time, making VoT the central monetization input for any congestion-relief valuation (Directorate-General for Mobility and Transport, 2020, p. 105).

When a tariff increase reduces entries, it compresses the average delay experienced by remaining users, and the welfare gain from that compression is measured by multiplying the avoided transit volume by the assumed time saving and the appropriate VoT. The Vademecum's broader framework for social welfare appraisal — which treats reduced travel time as a direct benefit to society rather than merely a fiscal outcome — reinforces the appropriateness of this approach for a partial but grounded policy evaluation (European Commission & JASPERS, 2021).

The valuation formula is:

$$Benefit = \Delta Q \times \Delta t \times VoT$$

where ΔQ is the estimated reduction in daily recorded entries, Δt is the assumed average time saving per avoided entry, and VoT is the value of time.

Two treatment effects are monetized: the short-run shock effect from the event-study and the

longer-run persistent effect from the filtered post-shock model. Because the exact trip-purpose composition of Area C entrants cannot be observed directly from the gate-count data, the analysis uses Italy-specific short-distance urban car values from the Handbook and applies a midpoint central scenario, along with conservative and upper-bound time-saving assumptions. This scenario design defines a bounded, transparent welfare range rather than a spurious point estimate. The full monetization results and scenario grid are reported in Section 3.6.

2.6 Methodological Validity, Limitations, and Delimitations

The validity of the empirical strategy rests on several assumptions that should be stated clearly. The first is that the 30 October 2023 tariff increase represents a sufficiently clear and behaviorally meaningful intervention date for identifying a local structural change in the entry series. This assumption is supported by the institutional logic of congestion pricing and by the descriptive behavior of the data.

The second concerns the choice of the final pre-treatment weekly bin as the omitted baseline in the event study, which ensures post-treatment coefficients are interpreted relative to the most relevant immediately preceding traffic conditions rather than a distant average where seasonality matters.

The third concerns the use of day-of-week fixed effects, which are necessary to prevent routine weekly variation in urban mobility from being mistaken for a policy response.

Validity is further strengthened through model triangulation. The event-study establishes the timing and shape of the immediate response. The longer-run level model evaluates persistence under regular traffic conditions. The log model checks functional-form sensitivity. The placebo test examines whether similar patterns emerge without the true intervention.

However, unobserved confounders could still influence results, and discussing their potential impact would clarify the robustness of the findings.

Limitations

The methodology has clear limits that define the range within which its findings can be responsibly interpreted.

Absence of an external control group. Because no untreated comparison area with sufficiently similar institutional conditions and equally rich public entry-count data was identified, the empirical strategy relies on the internal temporal structure of the Area C series.

The thesis, therefore, identifies a temporally concentrated change consistent with a policy-induced behavioral response, but it does not directly observe a perfect untreated counterfactual. Future research could strengthen external validity by identifying a comparable urban setting with similarly detailed open-entry count data.

Outcome variable scope. The dataset records entries into Area C, not the full downstream equilibrium of traffic conditions across Milan. While the thesis identifies an access-response effect at the charging perimeter, its limited scope precludes definitive attribution of changes to specific behavioral mechanisms such as mode shift or rerouting. Clarifying this helps readers understand the limits of causal inference regarding behavioral channels.

Gate universe. The municipal count files consistently list 40 observed gate IDs, rather than the 43 access points referenced in the official Area C documentation. Because the analytical files do not allow a clean reconstruction of all gate functions, the estimates apply to the empirically recorded monitored system as represented in the open-data files rather than to a fully reconstructed official gate taxonomy. A later extension could test whether results are materially unchanged when identifiable special-access or public-transport-specific gates are excluded or treated separately.

Partial economic valuation. The monetized congestion-relief estimates are deliberately partial. A continuous Area C-specific travel-time panel was not available for the full 2021–2024 study period, so the value-of-time exercise relies on transparent scenario assumptions about

minutes saved per avoided transit rather than directly observed time savings. The resulting welfare range should be read as a bounded estimate rather than an exact social total. This is a methodological delimitation designed to keep the valuation aligned with the evidentiary strength of the available data. Future work could improve precision by combining gate-based identification with treatment-area-specific speed or travel-time data.

Fiscal balance and environmental co-benefits were not estimated. The thesis does not estimate the fiscal balance of the October 2023 tariff increase, nor does it monetize environmental co-benefits. Both omissions are deliberate: the empirical design is strongest for identifying the access-response effect and translating it into a congestion-relief welfare estimate; it is not built around a full financial or environmental appraisal. A future extension could pair the welfare valuation developed here with a fiscal module and an emissions module if comparable annual data become available.

External Validity

Finally, the thesis has limited external validity. Milan is analytically informative because it combines a mature congestion-pricing regime, electronically monitored access control, and repeated administrative traffic data. However, not all cities share the same institutional design, policy history, or monitoring infrastructure. The study, therefore, aims for analytical rather than universal generalization: it tests whether an IoT-enabled demand-management system can generate an identifiable, measurable, and economically interpretable behavioral response under real urban conditions. Chapter 3 turns from the design itself to the evidence produced by that system, beginning with the descriptive structure of the Area C entry series and the positioning of the October 2023 reforms within it.

Chapter 3: Data Analysis and Impact Assessment

3.1 Area C as the Empirical IoT Traffic System Under Study

The empirical backbone of this thesis is the Area C gate system in central Milan.

As discussed in the preceding chapters, Area C is not only a congestion-charge regime but also an IoT-enabled traffic-monitoring infrastructure. Each access point is equipped with camera-based detection technology that records vehicle transits and classifies them within a structured administrative system. This makes Area C especially suitable for empirical analysis, because the policy generates a high-frequency, machine-recorded stream of mobility data rather than relying on occasional surveys, ad hoc observation, or irregular manual counts.

From an analytical perspective, the importance of this infrastructure lies not simply in digitalization itself but in the integration of sensing, recording, classification, and rule-based activation. The gate network continuously monitors movement along the entry perimeter of the charging zone, while the administrative data system identifies whether the charging regime was operative at the time of observation. In this sense, Area C functions as more than a toll boundary. It is a digitally mediated regulatory architecture in which measurement and enforcement are embedded in the same system. This is the sense in which the case is treated here as an IoT-driven smart traffic system: the infrastructure does not merely support regulation but also produces the evidence through which the effects of regulation can be evaluated.

The key variable used in the raw files is the `areac` field. As noted in the descriptions of the municipality's datasets, when `areac = 1`, the dataset indicates that Area C was active at the moment of monitoring; when `areac = 0`, the same gate system was recording traffic outside the active charging regime. This distinction is important because it allows the infrastructure to function as both an access-control system and a measurement system.

In other words, the same IoT network can distinguish between traffic recorded during the operative charging period and traffic recorded outside it. For this reason, the core empirical analysis focuses on observations collected during Area C's active period, since these are the periods when the charging regime can plausibly shape behavior. Table 3.1 illustrates this monitoring structure by comparing average daily transits recorded during Area C's active and inactive periods across three selected months of 2023.

Table 3.1. IoT monitoring box (*areac=1 or 0*)
 (Source: Comune di Milano open-data platform, generated using RStudio)

Month	Avg daily transits when Area C is inactive	Avg daily transits when Area C is active	Share of observations with areac=1
May-23	56748	107455	0.493
August-23	28407	49239	0.495
October-23	59648	105134	0.486

Table 3.1 is not intended as a causal comparison between active and inactive periods. Rather, it demonstrates the logic of the system’s data architecture. The gate network continuously records mobility while also flagging whether the charging regime is operational, thereby allowing the same infrastructure to serve as both a regulatory boundary and a measurement device. This feature is central to the thesis question because the study is concerned not with a generic traffic policy in Milan but with the role of an IoT-driven smart traffic system in reducing economic losses from congestion. The empirical contribution of the thesis depends precisely on the fact that Area C produces a structured and temporally fine-grained administrative record of behavior at the charging perimeter.

The broader institutional context also supports this interpretation. The Milan mobility reporting system describes Area C in 2023 as operating through 40 gates over 247 active

days, with an average of 74,673 daily transits during active charging hours (AMAT, 2023). Official Area C documentation also indicates that some gates are reserved for public transport or other special-access uses. However, the analytical count files do not contain a direct gate-type variable, and some special-access locations appear merged or absent in the observable gate-ID universe. The main analysis, therefore, is conducted on the full observed aggregate of operational gate IDs. Establishing Area C as an IoT-based monitoring and access-control system is important for the remainder of the chapter. Once the logic of the data-generating infrastructure is clear, the next step is to examine how traffic evolved over time within that system and to locate the October 2023 policy changes within the broader temporal pattern of observed transits.

3.2 Descriptive traffic patterns and policy timing

The analysis uses thirty-six months of daily Area C gate-count data spanning December 2021 to November 2024. Aggregating the raw 30-minute observations to the daily level reveals a clear and recurring seasonal structure in traffic within the zone.

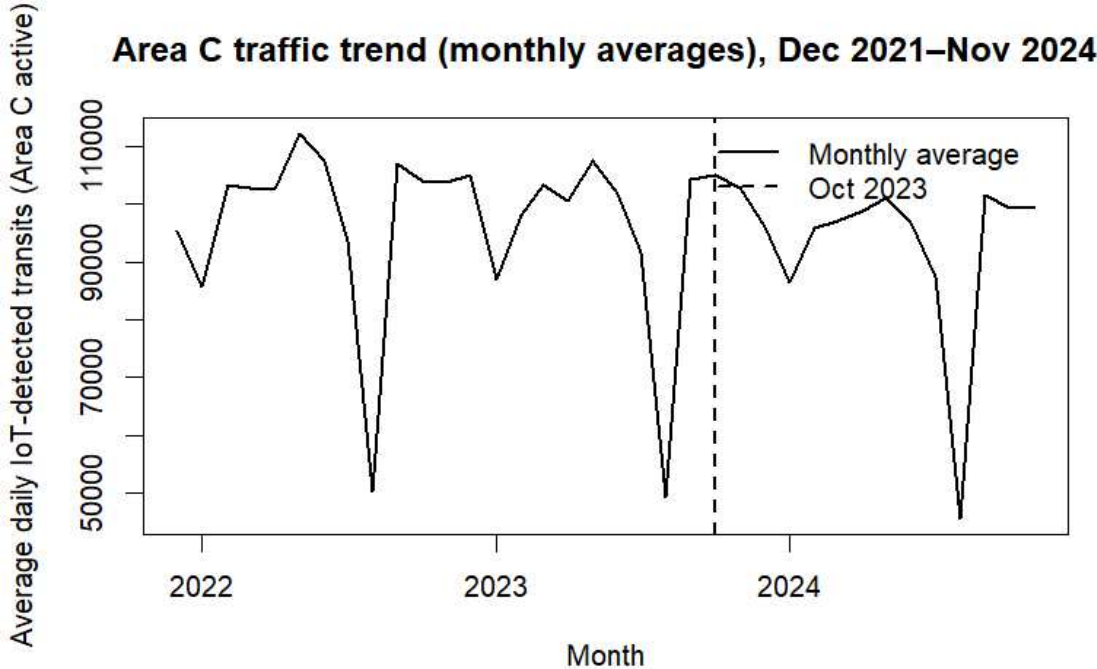


Figure 3-1. Monthly Average Area C traffic, Dec 2021–Nov 2024
 (Source: Author’s calculations based on Area C gate data)

Figure 3.1 reports monthly averages of daily Area C transits and shows that mobility in central Milan is not stable throughout the year. The most prominent recurring feature is the strong August trough, consistent with Italy's seasonal holiday pattern. Outside that period, average daily traffic tends to recover and stabilize at markedly higher levels.

This descriptive evidence matters for two reasons. First, it confirms that simple before-versus-after comparisons over long horizons would be misleading without controlling for seasonality. Second, it motivates the use of local policy windows, in which the objective is to isolate the immediate behavioral response to a discrete intervention rather than to compare across widely different times of the year.

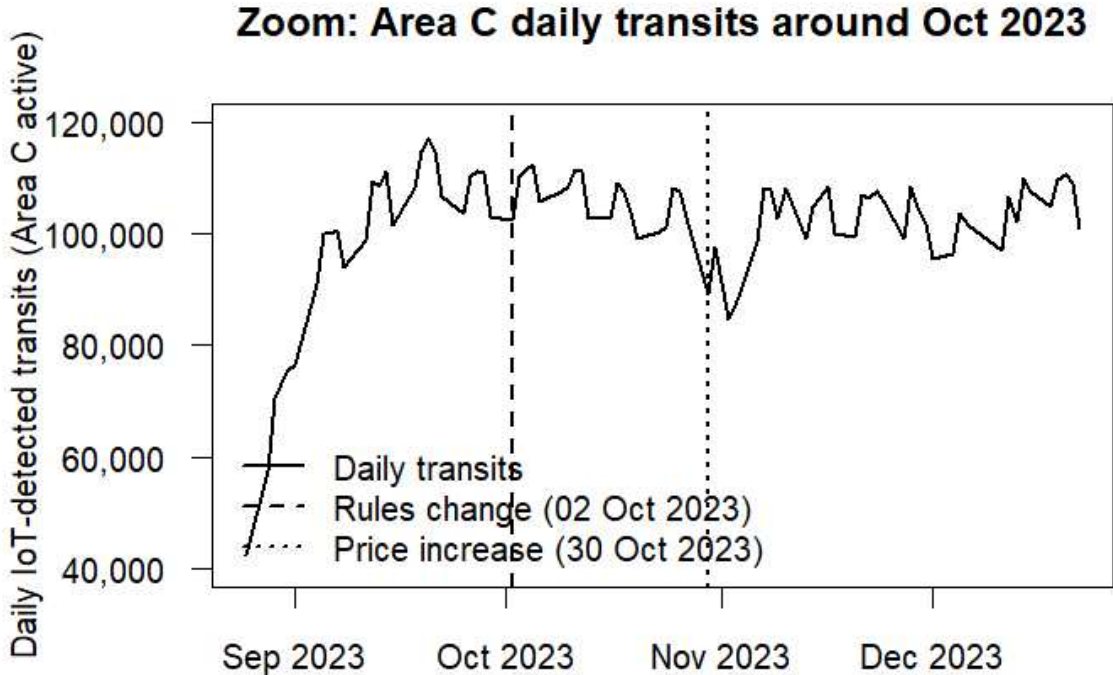


Figure 3-2. Zoom of daily Area C transits around October 2023
 (Source: Author's calculations based on Area C gate data)

Figure 3.2 narrows the focus to the period from late August to late December 2023 and marks the two main policy changes examined in this chapter: the **2 October 2023 rules change** and the **30 October 2023 price increase**. The visual pattern suggests that the first intervention did not produce a sharp and obvious break in aggregate traffic, whereas the second is followed by a more noticeable decline. This visual impression is not treated as causal proof on its own, but

it provides a clear motivation for centering the formal empirical strategy on the 30 October tariff increase. In other words, the descriptive evidence suggests that the price increase is the more plausible candidate for an immediate impact on traffic. The descriptive evidence is also consistent with the broader Milan mobility context. The 2023 mobility report characterizes congestion in Milan using TomTom-based indicators and notes that weekday congestion exhibits a clear peak-hour structure, while the average peak congestion level in 2023 increased relative to 2022 (AMAT, 2023). This broader city-level evidence supports the idea that reducing entries into a dense central area can plausibly generate time savings and congestion relief, even though the core empirical analysis here relies on Area C gate counts rather than a citywide speed series. Both the monthly trend and the local October 2023 zoom suggest that the tariff increase is the more likely source of a discrete traffic response. The next section formalizes this intuition through a local event-study / interrupted time-series design.

3.3 Empirical Design

The empirical strategy of this thesis is designed around the strengths and limits of the available Area C data. Rather than relying on a treated-versus-control comparison, the analysis uses a local event-study and interrupted time-series logic centered on the 30 October 2023 tariff increase, complemented by descriptive analysis and longer-run robustness checks. This design was chosen because no comparable untreated area with equally rich public data was identified. In contrast, the Area C gate system provides sharp policy data and a high-frequency internal record of observed entries. In this context, a local timing-based design offers the most credible way to evaluate whether traffic changed discontinuously and persistently around the reform.

The empirical analysis proceeds in four steps. First, the chapter presents descriptive time-series evidence using monthly averages and a daily zoom around the October 2023 policy period. This step is intended to place the reform within the broader temporal structure of Area C traffic and to distinguish the late-October tariff increase from other seasonal patterns. Second, the main short-run estimate is obtained from a local event-study around 30 October 2023, using weekly bins with a symmetric ± 35 -day window.

Third, a longer-run post-treatment specification is estimated on the wider daily series after excluding periods that are structurally atypical in Milan's traffic calendar. Fourth, additional validation is provided through a log-transformed specification and a placebo exercise around October 2022. The 30 October 2023 tariff increase is treated as the principal shock for identification purposes.

Although the October 2023 reform period also included a rules change on 2 October, the later tariff increase provides the cleaner policy break because it generated a clearer behavioral response in the data and is more straightforward to interpret as a price signal. As a result, the event study is centered on 30 October rather than on the earlier administrative adjustment. This choice does not deny the relevance of the earlier change, but it reflects the empirical objective of isolating the most identifiable discontinuity in observed entries.

For the main short-run analysis, the dependent variable is the daily number of Area C transits during active charging periods. The event-study groups days into weekly bins and omits the final pre-treatment bin as the reference category. Day-of-week fixed effects are included because traffic within Area C follows a clear weekly rhythm, and failing to account for this pattern risks conflating the policy shock with routine variation across weekdays.

Weekly binning is preferred to a fully day-by-day specification because it smooths excessive short-term noise while still preserving close temporal proximity to the reform.

The resulting coefficients thus describe how average daily transits in each relative week differ from the last pre-treatment week. The longer-run specification serves a different purpose. While the event-study captures the immediate response to the tariff increase, the wider daily model assesses whether part of that response remained visible after the initial shock period had passed. To reduce the influence of structurally unusual traffic periods, the specification excludes all August observations as well as the period from 23 December to 6 January. These exclusions reflect the descriptive evidence presented earlier in the chapter: August is characterized by a pronounced seasonal trough, while the late-December and early-January interval is affected by holiday-related disruptions that would make the post-treatment coefficient harder to interpret as a regular behavioral adjustment. For this reason, the longer-run model provides a cleaner estimate of persistence under more typical traffic conditions.

The design does not claim the certainty of a laboratory experiment, nor does it eliminate every possible alternative explanation that could coincide with the reform window. Its value lies elsewhere. By exploiting a sharply timed tariff increase within a high-frequency administrative dataset, the strategy asks whether Area C entries changed in a temporally concentrated and statistically meaningful way around the policy shock, and whether that change remained detectable under alternative specifications. In a case where rich internal measurement exists but a credible external control group does not, this form of bounded causal inference is both empirically disciplined and substantively appropriate.

3.4 Main Empirical Results

The first set of results concerns the aggregate response of Area C traffic to **the 30 October 2023** price increase.

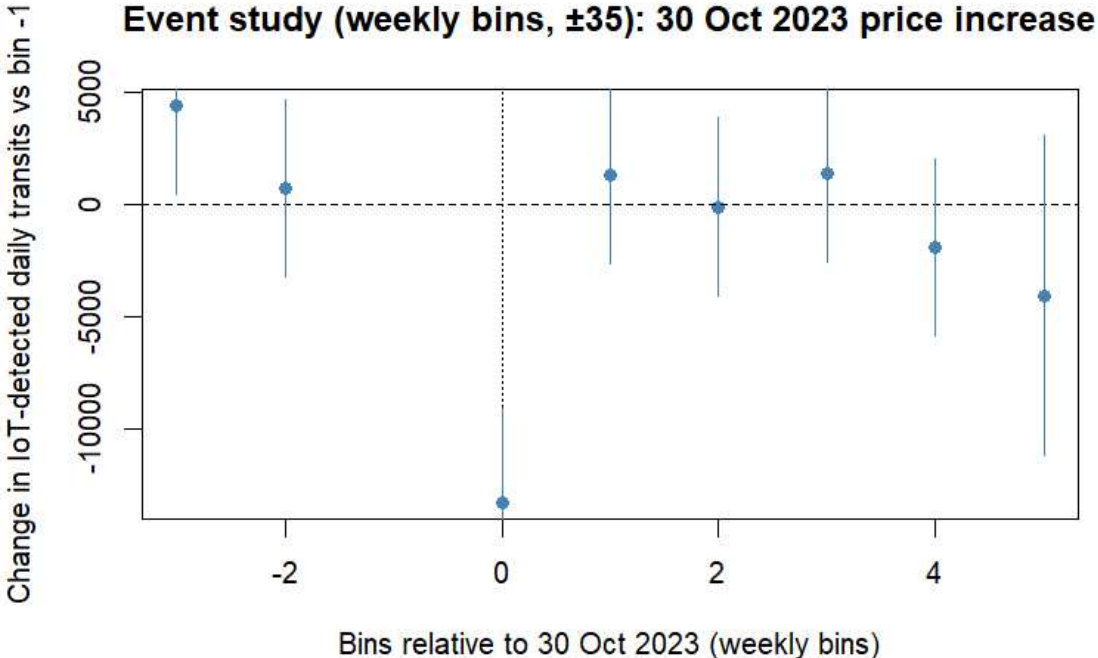


Figure 3-3. Weekly-binned event study for total transits around 30 Oct 2023 (Source: Author’s calculations based on Area C gate data)

Figure 3.3 presents the weekly-binned event-study estimates for total daily transits within a ± 35 -day window around the tariff change. The x-axis shows weekly bins relative to 30 Oct 2023, showing the estimated change in daily IoT-detected transits relative to bin -1. As: 0 = the week containing the price increase, -1 = the omitted reference week just before the shock, -2, -3 = earlier pre-treatment weeks, 1, 2, 3, etc. = post-treatment weeks.

The most important feature of the figure is the shock-week estimate, which is sharply negative relative to the omitted pre-treatment week. This visual pattern indicates an immediate and substantial reduction in traffic in Area C following the tariff increase.

*Table 3.2. Summary Statistics of Traffic Data for Area C
(Source: Author's calculations based on Area C gate data)*

Result	Estimate	p-value	Baseline	% change
Short-run price shock (30 Oct 2023) - Total transits	-13314.229	<0.001	103858.6	-12.82
Short-run price shock (30 Oct 2023) - Payer transits	-3393.95	0.002	40882.4	-8.3
Price shock - Resident payers (class 3)	-186.429	0.297		
Price shock - Service payers (class 4)	-224.964	0.005		
Price shock - Other payers (class 5)	-2982.557	0.001		
Long-run post-price effect	-4613.073	<0.001		
Placebo window (Oct 2022)	-3984.215	0.036		
Long-run post-price effect (log specification)	-0.045	<0.001		-4.41

Table 3.2 summarizes the corresponding estimate: the short-run effect implies a reduction of approximately **13,314 daily transits**, equivalent to about **12.8%** relative to the pre-treatment baseline. The estimate is a large effect for a mature congestion-charge regime, indicating that the October 2023 tariff increase was not merely symbolic but behaviorally consequential; it had a measurable, immediate effect on traffic entering the central zone.

The analysis then turns from total traffic to the composition of the response. A decline in overall entries can arise either because all users reduce traffic proportionally or because specific driver categories are more price-sensitive than others.

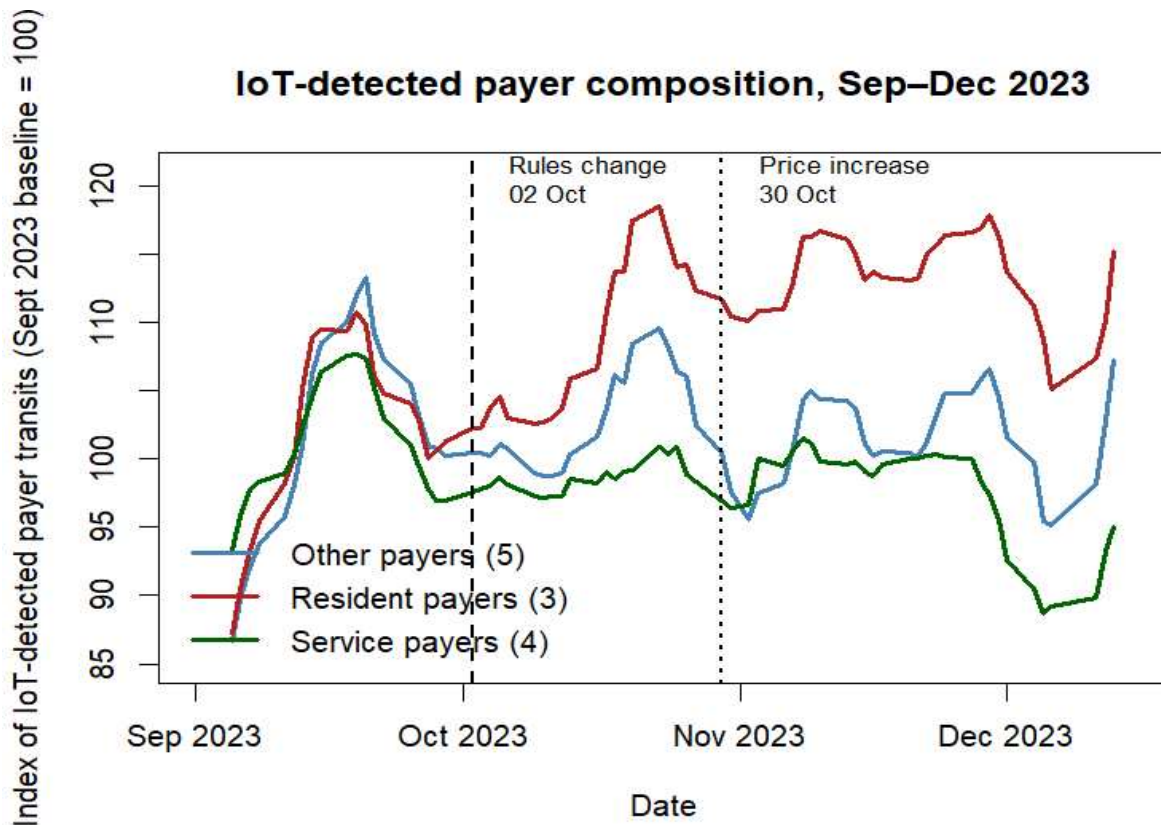


Figure 3-4. Payer composition around October 2023
(Source: Author's calculations based on Area C gate data)

Figure 3.4 addresses this question by plotting indexed payer series for three groups: resident payers, service payers, and other payers. The figure shows that the three categories do not move in the same way after the October 2023 changes. Resident payers remain relatively elevated compared with their September baseline, whereas other payers and service payers weaken more noticeably after the price increase.

The regression results in Table 3.2 reinforce this visual pattern. The short-run payer effect is approximately **-3,394 daily transits**, equivalent to about **-8.3%** relative to the pre-shock payer baseline. Within this broader payer response, the largest and most statistically significant adjustment was observed among **other payers**, while **service payers** also decline significantly. By contrast, the effect for **resident payers** is not statistically significant.

It indicates that residents, who are more likely to have regular, restricted, or permit-related access needs, were less receptive to the tariff increase than users with more flexibility or price

sensitivity. This mechanism strengthens the interpretation of the aggregate traffic effect. The October 2023 tariff increase did not simply coincide with a random decline in traffic; it appears to have changed the behavior of the categories of drivers most likely to respond to a stronger cost signal. Aligning with the broader composition evidence available for area C in the 2024 mobility report, which indicates that more than 85% of area C entries are made by vehicles for passenger transport, while only 1.3% are buses.

This correlation helps explain why the aggregate pricing response is best interpreted as a change in ordinary urban vehicle behavior rather than as a shift driven by public-transport-specific gates or bus flows. The aggregate and composition results suggest that the October 2023 tariff increase induced a genuine and economically meaningful reduction in Area C traffic, especially among the groups most likely to be price-sensitive. The next section tests whether these results remain credible under alternative specifications.

3.5 Robustness and Validation

The main event study identifies a substantial short-run decline in Area C traffic following the 30 October 2023 tariff increase. The next question is whether this effect was confined to the immediate implementation window or whether it persisted beyond the initial shock. To address that issue, the chapter estimates a longer-run post-treatment model on the wider daily series while excluding August and the period from 23 December to 6 January. These exclusions remove the most structurally atypical traffic periods in the Milan calendar and therefore reduce the risk that the post-treatment estimate is driven by seasonal distortions rather than by a more durable behavioral adjustment.

The longer-run estimate remains negative and strongly statistically significant. As reported in Table 3.2, the post-price effect in the filtered specification is approximately $-4,613$ daily transits. Indicating that the October 2023 tariff increase did not produce only a brief

implementation shock. Even after the most immediate response had passed, observed entries into Area C remained below their earlier level. The longer-run model, therefore, supports the interpretation that at least part of the reaction to the tariff increase reflected a more persistent change in access demand.

A second validation step concerns functional form. A log-transformed version of the same longer-run specification yields a coefficient of approximately **-0.045**, corresponding to an implied reduction of about **-4.4 percent**. This estimate is also highly significant. The close alignment between the level-based and log-based results strengthens the credibility of the main finding, as it suggests that the result is not driven by a small number of unusually high or low-traffic days. In other words, the inference does not depend materially on whether the outcome is modeled in absolute counts or in approximate percentage terms.

A third validation step is the placebo exercise, centered on October 2022. The purpose of this test is not to prove that no October-related movement ever occurs in Area C traffic, but to assess whether the 2023 finding simply reflects a recurring seasonal pattern that would have appeared around the same date even in the absence of the tariff reform. The placebo estimate is negative but weaker than the true 2023 treatment effect and less compelling than the combined evidence from the main event-study, the longer-run specification, and the log robustness check. This matters because it increases confidence that the 2023 result is not reducible to ordinary seasonality alone. At the same time, the placebo is not literally zero, so the exercise should be interpreted cautiously. It supports the substantive interpretation of a real policy effect without implying that every alternative timing-based concern has been fully eliminated.

Taken together, the event-study estimates, the longer-run filtered model, the log specification, and the placebo test point in the same direction. The evidence consistently indicates that the 30 October 2023 tariff increase was followed by a genuine reduction in Area C traffic, with a

large immediate response and a smaller but still meaningful persistent effect. This robustness pattern does not remove the observational nature of the study. Still, it substantially strengthens the case that the main finding reflects a policy-induced behavioral change rather than random fluctuation or a purely mechanical seasonal effect.

3.6 Economic Valuation Framework

3.6.1 Scope, appraisal basis, and valuation logic

To translate the estimated reduction in Area C traffic into an economic quantity, this thesis applies a **value-of-time (VoT)-based partial welfare valuation** of congestion relief. The purpose is not to estimate the overall financial performance of Area C as a municipal charging scheme, but to quantify the economic value of reduced congestion associated with the observed decline in entries after the October 2023 tariff increase. This distinction is fundamental to the interpretation of the results. The monetization developed here concerns one channel of policy benefit, namely the value of reduced delay implied by lower traffic pressure at the entry perimeter of a congested urban core. It does not measure revenues, operating margins, or the full social value of the scheme.

The valuation framework combines two complementary European sources. First, the *Handbook on the external costs of transport* is used as the transport-specific source because it provides methodologies, input values, and output values for external-cost estimation, including congestion and time valuation. Second, the *Economic Appraisal Vademecum 2021–2027* provides the broader logic of transparent, structured economic appraisal in public policy analysis. Overall, these sources support a disciplined but bounded valuation strategy: one that links the best-identified empirical effect in the thesis to a monetary interpretation without overstating the evidentiary base.

This choice of framework also clarifies what the monetization in this chapter is, and what it is not. It is a **partial welfare valuation** of congestion relief, not a full social cost-benefit analysis. A fuller CBA would require treatment-area-specific inputs on changes in travel time, environmental effects, noise, accidents, vehicle operating costs, and attributable operating and enforcement costs. Since a continuous public microdata panel of Area C-wide travel times over the full 2021–2024 period was not identified, a narrower but more robust welfare-based valuation is preferred. The Handbook itself notes that a more exact estimation requires case-specific input parameters and evaluation models. (Directorate-General for Mobility and Transport (European Commission) et al., 2020).

3.6.2 Valuation Assumptions and Scenario Design

The monetary value of congestion relief is estimated using the following relationship:

$$Benefit = \Delta Q \times \Delta t \times VoT$$

where ΔQ is the estimated reduction in daily transits, Δt is the assumed average time saving per avoided transit, and VoT is the value of time.

The valuation uses two treatment effects. The first is the **short-run shock effect**, taken from the weekly-binned event-study estimate around the **30 October 2023** tariff increase. The second is the **long-run effect**, taken from the longer-filtered specification that excludes August and the late-December/early-January holiday period. This distinction is analytically useful because it separates the immediate behavioral response from the more persistent adjustment that remained after the initial price shock had passed.

For the value of time, the analysis uses the Handbook’s Italy-specific short-distance urban car values, summarized in Table 3.3 below. The urban values are chosen because Area C is a central-city congestion-charge zone, and the mechanism being valued is congestion relief within Milan’s urban road network rather than inter-urban travel.

Table 3.3. Value-of-time inputs used in the valuation

Trip purpose	Italy short-distance urban VoT (€ / hour)	Source
Personal	5.9	European Commission Handbook, Table 87
Commuting/business	12.8	European Commission Handbook, Table 87
Central scenario (midpoint)	9.35	The author's calculation

The Handbook distinguishes between short-distance urban trips (less than 32 km) and long-distance inter-urban trips (more than 32 km), and for the purposes of Area C the urban short-distance values are the more appropriate reference. The relevant Italy-specific values are €5.9/hour for personal trips and €12.8/hour for commuting or business trips. The central value-of-time scenario (€9.35/hour) is set at the midpoint between those two values. This midpoint is used as a neutral benchmark, not because the study observes an exact trip-purpose composition among Area C entrants, but because a more precise purpose-weighted average is unavailable.

The second valuation input is the time saving per avoided transit. Because the exact travel-time gain per avoided transit is not directly observed for the treated road network, the monetization analysis uses three transparent scenario assumptions: **0.5 minutes**, **1.0 minute**, and **2.0 minutes** saved per avoided transit. These are interpreted as **conservative**, **baseline**, and **upper-bound** assumptions about the marginal congestion-relief effect of lower traffic volumes in a dense urban core. Values above 2 minutes were not used because they would imply a stronger network effect than can be credibly justified with the available evidence. The scenarios are therefore intended to define a bounded, transparent welfare range rather than a spurious point estimate.

3.6.3 Monetized congestion-relief results

*Table 3.4. Monetized congestion-relief benefits under the central VoT scenario
(Source: Author's calculations based on Area C gate data and scenario assumptions)*

Effects Type	Short-run shock effect			Long-run effect		
Avoided Transit/day	13314			4613		
Panel of VoT Assumption (Central): 9.35 €/hour						
Time-saving scenarios	Low	Central	High	Low	Central	High
Minutes saved per avoided transit	0.5	1	2	0.5	1	2
Daily benefit (€)	1037	2075	4150	359	719	1438
Weekly benefit (€)	5187	10374	20748	1797	3594	7189
Annual benefit (€)	238602	477204	954408	82670	165340	330680

Table 3.4 reports the main valuation results under the **central value-of-time scenario (€9.35/hour)** and separates the **short-run shock effect** from the **long-run effect**. Under the **baseline assumption of 1 minute saved per avoided transit**, the short-run tariff shock implies monetized congestion-relief benefits of approximately **€2,075 per day**, **€10,374 per week**, and **€477,204 per year**. Under the same assumptions, the longer-run effect implies benefits of approximately **€719 per day**, **€3,594 per week**, and **€165,340 per year**.

The alternative time-saving assumptions define a meaningful range around those central estimates. Under the **upper-bound scenario of 2 minutes saved per avoided transit**, the corresponding annualized benefits rise to **€954,408** in the short run and **€330,680** in the longer run. Under the **conservative 0.5-minute scenario**, the annualized benefits remain positive but smaller, at approximately **€238,602** in the short run and **€82,670** in the longer run. The appendix reports the fuller scenario grid, including low, central, and high value-of-time cases, but the main text focuses on the midpoint scenario in order to present a restrained and interpretable central estimate. These values suggest that the October 2023 tariff increase generated economically meaningful congestion-relief benefits even under cautious assumptions. Concurrently, the variation across scenarios makes clear that the estimated

monetary value remains sensitive to the assumed time saving per avoided transit, explaining why the chapter treats the exercise as a bounded welfare valuation rather than as a single fully observed social total. The purpose of monetization here is not to create a false impression of precision, but to translate the traffic effect into an economic order of magnitude.

3.6.4 Interpretation, limitations, and extension toward fuller appraisal

The monetized values in Table 3.4 above should not be confused with the financial performance of Area C as a charging scheme. Revenues collected from access charges and the system's operating costs are accounted for in a fiscal or budgetary perspective, whereas the estimates reported here are evaluated from a welfare-oriented perspective that values reduced congestion and delay. The two should therefore not be mechanically combined into a single total. Separating them is analytically important because they answer different questions: one concerns the budgetary operation of the scheme, while the other concerns the social value of reduced traffic pressure. Historical municipal material suggests that Area C has generated substantial revenues relative to its operating costs, but those figures reflect the scheme's overall financial performance rather than the marginal welfare effect of the 30 October 2023 tariff increase estimated in this thesis (Sevino, 2019). The main limitation of the current valuation is the lack of a continuous treatment-area-specific travel-time panel for the full study period. Thus, the monetization necessarily relies on scenario assumptions for the time saved per avoided transit. This limitation is handled transparently by keeping the assumptions modest and by reporting a full scenario grid in the appendix. A fuller future appraisal could extend the analysis by incorporating environmental co-benefits, vehicle operating costs, and, if reliable data become available, the scheme's fiscal balance. The next chapter interprets these results in relation to urban traffic management, policy design, and the broader role of IoT-driven smart mobility systems.

Chapter 4: Discussion and Policy Recommendations

This chapter interprets the empirical findings and draws out their policy implications. Chapter 3 established that the 30 October 2023 Area C tariff increase was associated with a substantial short-run decline in daily transits, a smaller but persistent longer-run reduction, and a differentiated response across payer categories, translating into non-trivial monetized congestion-relief benefits under a transparent value-of-time framework. Three connected questions follow. First, what do the results reveal about the role of IoT-driven smart traffic systems in managing congestion? Second, what do the findings suggest about the behavioral effects of pricing within a mature urban congestion-charge regime? Third, how can these results inform future policy design in Milan and comparable European cities? The central argument is that Area C should be understood not only as a charging mechanism but as a data-generating and behavior-shaping smart mobility system whose value lies in combining pricing, continuous monitoring, and adaptive governance into a single institutional design.

4.1 Behavioral Findings and Economic Interpretation

The central empirical finding is that the October 2023 tariff elicited a clear, economically meaningful behavioral response. The marked short-run fall and its partial persistence beyond the initial shock window indicate that the reform did more than generate a brief implementation shock: it changed access behavior within an already mature congestion-charging regime. This is analytically important because Area C had been operating for more than a decade before the 2023 reform. The results, therefore, do not capture the effect of introducing congestion pricing into a previously unregulated environment; they identify residual elasticity within an established system in which many users had already incorporated the charge into their routine mobility decisions. In that sense, the findings strengthen the interpretation of Area C as an active demand-management instrument rather than a merely symbolic or revenue-generating measure.

Before situating this result in the wider literature, it is important to acknowledge that the longer-run persistence of the effect, while statistically robust, cannot be fully decomposed with the available data. At least three mechanisms may contribute: genuine habit adjustment by users who restructured trips in response to the higher tariff; fleet-composition effects linked to the wider Area B regulatory environment, which may have reduced the pool of vehicles able or willing to enter the core; and broader post-pandemic mobility shifts that, although partially addressed by the placebo exercise, cannot be ruled out completely. The current design identifies the combined temporal change rather than separately isolating these channels. That limitation does not invalidate the result, but it does bound its interpretation and reinforces the value of the corridor-based difference-in-differences extension proposed in Chapter 5.

Set against the existing evidence on Area C, the October 2023 result is especially noteworthy. Beria, Tosi, and Nuccio (2017) show that the original introduction of Area C in January 2012 generated a sustained reduction of approximately 22 to 23 percent in average daily entries relative to the 2011 baseline, equivalent to roughly 33,000 fewer vehicles per day, and that this effect remained broadly stable over the following four years. The present study measures something different: not the initial shock of introducing a congestion charge, but the marginal behavioral response to a later tariff increase within a system that had already reached a mature equilibrium. In this context, the estimated 12.8 percent short-run reduction is not simply a smaller version of the 2012 effect. It shows that even after prolonged exposure to congestion pricing, a meaningful share of demand remains price-sensitive, and that tariff adjustment within an established regime retains genuine behavioral force. That finding is directly relevant to cities operating mature charging systems and considering whether incremental reforms still matter.

The payer decomposition sharpens this interpretation. The strongest adjustment appears among other payers, with service payers also showing a significant decline, while resident payers do not exhibit a statistically significant short-run response. This pattern is substantively plausible and theoretically consistent with the planning perspectives developed in Chapter 1. It suggests that congestion pricing operates most strongly on flexible and discretionary access, while having weaker immediate effects on users whose mobility is more routine, spatially constrained, or harder to substitute away from. In this sense, the tariff increase did not act as a uniform deterrent across all entrants; it worked selectively, and that selectivity matters for both interpretation and policy design.

As discussed in Section 1.4.4, this heterogeneity also has a distributional meaning. Recent work on transport-related social exclusion emphasizes that unequal mobility outcomes arise not simply from price levels, but from the interaction between cost burdens, limited modal alternatives, spatial mismatch, and unequal ability to reach socially valued opportunities (Bantis & Haworth, 2020; Luke, 2024). The weaker resident response is consistent with that interpretation: some users appear to absorb the higher tariff rather than substitute away from the trip because their mobility is more constrained or less easily reorganized. This implies that future tariff reforms should be evaluated not only in terms of aggregate elasticity but also in terms of which user groups face the least viable alternatives and therefore bear the greatest adjustment burden. Exemption design and category rules cannot, therefore, be treated as purely administrative details; they form part of the equity architecture of the policy itself. This segmentation of the demand response also has a specific implication for the welfare calculation developed in Chapter 3. The congestion relief generated by the reform benefits all remaining users, including those who continue entering the zone at the higher tariff. The cost of adjustment, however, is borne disproportionately by those who responded by canceling, rescheduling, or redirecting trips. In this sense, the aggregate monetized welfare estimate

partially conceals a redistribution from more flexible to less flexible users. A fuller distributional appraisal, decomposing welfare gains and costs by payer category, would be needed to make that asymmetry explicit.

The robustness analysis reinforces the main interpretation without eliminating all uncertainty. The persistence of the effect in the filtered longer-run model and in the log specification reduces the likelihood that the result is driven by outcome scale or a small number of unusual days. The placebo exercise around October 2022 yields a weaker break than the true 2023 treatment effect, supporting the interpretation of a genuine policy response rather than a purely recurring seasonal artifact. At the same time, the placebo is not literally zero, so the evidence should still be read as bounded causal inference rather than definitive experimental proof. That balance between robustness and restraint is important to the credibility of the chapter as a whole.

The monetization in Chapter 3 translates this traffic reduction into a bounded partial welfare valuation of congestion relief. Under the central value-of-time scenario and the baseline assumption of one minute saved per avoided transit, the short-run tariff shock implies congestion-relief benefits of approximately €2,075 per day and €477,204 per year, while the longer-run effect implies approximately €719 per day and €165,340 per year. These estimates are deliberately partial. They do not represent municipal profit, nor the full social value of Area C. Rather, they identify one channel of economic benefit, reduced delay, and separate it from fiscal performance, environmental co-benefits, and distributional effects. The sensitivity of the estimates to the assumed time saving confirms that the precise monetary magnitude is uncertain, but the sign and substantive importance of the welfare effect are robust across all scenarios. For context, Beria, Tosi, and Nuccio (2017) report that Area C generated annual revenues of €20-€29 million in its early years, with a substantial share reinvested in public transport and sustainable mobility. This suggests that the full economic significance of the

scheme operates on a scale considerably larger than the congestion-relief channel estimated here, and that a more comprehensive appraisal would need to incorporate those wider downstream effects as well.

Taken together, the findings show that the October 2023 tariff increase reduced traffic entering the central zone, that this response was segmented by user type, and that the resulting welfare gains are real but incomplete in both economic and planning terms. The main interpretive conclusion is therefore not simply that pricing “worked,” but that it worked in a selective and institutionally meaningful way: it retained behavioral force within a mature charging regime, generated measurable congestion-relief benefits, and at the same time exposed the importance of accessibility conditions and user heterogeneity in shaping who adjusts, who benefits, and who bears the burden of reform.

4.2 Area C as Smart Regulatory Infrastructure

This thesis treats Area C not merely as a congestion charge but as an IoT-driven smart traffic system whose regulatory capacity depends on the integration of sensing, classification, enforcement, and data generation. The ANPR infrastructure does not simply detect entry into the charging zone: it records plate numbers at the gate, cross-checks them against exemption lists, payment records, and vehicle-registry information, and routes unmatched or non-compliant entries into the enforcement process. In this regard, the system is not only a charging boundary but a digitally mediated regulatory architecture in which measurement and rule application are operationally embedded in the same infrastructure (Beria et al., 2017; Gibson & Carnovale, 2015; Percoco, 2013). The October 2023 reform illustrates the analytical importance of that architecture. The price signal altered incentives, but the gate-generated administrative record is what made it possible to estimate both the immediate and the more persistent behavioral response with a level of temporal precision that periodic

manual counts or survey-based evidence could not provide.

This is the sense in which Area C functions as smart regulatory infrastructure rather than simply as a priced access rule. A smart mobility system, in the meaning relevant to this thesis, is one in which intervention and observation are joined: the same infrastructure that changes behavior also records the consequences of that change. This matters because it gives the policy adaptive capacity. A congestion charge judged only by revenues or headline traffic counts remains a static instrument; a congestion charge that continuously produces structured behavioral evidence can be recalibrated, justified, and improved through its own operation. The significance of Area C, therefore, lies not only in its ability to deter some trips but in its ability to make the effects of deterrence empirically visible and economically interpretable. That feature is what connects the IoT framing developed in Chapter 1 to the quasi-experimental strategy implemented in Chapters 2 and 3.

A further distinction is important here: sensing capacity is not the same thing as evaluation capacity. Other cities operate technically comparable ANPR-based charging systems, but where the underlying monitoring record is not publicly accessible, the feedback loop between intervention and independent evaluation is weakened. The system may still enforce a charge, but its behavioral effects cannot be externally scrutinized in the same way, and tariff adjustment depends more heavily on internal or proprietary analysis. Area C is analytically valuable not only because it senses and enforces, but because the resulting record can be aggregated, examined, and used as evidence. In that sense, open administrative publication is not a secondary technical detail. It is part of what makes the system genuinely “smart” in governance terms, because it transforms monitoring infrastructure from a control device into an evaluative public asset.

4.3 Governance, Data Ethics, and Public Legitimacy

The analytical value of Area C's data-generation function carries a corresponding governance obligation. The same camera-based monitoring record that makes the empirical strategy of this thesis possible also raises questions of privacy, proportionality, and public trust.

In this context, the issue is not whether digital monitoring is useful, since Chapters 2 and 3 show that it is, but whether its evaluative benefits can be preserved without weakening its public legitimacy.

Three tensions are central. The first is **purpose limitation**: data collected for congestion charging and policy evaluation should remain subject to clear rules governing retention, reuse, and access. The second is **function creep**: infrastructure introduced for mobility regulation can gradually be extended to wider administrative or policing uses if those boundaries are not explicitly defined. The third is **cybersecurity**: both enforcement and policy evaluation depend on the reliability, continuity, and integrity of the monitoring system. For this reason, a smart traffic system should be judged not only by whether it changes behavior, but also by whether it is governed through privacy-conscious, proportionate, and operationally secure data practices (European Union Agency for Cybersecurity, 2020).

Milan's case also suggests that these issues are institutional, not merely technical. As Beria, Tosi, and Nuccio (2017) show, the transition from Ecopass to Area C was tied to public commitments on transparency, policy purpose, and revenue use. The open-data environment that enables this thesis can therefore be understood as part of the scheme's broader accountability logic rather than as an accidental administrative by-product. The broader implication is that cities seeking to build smart regulatory infrastructure should design governance credibility and data openness into the system from the outset, rather than treating them as optional additions after implementation.

4.4 Policy Implications for Milan

The findings suggest that Milan should treat Area C as an adjustable policy instrument rather than a completed institutional settlement. Five implications follow.

First, **pricing should remain a flexible and periodically recalibrated lever**. The October 2023 reform shows that tariff adjustments continue to influence access demand even within a mature regime: a €2.50 increase produced a statistically significant and partly persistent reduction in daily entries. This indicates that residual elasticity persists in the system and that periodic recalibration, guided by observed gate responses, is likely more effective than long periods of inaction followed by infrequent, large reforms.

Second, **tariff and exemption design should be more explicitly category-sensitive**. The strongest behavioral adjustment appears among the more flexible payer groups, while resident payers show a weaker measured response. This means that a single undifferentiated price signal does not affect all users in the same way. Future reforms should therefore pay closer attention to category-specific elasticity and to the extent to which different users face viable modal alternatives. In this respect, the social-exclusion framework discussed in Chapter 1 implies that exemption design is not a technical afterthought, but part of the distributive logic of the scheme itself.

Third, **pricing works best when embedded in a wider accessibility strategy**. Some of the observed reduction in entries may reflect trip suppression rather than desirable mode shift. For that reason, congestion pricing should remain linked to sustained investment in public transport access to the core, shared mobility, park-and-ride capacity, and broader proximity-enhancing urban policies. The planning perspectives developed in Section 1.4.4 suggest that the welfare gains from Area C would be larger and more evenly distributed if users in outer districts faced stronger non-car alternatives before reaching the charging perimeter.

Fourth, **the Area C data architecture should be treated as a strategic policy asset.** The gate system does more than enforce the charge: it generates continuous evidence on how users respond to policy change. That makes it possible to evaluate small reforms, detect behavioral adjustments, and recalibrate the instrument with far greater precision than would be possible under survey-based or intermittent monitoring. Milan should therefore treat Area C gate data not only as an enforcement record but also as a standing mobility observatory that can be integrated into future SUMP cycles alongside accessibility, modal share, and emissions indicators.

Fifth, **future appraisal should move beyond the congestion-relief channel estimated here.** The present thesis deliberately monetizes only one benefit dimension, reduced delay, and leaves aside fiscal performance, environmental co-benefits, and fuller distributional decomposition. A more complete future appraisal would combine the welfare logic developed in Chapter 3 with revenue, emissions, and accessibility evidence, allowing Milan to evaluate Area C not only as a pricing scheme but as a broader smart-mobility governance instrument.

4.5 Broader Implications for Smart Urban Mobility Policy

Although focused on Milan, the findings speak to broader European debates on smart urban traffic governance. Many cities face the same structural problem: dense urban cores generate substantial economic losses from congestion, yet blanket traffic bans and politically rigid restrictions are difficult to sustain. Digitally monitored congestion charging offers a more flexible alternative by combining access regulation, differentiated user treatment, and continuous behavioral observation. The Milan case suggests that the value of such systems lies not in digital technology alone, but in the combination of pricing, monitoring, and iterative policy adjustment. A smart mobility system, in the sense relevant to this thesis, is one that generates policy-relevant evidence through its own operation and thereby supports

recalibration over time. The Area C experience shows that the data-generation function of an IoT-based access-control system is not a secondary technical feature but a core governance asset. Cities investing in digitally monitored charging infrastructure should therefore design their systems with evaluation in mind from the outset: structured administrative records, consistent identifiers, and public data availability are the preconditions that make rigorous post-reform analysis possible. Where those conditions are absent, a system may still enforce a charge, but its effects become harder to scrutinize independently, and its adaptive capacity is correspondingly weaker.

A related lesson concerns political economy. The October 2023 reform did not involve introducing congestion charging into a previously free-access environment, but adjusting an existing charge within a system that had already achieved institutional and democratic legitimacy. This matters because the barriers to incremental tariff reform are lower than those associated with the introduction of a first-time scheme. In cases such as Edinburgh and Manchester, proposals for urban road charging encountered strong resistance before implementation (Gaunt et al., 2007; Beria et al., 2017). By contrast, the Milan case suggests that once pricing has become embedded in routine urban governance, users are more likely to respond behaviourally to tariff changes than to contest the legitimacy of the system itself. The broader lesson for European cities is therefore not simply that congestion charging can work, but that its long-run effectiveness depends on treating digital monitoring, public accountability, and policy adaptability as integral parts of the scheme rather than as optional additions.

4.6 Limitations

The methodological limitations and delimitations of this study are set out in full in Chapter 2, Section 2.7. In the context of interpreting the discussion above, two are particularly consequential. First, the absence of an external control group means the empirical strategy identifies a temporally concentrated change consistent with policy-induced behavioral response, but does not observe a perfect untreated counterfactual. The results should therefore be read as bounded causal inference rather than experimental proof. Second, the economic valuation is intentionally partial: it monetizes only the congestion-relief channel, and the resulting welfare range depends on scenario assumptions about time saved per avoided transit rather than on directly observed travel-time data. Both constraints are methodological delimitations rather than analytical failures, and they are handled through the multi-model robustness strategy described in Chapter 2.

Chapter 5: Conclusion

5.1 Main Findings and Contributions

This thesis examined whether an IoT-driven smart traffic system can be shown to reduce the economic loss associated with urban congestion, using Milan's Area C as the empirical case. The evidence indicates that the October 2023 tariff increase produced a statistically significant short-run reduction of approximately 12.8 percent in daily gate entries, followed by a smaller but still negative longer-run effect. The response was not uniform across user groups: it was concentrated among the more flexible payer categories, while resident payers showed no statistically significant short-run adjustment. Monetized through a value-of-time partial welfare framework, the estimated congestion-relief benefits range from approximately €165,000 in the longer run to €477,000 in the short-run shock period under central assumptions. These values should be read as bounded rather than exhaustive, since they capture one identified dimension of the reform's social value, reduced delay, while excluding environmental co-benefits, fiscal performance, and fuller distributional effects.

The thesis makes three linked contributions. Empirically, it shows that pricing retains meaningful behavioral force even within a mature congestion-charging regime, rather than only at the moment of first introduction. Methodologically, it demonstrates that open ANPR administrative records from a digitally monitored charging system can support a credible quasi-experimental evaluation of post-reform behavioral change. Conceptually, it argues that the significance of IoT-based traffic systems lies not only in operational optimization, but in their capacity to function as smart regulatory infrastructure: systems in which pricing, enforcement, and continuous data generation are integrated in a way that makes congestion management empirically observable, economically interpretable, and institutionally adaptable.

5.2 Directions for Future Research

Three directions for future research follow directly from the thesis's most consequential methodological constraints, as stated in full in Chapter 2, Section 2.7.

The first and most important concerns identification. The current empirical strategy relies on the internal temporal structure of the Area C series because no structurally comparable untreated corridor with equally rich open entry-count data was identified. A methodologically stronger design would apply a corridor-based difference-in-differences framework, comparing traffic flows on access corridors that feed directly into the Area C charging boundary against structurally similar corridors outside the perimeter, around the same reform date. This would provide a genuine treated-versus-control comparison and sharper causal identification than the interrupted time-series can achieve alone. It would also make it possible to distinguish rerouting effects from genuine demand suppression or mode shift, which the current gate-count structure cannot disentangle. The granular spatial data required for such a design, corridor-level flow counts with consistent pre- and post-reform coverage, are increasingly available from AMAT's monitoring systems and would represent the most direct methodological extension of this thesis.

The second concerns valuation precision. The welfare estimates rest on scenario assumptions about minutes saved per avoided transit because a continuous Area C-specific travel-time panel was not publicly available for the full study period. Combining the gate-based entry identification developed here with GPS-derived or loop-detector-based speed data for the Area C corridor and adjacent streets would allow the welfare valuation to rest on directly observed time savings rather than assumed ones, substantially improving both precision and credibility. This extension would also enable estimation of spillover effects on streets outside the charging perimeter, addressing the scope limitation that restricts the current analysis to the monitored boundary.

The third concerns the spatial and dimensional scope of the appraisal framework. This thesis focuses exclusively on Area C, the inner congestion-charge perimeter. A natural and analytically rich extension would apply a comparable design to Area B, Milan's Low Emission Zone, if equivalent open administrative records become available from its 185 ANPR gates. Area B operates on the same IoT infrastructure platform and addresses a different externality dimension, fleet composition, and emissions, rather than access volume and delay. A comparative study of behavioral responses to Area B's exclusion rules would allow examination of whether fleet-composition regulation and access pricing produce additive or substitutive effects on congestion and emissions across the metropolitan area, which is a genuinely open question with direct relevance to the design of layered demand-management architectures in other European cities. More broadly, extending the appraisal framework to incorporate emissions monitoring, fiscal balance data, and distributional decomposition alongside the congestion-relief channel would move the analysis toward the multi-dimensional evaluation that SUMP governance frameworks require and that the Area C and Area B IoT data infrastructures are, in principle, already capable of supporting.

References

- Afandizadeh Zargari, S., Memarnejad, A., & Mirzahosseini, H. (2021). Hourly Origin–Destination Matrix Estimation Using Intelligent Transportation Systems Data and Deep Learning. *Sensors*, 21(21), 7080. <https://doi.org/10.3390/s21217080>
- Afrin, T., & Yodo, N. (2020). A Survey of Road Traffic Congestion Measures towards a Sustainable and Resilient Transportation System. *Sustainability*, 12(11), 4660. <https://doi.org/10.3390/su12114660>
- Akbardin, J., Parikesit, D., Riyanto, B., & Mulyono, A. T. (2020). Development of Road Freight Transportation Distribution Model Based on Vehicle Transportation Inter Zone Requirement. *Transportation Research Procedia*, 48, 562–573. <https://doi.org/10.1016/j.trpro.2020.08.060>
- Alyavina, E., Nikitas, A., & Njoya, E. T. (2024). Mobility-as-a-Service and unsustainable travel behaviour: Exploring the car ownership and public transport trip replacement side-effects of the MaaS paradigm. *Transport Policy*, 150, 53–70. <https://doi.org/10.1016/j.tranpol.2024.03.001>
- AMAT. (2022). Report della Mobilità 2022.pdf (p. 73). <https://online.fliphtml5.com/AMAT-MI/oapr/index.html#p=1>
- AMAT. (2023). Report della Mobilità 2023.pdf (p. 74). https://online.fliphtml5.com/AMAT-MI/Report_Milano_2_DEF_-POvG/index.html#p=1
- AMAT. (2024a). Report della Mobilità 2024.pdf (p. 53). <https://datashare.amat-mi.it/index.php/s/AkbAstJcrTJWXNk>
- AMAT. (2024b, February 27). AMAT ha ottenuto la certificazione di un sistema avanzato di Intelligenza Artificiale per il monitoraggio della mobilità urbana milanese—Agenzia Mobilità Ambiente Territorio. AMAT. <https://www.amat-mi.it/it/notizie/122/>
- An, C., & Shen, J. (2025). Assessing the impact of Mobility-as-a-Service (MaaS) on sustainable urban travel behaviors: A systematic literature review. *Frontiers in Sustainable Cities*, 7. <https://doi.org/10.3389/frsc.2025.1645488>
- Anas, A. (2020). The cost of congestion and the benefits of congestion pricing: A general equilibrium

- analysis. *Transportation Research Part B: Methodological*, 136, 110–137.
<https://doi.org/10.1016/j.trb.2020.03.003>
- Andamento della popolazione 2014 – 2024. (n.d.). Portale Del Dato. Retrieved February 21, 2026, from
<https://dati.comune.milano.it/web/portale-del-dato/w/andamento-della-popolazione-2014-2024>
- Automobile Club d'Italia. (2025). Rapporto incidenti stradali 2024 (p. 16). Istituto Nazionale di Statistica.
- Bantis, T., & Haworth, J. (2020). Assessing transport related social exclusion using a capabilities approach to accessibility framework: A dynamic Bayesian network approach. *Journal of Transport Geography*, 84, 102673. <https://doi.org/10.1016/j.jtrangeo.2020.102673>
- Batley, R., Bates, J., Bliemer, M., Börjesson, M., Bourdon, J., Cabral, M. O., Chintakayala, P. K., Choudhury, C., Daly, A., Dekker, T., Drivyla, E., Fowkes, T., Hess, S., Heywood, C., Johnson, D., Laird, J., Mackie, P., Parkin, J., Sanders, S., ... Worsley, T. (2019). New appraisal values of travel time saving and reliability in Great Britain. *Transportation*, 46(3), 583–621.
<https://doi.org/10.1007/s11116-017-9798-7>
- Bento, A., Roth, K., & Waxman, A. (2024). The Value of Urgency: Evidence from Real-Time Congestion Pricing. *Journal of Political Economy Microeconomics*, 2(4), 786–851.
<https://doi.org/10.1086/731832>
- Beria, P., Tosi, L., & Nuccio, D. (2018). Four years of Milan's road charge: Effectiveness, acceptability and impacts (pp. 501–521). https://doi.org/10.1049/PBTR008E_ch19
- Bernal, J. L., Cummins, S., & Gasparrini, A. (2017). Interrupted time series regression for the evaluation of public health interventions: A tutorial. *International Journal of Epidemiology*, 46(1), 348–355. <https://doi.org/10.1093/ije/dyw098>
- Berrill, P., Nachtigall, F., Javaid, A., Milojevic-Dupont, N., Wagner, F., & Creutzig, F. (2024). Comparing urban form influences on travel distance, car ownership, and mode choice. *Transportation Research Part D: Transport and Environment*, 128, 104087.
<https://doi.org/10.1016/j.trd.2024.104087>
- Binsuwadan, J., De Jong, G., Batley, R., & Wheat, P. (2022). The value of travel time savings in

- freight transport: A meta-analysis. *Transportation*, 49(4), 1183–1209.
<https://doi.org/10.1007/s11116-021-10207-2>
- Büchel, B., & Corman, F. (2020). Review on Statistical Modeling of Travel Time Variability for Road-Based Public Transport. *Frontiers in Built Environment*, 6, 70.
<https://doi.org/10.3389/fbuil.2020.00070>
- Bueno-Suárez, C., & Coq-Huelva, D. (2020). Sustaining What Is Unsustainable: A Review of Urban Sprawl and Urban Socio-Environmental Policies in North America and Western Europe. *Sustainability*, 12(11), 4445. <https://doi.org/10.3390/su12114445>
- Buxo, T. L. (2021, March 9). Vianova, an EIT Urban Mobility supported startup improves the shared mobility services in Milan. EIT Urban Mobility.
<https://www.eiturbanmobility.eu/vianova-an-eit-urban-mobility-supported-startup-improves-the-shared-mobility-services-in-milan/>
- Chatziioannou, I., Nikitas, A., Tzouras, P. G., Bakogiannis, E., Alvarez-Icaza, L., Chias-Becerril, L., Karolemeas, C., Tsigdinos, S., Wallgren, P., & Rexfelt, O. (2023). Ranking sustainable urban mobility indicators and their matching transport policies to support liveable city Futures: A MICMAC approach. *Transportation Research Interdisciplinary Perspectives*, 18, 100788.
<https://doi.org/10.1016/j.trip.2023.100788>
- Chen, D., Yan, X., Liu, F., Liu, X., Wang, L., & Zhang, J. (2019). Evaluating and Diagnosing Road Intersection Operation Performance Using Floating Car Data. *Sensors*, 19(10), 2256.
<https://doi.org/10.3390/s19102256>
- Cheng, Z. (Aaron), Pang, M.-S., & Pavlou, P. A. (2020). Mitigating Traffic Congestion: The Role of Intelligent Transportation Systems. *Information Systems Research*, 31(3), 653–674.
<https://doi.org/10.1287/isre.2019.0894>
- Città Metropolitana di Milano. (2022). Progetto Sicurezza Milano Metropolitana: Dal 2019 -16% di incidenti, -18% di morti e -20% di feriti (p. 3) [Comunicato stampa]. Città Metropolitana di Milano.
<https://www.cittametropolitana.mi.it/portale/news/comunicati/Progetto-Sicurezza-Milano-Metropolitana-dal-2019-16-di-incidenti-18-di-morti-e-20-di-feriti/>

- Comune di Milano. (2025). Sette statistiche per la storia della città. Portale del Dato.
<https://dati.comune.milano.it/en/web/portale-del-dato/w/sette-statistiche-per-la-storia-della-citta-a?>
- Cordoş, N., Duma, I., Moldovanu, D., Todoruţ, A., & Barabás, I. (2025). An Overview of Intelligent Transportation Systems in Europe. *World Electric Vehicle Journal*, 16(7), 387.
<https://doi.org/10.3390/wevj16070387>
- Crova, F., Forello, A. C., Bernardoni, V., Calzolari, G., Canepari, S., Argentini, S., Costabile, F., Frezzini, M. A., Giardi, F., Lucarelli, F., Massabò, D., Massimi, L., Nava, S., Paglione, M., Pazzi, G., Prati, P., Rinaldi, M., Russo, M., Valentini, S., ... Vecchi, R. (2024). Assessing the role of atmospheric dispersion vs. Emission strength in the southern Po Valley (Italy) using dispersion-normalised multi-time receptor modelling. *Atmospheric Environment*, 316, 120168. <https://doi.org/10.1016/j.atmosenv.2023.120168>
- Diomede, D. G. (2024, February 26). RINA emette il primo certificato Trustable AI Management System—RINA Italy. RINA. RINA.org.
<https://www.rina.org/it/media/news/2024/02/27/trustable-ai-management-system>
- Directorate-General for Mobility and Transport (European Commission), CE Delft, Essen, Huib van, Wijngaarden, Lisanne van, Sutter, Daniel, Bieler, Cuno, Maffii, Silvia, Fiorello, Davide, Fermi, Francesca, Parolin, Riccardo, Schrotten, Arno, Brambilla, Marco, & El Beyrouy, Kareen. (2020). Handbook on the external costs of transport: Version 2019 – 1.1. Publications Office of the European Union. <https://data.europa.eu/doi/10.2832/51388>
- Direzione Mobilità e Trasporti - Comune di Milano. (2021, September). Disciplina per l'accesso di veicoli nella Zona a Traffico Limitato denominata "Area B" [MAYOR'S ORDINANCE]. Area Pianificazione e Programmazione Mobilità - Comune di Milano.
- Duranton, G., & Turner, M. A. (2011). The Fundamental Law of Road Congestion: Evidence from US Cities. *American Economic Review*, 101(6), 2616–2652.
<https://doi.org/10.1257/aer.101.6.2616>
- European Commission. (2019). GUIDELINES FOR DEVELOPING AND IMPLEMENTING A SUSTAINABLE URBAN MOBILITY PLAN (SECOND EDITION). Rupprecht Consult -

- Forschung & Beratung GmbH.
https://urban-mobility-observatory.transport.ec.europa.eu/system/files/2023-09/sump_guidelines_2019_second%20edition.pdf
- European Commission. (2021). COMMISSION STAFF WORKING DOCUMENT EVALUATION of the 2013 Urban Mobility Package. European Commission.
<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52021SC0047&qid=1775212310219>
- European Commission - Directorate-General for Regional and Urban Policy, & JASPERS. (2021). Economic Appraisal Vademecum 2021–2027: General principles and sector applications. European Commission.
https://ec.europa.eu/regional_policy/sources/guides/vademecum_2127/vademecum_2127_en.pdf
- European Union. (2023). COMMISSION RECOMMENDATION (EU) 2023/550 of 8 March 2023 on National Support Programmes for Sustainable Urban Mobility Planning. Official Journal of the European Union.
<https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX%3A32023H0550>
- European Union Agency for Cybersecurity (EU body or agency). (2020). Cybersecurity stocktaking in the CAM: Stakeholder mapping and stocktaking of connected and automated mobility (CAM) cybersecurity. Publications Office of the European Union.
<https://data.europa.eu/doi/10.2824/24902>
- Ewing, R., Tian, G., & Lyons, T. (2018). Does compact development increase or reduce traffic congestion? *Cities*, 72, 94–101. <https://doi.org/10.1016/j.cities.2017.08.010>
- Ewusie, J. E., Soobiah, C., Blondal, E., Beyene, J., Thabane, L., & Hamid, J. S. (2020). Methods, Applications and Challenges in the Analysis of Interrupted Time Series Data: A Scoping Review. *Journal of Multidisciplinary Healthcare*, 13, 411–423.
<https://doi.org/10.2147/JMDH.S241085>
- Gaunt, M., Rye, T., Allen, S., Ltd, F., & Sgm. (2007). Public Acceptability of Road User Charging: The Case of Edinburgh and the 2005 Referendum. *Transport Reviews*, 27, 85–102.

<https://doi.org/10.1080/01441640600831299>

- Gibson, M., & Carnovale, M. (2015). The effects of road pricing on driver behavior and air pollution. *Journal of Urban Economics*, 89, 62–73. <https://doi.org/10.1016/j.jue.2015.06.005>
- Greco, C. (2021, February 9). La mobilità intelligente avanza nei comuni italiani, a partire da Milano: IoT e videocamere per il monitoraggio [B2B Editorial Portal]. Digital4. <https://www.digital4.biz/executive/digital-transformation/mobilita-intelligente-nei-comuni-italiani-citta-metropolitana-di-milano-safety21/>
- Gu, C. (2019). Urbanization: Processes and driving forces. *Science China Earth Sciences*, 62(9), 1351–1360. <https://doi.org/10.1007/s11430-018-9359-y>
- Gu, Y., Fu, X., Liu, Z., Xu, X., & Chen, A. (2020). Performance of transportation network under perturbations: Reliability, vulnerability, and resilience. *Transportation Research Part E: Logistics and Transportation Review*, 133, 101809. <https://doi.org/10.1016/j.tre.2019.11.003>
- Hart, R., Kyriakopoulou, E., & Lu, T. (2024). Urban Transport Policies and Net Zero Emissions in the European Union. *Annual Review of Resource Economics*, 16(Volume 16, 2024), 187–206. <https://doi.org/10.1146/annurev-resource-101623-101611>
- Herzog, I. (2024). Microgeographic speed, reliability, and traffic externalities. *Economics of Transportation*, 39, 100371. <https://doi.org/10.1016/j.ecotra.2024.100371>
- Holden, E., Banister, D., Gössling, S., Gilpin, G., & Linnerud, K. (2020). Grand Narratives for sustainable mobility: A conceptual review. *Energy Research & Social Science*, 65, 101454. <https://doi.org/10.1016/j.erss.2020.101454>
- Hrelja, R., Olsson, L., Pettersson-Löfstedt, F., & Rye, T. (2020). TRANSIT ORIENTED DEVELOPMENT (TOD). A Literature Review.
- Ibraeva, A., Correia, G. H. D. A., Silva, C., & Antunes, A. P. (2020). Transit-oriented development: A review of research achievements and challenges. *Transportation Research Part A: Policy and Practice*, 132, 110–130. <https://doi.org/10.1016/j.tra.2019.10.018>
- Ison, S., & Rye, T. (2005). Implementing Road User Charging: The Lessons Learnt from Hong Kong, Cambridge and Central London. *Transport Reviews*, 25(4), 451–465. <https://doi.org/10.1080/0144164042000335788>

- Istituto Nazionale di Statistica (ISTAT). (2020). Day-time population for study and work [Experimental statistic]. Istituto Nazionale di Statistica (ISTAT).
<https://www.istat.it/wp-content/uploads/2020/11/Daytime-population-2016.pdf>
- Janjevic, M., & Winkenbach, M. (2020). Characterizing urban last-mile distribution strategies in mature and emerging e-commerce markets. *Transportation Research Part A: Policy and Practice*, 133(C), 164–196.
- Jayasooriya, S., & Bandara, Y. (2017). Measuring the Economic costs of traffic congestion. 2017 Moratuwa Engineering Research Conference (MERCon), 141–146.
<https://doi.org/10.1109/MERCON.2017.7980471>
- Knorr, D., Khoo, C. S. H., & Augustin, M. A. (2018). Food for an Urban Planet: Challenges and Research Opportunities. *Frontiers in Nutrition*, 4, 73. <https://doi.org/10.3389/fnut.2017.00073>
- Kriswardhana, W., & Esztergár-Kiss, D. (2023). A systematic literature review of Mobility as a Service: Examining the socio-technical factors in MaaS adoption and bundling packages. *Travel Behaviour and Society*, 31, 232–243. <https://doi.org/10.1016/j.tbs.2022.12.007>
- Kukkapalli, V. M., & Pulugurtha, S. S. (2020). Modeling the effect of a freeway road construction project on link-level travel times. *Journal of Traffic and Transportation Engineering (English Edition)*, 8(2), 267–281. <https://doi.org/10.1016/j.jtte.2019.11.002>
- Li, J., Li, H., Zhang, Y., & Luo, X. (2021). Evaluation of Joint Development of Park and Ride and Transit-Oriented Development Near Metro Stations in Chengdu, China. *Journal of Urban Planning and Development*, 147(2), 05021008.
[https://doi.org/10.1061/\(ASCE\)UP.1943-5444.0000675](https://doi.org/10.1061/(ASCE)UP.1943-5444.0000675)
- Li, J., Peng, L., Hou, K., Tian, Y., Ma, Y., Xu, S., Qiu, T. Z., Li, J., Peng, L., Hou, K., Tian, Y., Ma, Y., Xu, S., & Qiu, T. Z. (2023). Adaptive signal control and coordination for urban traffic control in a connected vehicle environment: A review. *Digital Transportation and Safety*, 2(2), 89–111. <https://doi.org/10.48130/DTS-2023-0008>
- Liang, L., Xu, M., Grant-Muller, S., & Mussone, L. (2019). Household travel mode choice estimation with large-scale data—An empirical analysis based on mobility data in Milan. *International Journal of Sustainable Transportation*, 15(1), 70–85.

<https://doi.org/10.1080/15568318.2019.1686782>

- Lieberthal, E. B., Serok, N., Duan, J., Zeng, G., & Havlin, S. (2024). Addressing the urban congestion challenge based on traffic bottlenecks. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 382(2285), 20240095.
<https://doi.org/10.1098/rsta.2024.0095>
- Liu, L., Zhang, M., & Xu, T. (2020). A conceptual framework and implementation tool for land use planning for corridor transit oriented development. *Cities*, 107, 102939.
<https://doi.org/10.1016/j.cities.2020.102939>
- Lu, H., Shang, Z., Ruan, Y., & Jiang, L. (2023). Study on Urban Expansion and Population Density Changes Based on the Inverse S-Shaped Function. *Sustainability*, 15(13), 10464.
<https://doi.org/10.3390/su151310464>
- Lu, J., Li, B., Li, H., & Al-Barakani, A. (2021). Expansion of city scale, traffic modes, traffic congestion, and air pollution. *Cities*, 108, 102974. <https://doi.org/10.1016/j.cities.2020.102974>
- Luke, R. (2024). Transport-related social exclusion and mobility in developing countries: The South African case. *Frontiers in Sustainable Cities*, 6. <https://doi.org/10.3389/frsc.2024.1199055>
- Marchetti, E., & Antonelli, C. (2024). Social aspects of low emission zones: Milan case study.
- Miller, D. L. (2023). An Introductory Guide to Event Study Models. *Journal of Economic Perspectives*, 37(2), 203–230. <https://doi.org/10.1257/jep.37.2.203>
- Moreno, C., Allam, Z., Chabaud, D., Gall, C., & Pratlong, F. (2021). Introducing the “15-Minute City”: Sustainability, Resilience and Place Identity in Future Post-Pandemic Cities. *Smart Cities*, 4(1), 93–111. <https://doi.org/10.3390/smartcities4010006>
- Mosquera, R. (2024a). Stuck in traffic: Measuring congestion externalities with negative supply shocks. *Regional Science and Urban Economics*, 108, 104013.
<https://doi.org/10.1016/j.regsciurbeco.2024.104013>
- Mosquera, R. (2024b). Stuck in traffic: Measuring congestion externalities with negative supply shocks. *Regional Science and Urban Economics*, 108, 104013.
<https://doi.org/10.1016/j.regsciurbeco.2024.104013>
- Moulin, L., & Urbano, V. M. (2025). Evaluating the effectiveness of two congestion limitation policies

- in Milan: Charge increase and vehicle type. *Transport Policy*, 165, 17–27.
<https://doi.org/10.1016/j.tranpol.2025.02.006>
- Nair, G. S., Bhat, C. R., Batur, I., Pendyala, R. M., & Lam, W. H. K. (2020). A model of deadheading trips and pick-up locations for ride-hailing service vehicles. *Transportation Research Part A: Policy and Practice*, 135, 289–308. <https://doi.org/10.1016/j.tra.2020.03.015>
- Oladimeji, D., Gupta, K., Kose, N. A., Gundogan, K., Ge, L., & Liang, F. (2023). Smart Transportation: An Overview of Technologies and Applications. *Sensors*, 23(8), 3880.
<https://doi.org/10.3390/s23083880>
- Orozco, L. G. N., Alessandretti, L., Saberi, M., Szell, M., & Battiston, F. (2023). Multimodal urban mobility and multilayer transport networks. *Environment and Planning B: Urban Analytics and City Science*, 50(8), 2038–2070. <https://doi.org/10.1177/23998083221108190>
- Ossokina, I. V., van Ommeren, J., & van Mourik, H. (2023). Do highway widenings reduce congestion? *Journal of Economic Geography*, 23(4), 871–900.
<https://doi.org/10.1093/jeg/lbac034>
- Otsubo, H., Gisches, E. J., & Rapoport, A. (2025). The Downs-Thomson Paradox with Endogenously Determined Departure Times. *Experimental Economics*, 28(3), 600–624.
<https://doi.org/10.1017/eec.2025.10025>
- Pandey, A., & Lehe, L. J. (2024). Congestive mode-switching and economies of scale on a bus route. *Transportation Research Part B: Methodological*, 183, 102930.
<https://doi.org/10.1016/j.trb.2024.102930>
- Paolini, F. (2009). A Country on four wheels: The car and society in Italy (1900-1974). *Tst Transportes Servicios Y Telecomunicaciones*.
https://www.academia.edu/75335997/A_Country_on_four_wheels_The_car_and_society_in_Italy_1900_1974_
- Papadakis, D. M., Savvides, A., Michael, A., & Michopoulos, A. (2024). Advancing sustainable urban mobility: Insights from best practices and case studies. *Fuel Communications*, 20, 100125.
<https://doi.org/10.1016/j.jfueco.2024.100125>
- Penfold, R. B., & Zhang, F. (2013). Use of interrupted time series analysis in evaluating health care

- quality improvements. *Academic Pediatrics*, 13(6 Suppl), S38-44.
<https://doi.org/10.1016/j.acap.2013.08.002>
- Percoco, M. (2013). Is road pricing effective in abating pollution? Evidence from Milan. *Transportation Research Part D: Transport and Environment*, 25, 112–118.
<https://doi.org/10.1016/j.trd.2013.09.004>
- Percoco, M. (2017). Cost Distribution and the Acceptability of Road Pricing: Evidence from Milan's Referendum. *Journal of Transport Economics and Policy*, 51(1), 34–46.
- Pi, M., Yeon, H., Son, H., & Jang, Y. (2021). Visual Cause Analytics for Traffic Congestion. *IEEE Transactions on Visualization and Computer Graphics*, 27(3), 2186–2201.
<https://doi.org/10.1109/TVCG.2019.2940580>
- Pouzoukidou, G., & Chatziyiannaki, Z. (2021). 15-Minute City: Decomposing the New Urban Planning Eutopia. *Sustainability*, 13(2), 928. <https://doi.org/10.3390/su13020928>
- Pucci, P. (2021). Spatial dimensions of electric mobility—Scenarios for efficient and fair diffusion of electric vehicles in the Milan Urban Region. *Cities*, 110, 103069.
<https://doi.org/10.1016/j.cities.2020.103069>
- Qin, X., Ke, J., Wang, X., Tang, Y., & Yang, H. (2022). Demand management for smart transportation: A review. *Multimodal Transportation*, 1(4), 100038.
<https://doi.org/10.1016/j.multra.2022.100038>
- Ren, Q., Man, K., Li, M., Gao, B., & Ma, J. (2019). Intelligent design and implementation of blockchain and Internet of things–based traffic system. *International Journal of Distributed Sensor Networks*, 15, 155014771987065. <https://doi.org/10.1177/1550147719870653>
- Romero, F., Gomez, J., Paez, A., & Vassallo, J. M. (2020). Toll roads vs. Public transportation: A study on the acceptance of congestion-calming measures in Madrid. *Transportation Research Part A: Policy and Practice*, 142, 319–342. <https://doi.org/10.1016/j.tra.2020.11.001>
- Rotaris, L., Danielis, R., Marcucci, E., & Massiani, J. (2010). The urban road pricing scheme to curb pollution in Milan, Italy: Description, impacts and preliminary cost–benefit analysis assessment. *Transportation Research Part A: Policy and Practice*, 44(5), 359–375.
<https://doi.org/10.1016/j.tra.2010.03.008>

- Saberi, M., Hamedmoghadam, H., Ashfaq, M., Hosseini, S. A., Gu, Z., Shafiei, S., Nair, D. J., Dixit, V., Gardner, L., Waller, S. T., & González, M. C. (2020). A simple contagion process describes spreading of traffic jams in urban networks. *Nature Communications*, 11(1), 1616.
<https://doi.org/10.1038/s41467-020-15353-2>
- Salvia, G., & Morello, E. (2020). Sharing cities and citizens sharing: Perceptions and practices in Milan. *Cities*, 98, 102592. <https://doi.org/10.1016/j.cities.2019.102592>
- Sánchez-Cambronero, S., Álvarez-Bazo, F., Rivas, A., & Gallego, I. (2021). Dynamic Route Flow Estimation in Road Networks Using Data from Automatic Number of Plate Recognition Sensors. *Sustainability*, 13(8), 4430. <https://doi.org/10.3390/su13084430>
- Serok, N., Havlin, S., & Blumenfeld Lieberthal, E. (2022). Identification, cost evaluation, and prioritization of urban traffic congestions and their origin. *Scientific Reports*, 12(1), 13026.
<https://doi.org/10.1038/s41598-022-17404-8>
- Sevino, V. (2019, June 12). Traffic Regulation Measures in Milan: Area C and Area B [Statistical]. GO SUMP training on financing sustainable mobility.
https://medcities.org/documents/Valentino+Sevino+-+Milan+_+High+Level+Trainings+_+Go+SUMP.pdf
- Shoup, D. (2018). Parking and the City. <https://doi.org/10.4324/9781351019668>
- Storme, T., De Vos, J., De Paepe, L., & Witlox, F. (2020). Limitations to the car-substitution effect of MaaS. Findings from a Belgian pilot study. *Transportation Research Part A: Policy and Practice, Developments in Mobility as a Service (MaaS) and Intelligent Mobility*, 131, 196–205. <https://doi.org/10.1016/j.tra.2019.09.032>
- Tang, J., & Heinemann, H. R. (2018). A resilience-oriented approach for quantitatively assessing recurrent spatial-temporal congestion on urban roads. *PLOS ONE*, 13(1), e0190616.
<https://doi.org/10.1371/journal.pone.0190616>
- Tang, J., Wan, L., Schooling, J., Zhao, P., Chen, J., & Wei, S. (2022). Automatic number plate recognition (ANPR) in smart cities: A systematic review on technological advancements and application cases. *Cities*, 129, 103833. <https://doi.org/10.1016/j.cities.2022.103833>
- Thomas, R., Pojani, D., Lenferink, S., Bertolini, L., Stead, D., & van der Krabben, E. (2018). Is

- transit-oriented development (TOD) an internationally transferable policy concept? *Regional Studies*, 52(9), 1201–1213. <https://doi.org/10.1080/00343404.2018.1428740>
- Tikoudis, I. (2023). Revisiting the Pigouvian tax in urban roads: Housing supply restrictions, leaking profits and spatial inequality. *Economics of Transportation*, 35, 100324. <https://doi.org/10.1016/j.ecotra.2023.100324>
- United Nations. (2025). *World Urbanization Prospects 2025: Summary of Results* (NO. 12 TR; p. 124). UN DESA/POP. <https://desapublications.un.org/publications/world-urbanization-prospects-2025-summary-resu>lts
- Wang, M., & Debbage, N. (2021). Urban morphology and traffic congestion: Longitudinal evidence from US cities. *Computers, Environment and Urban Systems*, 89, 1–11. <https://doi.org/10.1016/j.compenvurbsys.2021.101676>
- Xiao, D., Kim, I., & Zheng, N. (2024). Does built environment have impact on traffic congestion? —A bootstrap mediation analysis on a case study of Melbourne. *Transportation Research Part A: Policy and Practice*, 190, 104297. <https://doi.org/10.1016/j.tra.2024.104297>
- Yigitcanlar, T., Kamruzzaman, Md., Buys, L., Ioppolo, G., Sabatini-Marques, J., da Costa, E. M., & Yun, J. J. (2018). Understanding ‘smart cities’: Intertwining development drivers with desired outcomes in a multidimensional framework. *Cities*, 81, 145–160. <https://doi.org/10.1016/j.cities.2018.04.003>
- Yildirimoglu, M., Ramezani, M., & Amirgholy, M. (2021). Staggered work schedules for congestion mitigation: A morning commute problem. *Transportation Research Part C: Emerging Technologies*, 132, 103391. <https://doi.org/10.1016/j.trc.2021.103391>
- Yilmazkuday, H. (2025). Welfare costs of urban traffic through retail prices. *Transport Policy*, 162, 396–405. <https://doi.org/10.1016/j.tranpol.2024.12.017>
- Yin, C., Shao, C., & Wang, X. (2018). Built Environment and Parking Availability: Impacts on Car Ownership and Use. *Sustainability*, 10(7), 2285. <https://doi.org/10.3390/su10072285>
- Zang, Z., Xu, X., Qu, K., Chen, R., & Chen, A. (2022). Travel time reliability in transportation networks: A review of methodological developments. *Transportation Research Part C: Emerging Technologies*, 132, 103391. <https://doi.org/10.1016/j.trc.2021.103391>

Emerging Technologies, 143, 103866. <https://doi.org/10.1016/j.trc.2022.103866>

Zefreh, M. M., & Török, Á. (2020). DISTRIBUTION OF TRAFFIC SPEED IN DIFFERENT TRAFFIC CONDITIONS: AN EMPIRICAL STUDY IN BUDAPEST. *Transport*, 35(1), 68–86. <https://doi.org/10.3846/transport.2019.11725>

Zhang, C., Sabar, N. R., Chung, E., Bhaskar, A., & Guo, X. (2019). Optimisation of lane-changing advisory at the motorway lane drop bottleneck. *Transportation Research Part C: Emerging Technologies*, 106, 303–316. <https://doi.org/10.1016/j.trc.2019.07.016>

Zhu, F., Lv, Y., Chen, Y., Wang, X., Xiong, G., & Wang, F.-Y. (2020). Parallel Transportation Systems: Toward IoT-Enabled Smart Urban Traffic Control and Management. *IEEE Transactions on Intelligent Transportation Systems*, 21(10), 4063–4071.

<https://doi.org/10.1109/TITS.2019.2934991>

Appendices

This appendix reports supplementary valuation results and additional diagnostic figures supporting the empirical analysis of the October 2023 Area C interventions.

The monetization exercise uses two estimated reductions in daily transits: 13,314 avoided transits per day for the short-run shock effect, and 4,613 avoided transits per day for the longer-run post-price effect.

Table A-1. Full scenario monetization grid for estimated Area C traffic reductions after the 30 October 2023 price increase

Effects Type	Short-run shock effect			Long-run effect		
Avoided Transit/day	13314			4613		
Panel of VoT Assumption (Central): 9.35 €/hour						
Time-saving scenarios	Low	Central	High	Low	Central	High
Minutes saved per avoided transit	0.5	1	2	0.5	1	2
Daily benefit (€)	1037	2075	4150	359	719	1438
Weekly benefit (€)	5187	10374	20748	1797	3594	7189
Annual benefit (€)	238602	477204	954408	82670	165340	330680
Panel of VoT Assumptions (Low): 5.9 €/hour						
Time-saving scenarios	Low	Central	High	Low	Central	High
Minutes saved per avoided transit	0.5	1	2	0.5	1	2
Daily benefit (€)	655	1309	2618	227	454	907
Weekly benefit (€)	3273	6546	13092	1134	2268	4536
Annual benefit (€)	150562	301123	602247	52166	104332	208665
Panel of VoT Assumptions (High): 12.8 €/hour						
Time-saving scenarios	Low	Central	High	Low	Central	High
Minutes saved per avoided transit	0.5	1	2	0.5	1	2
Daily benefit (€)	1420	2840	5681	492	984	1968
Weekly benefit (€)	7101	14202	28404	2460	4921	9841
Annual benefit (€)	326642	653285	1306570	113174	226348	452696

Notes: This table reports a sensitivity exercise based on two estimated treatment effects derived from the Area C analysis: the short-run shock effect and the longer-run post-price effect. Monetary values are computed by combining three assumptions on minutes saved per avoided transit (0.5, 1.0, and 2.0 minutes) with three urban value-of-time scenarios (€5.90, €9.35, and €12.80 per hour). Weekly values assume 5 charged weekdays. Annual values assume 230 charged days per year. Values are rounded to the nearest euro. The table should be interpreted as a partial welfare valuation of congestion-related time savings rather than as a full cost-benefit analysis, since it does not include full system costs, wider externalities, or direct observed travel-time data. Source: Author’s calculations based on Area C IoT gate data and scenario assumptions.

Source: Author’s calculations based on Area C IoT gate data and scenario assumptions.

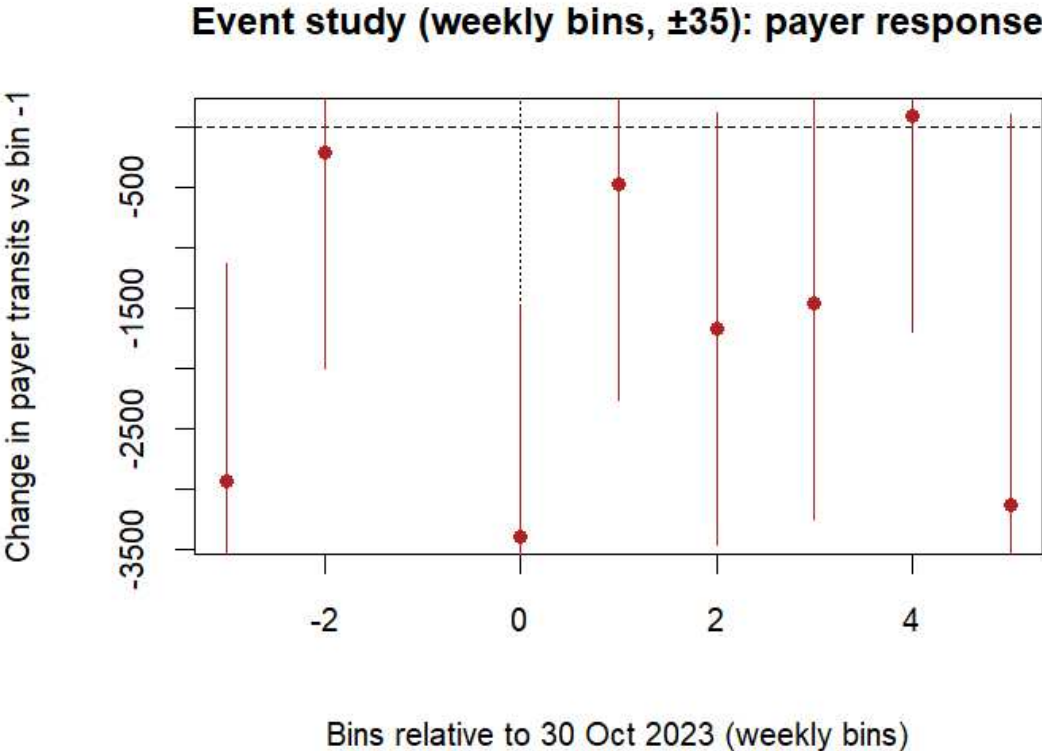


Figure A-1. Event-study estimates for Area C payer transits around the 30 October 2023 price increase

Notes: The figure reports weekly-binned event-study coefficients for payer transits relative to the omitted pre-shock reference bin within a ± 35 -day window around 30 October 2023. Vertical bars represent 95% confidence intervals. The figure is included as an appendix diagnostic to show the timing and short-run dynamics of the payer response to the price increase.

Source: Author’s calculations based on Area C gate data.

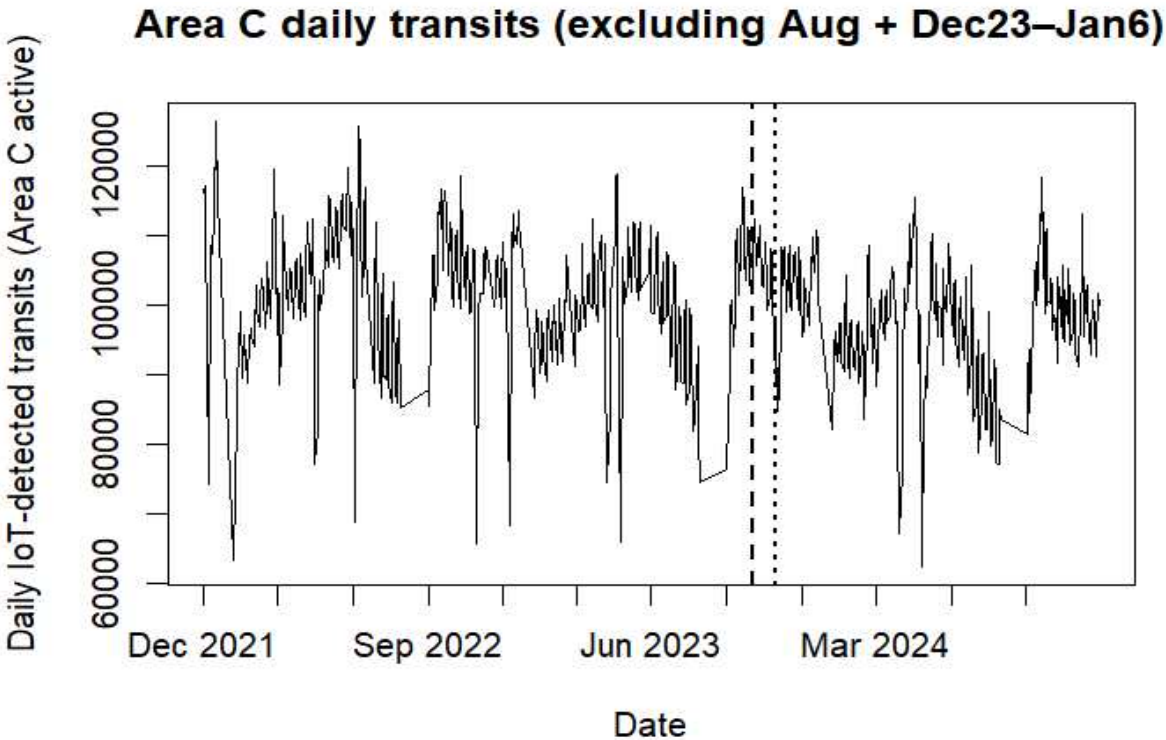


Figure A-2. Daily Area C transits after filtering August and the year-end holiday period

Notes: The series excludes all days in August and the interval from 23 December to 6 January, consistent with the holiday-exclusion rule used in the long-run robustness analysis. Vertical lines indicate the 2 October 2023 rules change and the 30 October 2023 price increase. The figure is reported as a transparency check for the cleaned daily series used in the robustness analysis.

Source: Author’s calculations based on Area C gate data.

Data and code availability:

The R script, all the data inputs, and the processed analysis outputs used to generate the appendix tables and figures were maintained in a separate replication folder during thesis preparation and can be made available by the author upon request.
