Master Thesis

Shared Autonomous Electric Vehicle: A Simulation Study on Operational Efficiency and Service Dynamics

Belodedova Elena

Date of Birth: 30.11.2000
Student ID: 12338844

Subject Area: Digital Economy
Supervisor: Amin Anjomshoaa

Date of Submission:

Institute for Data, Process and Knowledge Management, Vienna University of Economics and Business, Welthandelsplatz 1, 1020 Vienna, Austria

DEPARTMENT FÜR WIRTSCHAFTSINFORMATIK UND OPERATIONS MANAGEMENT DEPARTMENT OF INFORMATION SYSTEMS AND OPERATIONS MANAGEMENT

Contents

List of Figures

List of Tables

Abstract

Shared autonomous electric vehicles (SAEV) represent a transformative shift in urban mobility, promising enhanced flexibility and sustainability. This study investigates the impact of integrating SAEVs into a taxi-hailing aggregator, focusing on their operational efficiency and environmental benefits. Based on New York City trip data from high-volume for-hire vehicle (FHV) services, a simulation is conducted to model how such vehicles would operate, exploring fleet performance, environmental impact, and revenue while considering constraints such as charging infrastructure, fleet size, and real-time demand fluctuations in an urban setting. The simulation evaluates key metrics, including CO2 emissions, utilization, completed trips rate, and net revenue, comparing them with traditional 15-hour driver-based taxis and 24-hour hailing service. Our results confirm the initial expectations, demonstrating that SAEVs outperform traditional taxi services across all key metrics, irrespective of fleet size. The simulation reveals that, on average, net revenue increased by over 250%, while CO2 emissions were reduced by more than 100% compared to the benchmark services. In comparison to 24-hour hailing services, SAEVs clearly offer advantages in terms of revenue and environmental impact, though they do exhibit a decline in operational efficiency. Importantly, this decline can be mitigated through appropriate scaling of the SAEV fleet.

1 Introduction

The taxi industry represents a vast and rapidly expanding market. In major cities like Beijing, the number of operational taxis exceeds 74,000 [28]. Globally, the ride-hailing segment is projected to generate approximately US$172.50 billion in revenue in 2024, with a compound annual growth rate (CAGR) of 4.33% expected between 2024 and 2029. By 2029, the market is forecasted to reach over US$212.75 billion, with the number of users surpassing 2.3 billion [22]. These figures highlight not only the economic scale of the industry but also emphasize the pressing need for innovation to ensure its sustainability, operational efficiency, and environmental responsibility.

Despite its economic value, the taxi sector imposes a disproportionately high environmental burden. Although taxis represent a small fraction of the total vehicle population, they contribute approximately 16.78 million tons of CO2 emissions annually in China alone [28]. This amount is nearly double that of the national railway system, four times that of public buses, and 15 to 25 times higher than emissions from private vehicles on a per-unit basis [28]. These statistics underscore the urgent need for decarbonization within the urban mobility sector.

One of the widely proposed solutions to reduce emissions is carpooling. As a result, recent research efforts have increasingly focused on shared ride models, which are believed to reduce traffic congestion, lower environmental impact, and improve transportation efficiency. However, empirical studies show that carpooling often has a limited effect on overall platform revenue. This is due not only to the operational difficulty of effectively matching passengers, but also to cannibalization effects, where lower-cost pooled rides draw demand away from higher-priced solo rides [14]. In other words, while carpooling may expand the market and improve vehicle utilization, it can reduce revenue from traditional services, thereby limiting its profitability. As a result, carpooling is not well suited for growing profits, which is usually the main goal for companies.

This tension between traditional ride-hailing services and carpooling explains why major platforms like Uber and Lyft have scaled back or even discontinued their carpooling offerings in certain markets. Although carpooling contributes to reducing emissions by improving driver efficiency, it does not generate sufficient revenue to be a sustainable business model in its current form.

At first glance, it may seem that pricing mechanisms could solve the revenue challenges of carpooling. In reality, this is not the case. On ride-hailing platforms, like in other two-sided marketplaces, pricing is designed mainly to balance demand and supply [5], rather than to maximize revenue per trip.

Price adjustments also directly affect service quality, for example by influencing passenger waiting times or driver earnings, both of which are critical for platform stability. As a result, simply changing prices cannot overcome the economic limits of carpooling without risking declines in efficiency or service quality. This highlights the importance of broader optimization approaches, where key indicators such as waiting time, fare levels, and pickup reliability are jointly considered.

At the same time, the taxi and ride-hailing industries are reaching a structural ceiling in terms of revenue growth. Most of the existing optimization tools-such as route-matching algorithms, demand forecasting, or incentive schemes for drivers-have already been implemented by major platforms, but their contribution to revenue growth has largely plateaued.

In some markets, the main barrier is not dispatch efficiency anymore, but simply the lack of drivers. Russia already shows signs of this: there are not enough active drivers to grow further. Although driver's licenses are split almost evenly between men and women, female participation in ride-hailing remains very low. To cope with this, companies even run campaigns to attract more women into driving. For example, Uber have introduced campaigns to recruit more female drivers [17]. These efforts illustrate that the industry is pressing against the boundaries of its current model, with little room left for organic growth. As a result, firms and policymakers are increasingly looking toward disruptive innovations to unlock the next wave of transformation in urban mobility.

One potential solution, though costly and still under development, is the integration of Shared Autonomous Electric Vehicles (SAEVs), which may offer a way to address both environmental and revenue-related challenges simultaneously. The concept of SAEVs combines three elements. First, "Shared" - the vehicle is not owned privately, but operated within car-sharing or ride-hailing systems, ensuring higher utilization rates. Second, "Autonomous" - the vehicle can operate without a human driver, which removes one of the largest cost components in traditional ride-hailing services and therefore creates substantial potential for revenue growth. And third, "Electric" - the vehicle runs on electricity and relies on charging infrastructure, thereby reducing emissions compared to conventional fuel-powered fleets. Taken together, these features make SAEVs a promising technological innovation: they not only offer environmental benefits, but also directly address the profitability issue that carpooling alone cannot solve, primarily by eliminating driver-related costs and enabling more flexible fleet operations.

The convergence of vehicle electrification and autonomous driving is expected to reshape urban mobility in near future. Electrification, supported by breakthroughs in electric motor efficiency and battery technology, offers substantial benefits such as reduced greenhouse gas emissions, improved energy efficiency, and lower operational costs [13].

However, transitioning to such transformative systems comes with several challenges. These include economic barriers such as high upfront costs and limited adoption incentives, environmental concerns surrounding battery recycling, renewable energy sourcing, and infrastructure constraints such as insufficient charging networks and urban-rural disparities [15]. Addressing these challenges requires detailed insights into the operational dynamics and impacts of integrating SAEVs within urban settings.

This research adopts a simulation-driven approach to explore the integration of SAEVs within a taxi-hailing platform. The simulation is designed to dynamically alter key operational parameters-such as fleet size and charging infrastructure - to assess their impact on performance metrics, including fleet efficiency, revenue, CO2 emissions and completed rides. Furthermore, by comparing SAEVs operations to traditional taxi services and taxi-hailing platform, the study evaluates potential revenue uplift and environmental benefits. Through this approach, the research aims to provide actionable strategies to optimize SAEV deployment and shed light on existing barriers. While environmental impact and operational efficiency have been widely estimated in previous studies [11, 20, 28, 6, 4], this research places a specific focus on the economic impact, as financial concerns remain one of the primary obstacles to SAEV adoption and a less-explored aspect of autonomous mobility.

2 Theoretical Background

2.1 Urban mobility platforms

In recent years, app-based ride-hailing platforms such as Uber and Lyft have significantly changed how urban transportation operates. Unlike traditional taxi services, which use fixed pricing and limited dispatching, these platforms connect riders and drivers through real-time algorithms. They act as two-sided marketplaces that continuously manage supply and demand using location data, pricing models, and dispatch algorithms.

Unlike traditional taxi services, platforms such as Uber have redefined the logic of urban transport by treating it not just as a logistics problem, but as a two-sided marketplace. Rather than simply building an app and setting high prices to maximize short-term profits, taxi-hailing platforms recognized a core dynamic: the value of the platform to one group (drivers or riders) depends on the sustained participation of the other [18]. This interdependence creates what is known as a network effect, where growth on one side reinforces growth on the other. To achieve long-term profitability, the platform must first reach and maintain marketplace efficiency - a state in which rides are matched quickly, wait times are minimized, and prices remain competitive. This goal is achieved through the coordinated use of dispatching algorithms and dynamic pricing mechanisms.

Dispatch algorithms on ride-hailing platforms no longer follow simplistic first-in-first-out rules but instead rely on batching strategies, where multiple rider requests and available drivers are collected over short time windows and then matched via optimization algorithms [27]. This batching approach leads to better spatial matching, reduced rider wait times, and improved vehicle utilization. On the pricing side, platforms apply dynamic pricing (DP) mechanisms - often called surge pricing - that adjust fares in real time based on localized supply-demand imbalances. These pricing strategies are grounded in market equilibrium models, which aim to set prices that balance rider participation, driver supply, and en route time [5]. When jointly optimized, matching and pricing mechanisms not only improve platform revenue and capacity utilization but also mitigate structural inefficiencies such as the Wild Goose Chase, where drivers chase fleeting demand signals across the city, leading to increased empty travel and poor rider experience. Empirical results based on Uber data demonstrate that joint optimization of dispatching and pricing algorithms significantly improves platform performance: it reduces price volatility, increases driver utilization - the share of time drivers spend actively transporting passengers - by over 10%, and enhances overall welfare (defined as rider utility minus driver compensation costs) by approximately 2% [27]. These findings highlight the operational advantages of integrated algorithmic control.

In addition to algorithmic innovations, platform-based models also improve operational efficiency compared to traditional taxis. A study by Cramer and Krueger (2016) found that Uber drivers spend approximately 30% more time and 50% more distance on trips with passengers compared to taxi drivers in the same cities [8]. These gains are attributed to superior matching algorithms, higher platform density, and the use of flexible labor rather than fixed vehicle supply. The results suggest that even before the deployment of automation, digital ride-hailing already offers significant improvements in asset utilization and supply responsiveness.

However, while these early innovations brought substantial benefits, further optimization has become increasingly difficult. As noted by Yan et al. (2020), the most accessible improvements - such as real-time routing and surge pricing - have already been implemented and scaled [27]. New proposals for improving efficiency tend to focus on fine-tuning existing mechanisms or integrating them more closely, such as combining batching with pricing algorithms. Although these strategies yield measurable gains, they are often modest in scale and fall short of the disruptive impact seen when ride-hailing platforms first emerged. This highlights a growing challenge in the field: it is no longer sufficient to optimize isolated components, and at the same time, it is increasingly difficult to develop novel mechanisms that meaningfully outperform the current system. As a result, research and innovation in this space are shifting toward marginal improvements rather than transformative breakthroughs.

2.2 Service Quality Indicators and Optimization

Ride-hailing platforms are inherently multi-objective systems, where economic outcomes and service quality are jointly shaped. In the SAEV context, four families of performance indicators are particularly relevant: (i) *utilization efficiency* (share of time vehicles spend serving passengers rather than idling or charging), (ii) *completed trip rate* (the proportion of rider requests successfully served), (iii) *net revenue* (aggregate economic return from trips), and (iv) *environmental impact* (measured through operational CO_2 emissions). These dimensions reflect both user-facing service quality and system-level sustainability.

In operational terms, pricing in conventional ride-hailing is often used as a rapid lever to balance demand and supply [5], while fleet sizing and vehicle allocation function as slower structural levers that determine long-term availability. In SAEV fleets, the absence of driver churn reduces the role of dynamic pricing, but charging downtime introduces new availability constraints. As a result, utilization, completed trip rate, and revenue become more directly dependent on charging infrastructure and repositioning strategies.

For this reason, results are reported jointly across economic, operational, and environmental dimensions. Emphasis is placed on identifying trade-offs, such as higher utilization versus reduced service coverage, rather than on optimizing a single metric in isolation.

2.3 Minimum Fleet Problem

As algorithmic innovation within ride-hailing platforms reaches a plateau, researchers have increasingly turned to questions of system-level design and resource allocation, with a particular focus on how to deploy fleets most efficiently. A central question in this line of work is the Minimum Fleet Problem (MFP) - how many vehicles are needed to meet a given level of demand while maintaining service quality. This problem lies at the intersection of operational efficiency and cost minimization.

One of the most influential studies in this domain, conducted by Spieser et al. (2014), demonstrated that in an idealized urban setting like Singapore, a centrally dispatched fleet of autonomous vehicles could satisfy total demand with up to 60% fewer vehicles than the current taxi system, even without ride-sharing [21]. This finding illustrated the enormous efficiency potential of autonomy and centralized control. However, their model relied on strong assumptions - such as full knowledge of future demand, negligible traffic, and unlimited computing power - that do not hold in practice.

Further advancing the field, Vazifeh et al. (2018) used large-scale mobility data from New York City to solve the MFP under real-world urban dynamics [26]. Their approach combined origin-destination flows with time-varying demand to determine the smallest fleet size capable of serving all trips without delay. Remarkably, they found that only around 30% of the current taxi fleet would be needed if trips were optimally assigned across time and space. Their model provides a data-driven, network-aware lower bound, revealing the extent of overcapacity in traditional systems and highlighting the potential efficiency gains from algorithmically managed fleets.

2.4 Carpooling and Share-rides

With much of the operational optimization potential already exploited, platforms are increasingly turning to product innovation as a means to grow revenue. One prominent example is the introduction of shared ride options, such as carpooling and ride-pooling, which aim to increase vehicle occupancy and expand service offerings. However, as the following section will show, the economic impact of these models is mixed, and their effect on platform revenue is far from straightforward.

Although shared rides may increase the total number of passengers served per hour, research has shown that this does not necessarily translate into higher platform revenue. Lin et al. (2022) demonstrate that carpooling introduces a cannibalization effect, where some riders who would otherwise pay for a full-price solo ride switch to a discounted shared option [14]. As a result, the marginal increase in rider volume is often offset by lower average fares, limiting the revenue potential of shared mobility offerings. This substitution effect is particularly pronounced in mature ride-hailing markets, where rider expectations for convenience and speed remain high. Thus, while shared rides improve vehicle utilization and reduce average cost per trip, they may also reduce per-ride profitability for the platform.

Additionally, carpooling requires more complex matching algorithms, spatial forecasting, and pricing coordination to ensure that detours remain acceptable and rider satisfaction is maintained. These systems add complexity to operations but do not always lead to higher profits. As Lin et al. (2022) point out, the benefits of shared rides - such as lower prices or better access - often go to the riders, not the platform. In many cases, platforms earn less per trip when riders choose carpooling instead of more expensive solo rides [14]. As a result, Uber and Lyft have discontinued their carpooling services, despite gains in other areas such as reduced emissions or improved vehicle utilization. This reflects a core challenge in mobility innovation: not all efficiency improvements translate into economic sustainability.

2.5 Integration and Impication of Shared Autonomous Electric Vehicles

The emergence of Shared Autonomous Electric Vehicles (SAEVs) offers a promising path forward for transforming urban mobility systems. By combining vehicle automation with electrification and shared ride capabilities, SAEVs aim to address three critical dimensions of modern transportation: labor cost reduction, emission mitigation, and fleet-level operational efficiency. Unlike traditional taxi services or human-operated ride-hailing platforms, SAEV systems allow for full centralization of dispatch, dynamic repositioning, and optimized charging coordination - making them especially attractive for future urban deployment.

Several simulation-based studies have highlighted the potential advantages of SAEVs. Burns et al. (2012) pioneered early models demonstrating that SAEVs can reduce mobility costs per mile by over 75% compared to personal vehicles when scaled appropriately [4]. More recent research has extended this by incorporating demand variability, energy constraints, and urban traffic conditions. For instance, Chen et al. (2016) showed that electrified autonomous fleets, if managed effectively, can outperform conventional taxis on both environmental and economic metrics, particularly in cities with dense travel demand and charging infrastructure [6].

However, SAEVs also introduce new logistical complexities, especially due to their dependence on electric charging. Hyland and Mahmassani (2020) emphasized that fleet performance declines sharply when charging station congestion is not properly managed, with up to 20% effective capacity loss observed [12]. Similarly, Sumitkumar and Al-Sumaiti (2024) explored the interplay between SAEV deployment, energy consumption, and transportation economics. [23] Their study highlights that to realize the full potential of SAEVs, operational strategies must be aligned with energy management policies. They propose a simulation framework for assessing social and economic trade-offs, offering guidance on how to integrate SAEVs into urban transport in ways that support both sustainability and efficiency goals. This directly supports the approach taken in the present study, which examines system-level behavior under energy, infrastructure, and service constraints.

Beyond operational factors, researchers increasingly recognize the broader societal and policy implications of SAEV adoption. Almaskati et al. (2024) reviewed over 200 studies and highlighted that SAEVs can reduce emissions, improve accessibility, and dramatically cut parking demand [1]. However, they also warned of unintended consequences: increased vehicle miles traveled (VMT), reduced public transit use, and equity concerns if deployment is left unregulated. Their findings emphasize that technology alone cannot ensure sustainability - policy design, pricing schemes, and governance structures will determine the actual societal outcomes.

Expanding this perspective, Milakis et al. (2017) proposed a ripple-effect model of autonomous vehicle impacts, structured into three layers [16]. First-order effects, such as improved road safety and fuel economy, are mostly positive when SAEVs are electrified and shared. However, second-order effects, like induced demand due to lower travel costs and convenience, may partially offset sustainability gains. Third-order effects -including long-term shifts in land use, car ownership, and labor markets - remain speculative but potentially transformative. These layered outcomes suggest that SAEVs can either support or hinder broader mobility goals depending on their integration context.

2.6 Conceptual Discussion of SAEVs

Beyond the quantitative results of simulation, Shared Autonomous Electric Vehicles (SAEVs) raise broader conceptual questions concerning their role in urban mobility systems, social impacts, and integration into sustainable city planning. This section outlines several of these dimensions, linking the technical findings of this thesis to ongoing debates in transportation research and policy.

SAEVs can be understood not merely as a fleet optimization problem, but as a component of the emerging paradigm of Mobility-as-a-Service (MaaS). Under MaaS, different transport modes are integrated into a seamless digital platform, allowing users to select from public transit, ride-hailing, carsharing, and micro-mobility options (Cohen & Kietzmann, 2014; Shaheen & Cohen, 2020) [7, 19]. In this context, SAEVs could function either as competitors to public transport or as complementary feeders to high-capacity services. The simulation results in this thesis highlight that fleet and charging constraints strongly affect reliability; in a MaaS framework, these constraints would influence whether SAEVs are positioned as a primary mode or a last-mile connector.

One of the most debated aspects of SAEV adoption concerns its impact on employment in the transport sector. Conventional taxi and ride-hailing services employ hundreds of thousands of drivers worldwide. Replacing them with autonomous fleets would reduce labor costs, as reflected in the revenue gains observed in our results, but could also trigger significant job displacement. Fagnant and Kockelman (2015) argue that automation may create new opportunities in vehicle maintenance, fleet management, and charging infrastructure, yet the scale of job losses among drivers is likely to dominate in the short term [10]. This raises questions of social equity and the need for policy interventions such as retraining programs.

Equity has emerged as a central concern in the autonomous mobility literature (Milakis et al., 2017) [16]. SAEVs, if deployed primarily in dense, affluent districts, may exacerbate spatial and social inequalities. The results of this study confirm that charging station distribution significantly affects service availability. Without targeted policy, disadvantaged neighborhoods may face worse service reliability. Almaskati et al. (2020) emphasize that regulatory frameworks must account for spatial justice, ensuring that SAEV systems do not deepen existing transport deserts but instead expand accessibility [1].

Tirachini and Antoniou (2020) highlight the ambiguous relationship between autonomous vehicles and public transport [25]. On one hand, SAEVs could draw passengers away from buses and metro systems, undermining fare revenues and increasing congestion. On the other, they could complement mass transit by providing efficient first- and last-mile connectivity. The simulation framework developed here does not explicitly model intermodal substitution, but the observed sensitivity of waiting times and missed trips suggests that SAEVs are unlikely to replace high-capacity transport modes in peak demand scenarios without substantial oversupply. This points toward a complementary rather than substitutive role in most urban contexts.

Electrification is often justified by its environmental benefits, yet the sustainability of SAEVs must be evaluated holistically. Bauer et al. (2021) show that charging congestion and grid constraints may limit the practical benefits of electrification [2], while Silva et al. (2022) argue for life-cycle assessments that include battery production and recycling [20]. The present study focuses on operational emissions, confirming substantial reductions relative to internal combustion engine fleets. However, broader sustainability considerations, such as rebound effects from induced demand, remain important for long-term evaluation.

Finally, SAEVs are often positioned as part of a "smart city" vision, where data-driven mobility integrates with energy, land use, and environmental systems. Realizing this vision requires governance frameworks that ensure interoperability, data transparency, and public accountability. Without these, SAEV deployment risks being driven primarily by private platforms, with outcomes that may conflict with collective urban goals. The simulation results in this thesis provide a technical perspective on operational constraints, but achieving societal benefits will depend on governance choices far beyond fleet optimization.

2.7 Policy and Societal Implications

The results of the simulation not only provide technical insights into the performance of SAEVs but also carry broader implications for urban transport policy, sustainability, and social equity. This section discusses the potential policy relevance of the findings and highlights issues that extend beyond the technical scope of the model.

A first implication concerns the relationship between SAEVs and existing public transport networks. Studies such as Milakis et al. (2017) argue that autonomous vehicles may either complement or compete with public transit, depending on pricing and integration strategies [16]. If SAEVs primarily replace bus and metro trips, overall congestion and emissions reductions may be limited. Conversely, if they are strategically integrated as feeders to high-capacity transit, SAEVs could enhance accessibility and reduce reliance on private car ownership.

Equity is another crucial dimension. Previous research emphasizes that transport innovations often disproportionately benefit affluent urban cores, while peripheral or low-income neighborhoods face reduced access (Almaskati et al., 2020) [1]. Our results confirm that charging infrastructure availability strongly influences service reliability. If stations are clustered in commercial districts, residents in underserved areas may experience longer waiting times and higher rates of missed trips. Policymakers should therefore consider subsidies or zoning regulations to ensure equitable distribution of charging facilities and SAEV coverage.

The environmental benefits of SAEVs depend not only on vehicle technology but also on the decarbonization of electricity grids. Bauer et al. (2021) highlight that without sufficient renewable penetration, large-scale electrification may shift emissions from tailpipes to power plants [2]. This creates an urgent need for joint transport-energy planning: investments in SAEVs must be coordinated with investments in clean energy generation, grid capacity, and demand management.

Market competition also plays a central role. As Vazifeh et al. (2018) and Spieser et al. (2014) demonstrate, coordinated dispatch dramatically reduces the required fleet size [21], [26]. However, in unregulated markets with multiple competing platforms, such coordination may be absent. Our results suggest that without regulation, oversupply of vehicles could reduce utilization, while undersupply would harm service reliability. Policymakers may thus need to explore regulatory frameworks, such as fleet caps, congestion pricing, or incentives for pooling, to balance efficiency and competition.

The COVID-19 pandemic has highlighted the vulnerability of urban transport systems to sudden demand shocks. SAEV fleets, unlike conventional taxis, cannot flexibly scale driver supply to meet surges. This implies that resilience planning-through redundant capacity, dynamic pricing, or integration with emergency mobility plans-will be essential. Research such as Almaskati et al. (2020) stresses that future mobility systems must be stress-tested not only for average conditions but also for crisis scenarios [1].

2.8 Previous Simulation Modelling in SAEV research

Taken together, the literature indicates that the success of SAEVs hinges on more than technological viability. Effective deployment requires careful coordination across infrastructure, operations, policy, and user behavior. To investigate these interactions systematically, researchers have increasingly turned to simulation modeling as a primary tool for evaluating SAEV systems under realistic urban conditions. Given the scarcity of large-scale real-world deployments, simulation enables controlled exploration of how different fleet configurations, charging logistics, and policy levers affect system performance and societal outcomes.

Yet many studies model SAEVs in isolation, assuming either full demand compliance, optimal routing, or instantaneous vehicle rebalancing. For example, the influential study by Vazifeh et al. (2018) assumes perfect coordination, no charging constraints, and no idle time [26]. Similarly, energy-related simulations often impose simplified rules for recharging, treating station availability and queuing as static or exogenous variables [6].

A particularly relevant contribution is the simulation study by Bauer et al. (2021), which evaluates the cost, energy use, and environmental impact of SAEVs operating in Manhattan [2]. Their model integrates vehicle relocation strategies and detailed battery degradation dynamics, and places strong emphasis on optimizing charging infrastructure and life-cycle emissions. However, their simulation relies solely on yellow taxi data from 2015 and assumes a nearest-vehicle dispatch algorithm. It does not incorporate batching or modern ride-hailing strategies, nor does it benchmark SAEV performance against existing human-operated services. As a result, while the technical modeling is detailed, the system-level conclusions may be limited in transferability to today's platform-based mobility landscape.

Another major limitation is the lack of comparative evaluation. While many models analyze SAEV systems in a vacuum, real-world policy decisions require benchmarking against existing services, such as taxis or driver-based ride-hailing. Without such comparisons, it is difficult to assess the true marginal value of automation or electrification. Moreover, previous studies rarely explore performance degradation under constrained infrastructure, leaving unclear how sensitive system efficiency is to real-world bottlenecks.

The preceding literature review highlights how platform-based ride-hailing systems revolutionized urban mobility by treating transportation as a two-sided marketplace governed by algorithmic pricing and dispatch. Early innovations such as batching and surge pricing substantially improved efficiency, utilization, and welfare. However, as these strategies matured, further performance gains became harder to achieve, leading to diminishing returns from algorithmic optimization alone. This prompted researchers to shift focus toward system-level questions such as fleet sizing, capacity utilization, and the integration of shared and autonomous vehicle technologies.

The literature on SAEVs identifies the potential of combining autonomy, electrification, and sharing to address labor costs, emissions, and operational efficiency. However, empirical and simulation-based studies often evaluate SAEVs in isolation, under ideal dispatching or simplified charging scenarios. Many fail to benchmark performance against conventional services or explore how fleet behavior degrades under infrastructure constraints. Additionally, the broader societal consequences - such as equity concerns, transit substitution, and induced demand - are often acknowledged but not directly incorporated into operational models.

This thesis addresses several of these gaps by developing a simulation framework that compares SAEVs against traditional taxi systems using real-world NYC trip data. The model explicitly varies two critical levers: fleet size and charging infrastructure capacity. It captures real-world system frictions such as charging delays, idle time, and rejected trips, rather than assuming optimal rebalancing or omniscient control. By benchmarking SAEVs against driver-operated services, the simulation reveals where SAEVs outperform and underperform across key performance dimensions. Rather than pursuing theoretical optima, the model is used to stress-test deployment scenarios, enabling more policy-relevant insights into how SAEVs can function under urban constraints.

3 Method

3.1 Data Description

This study uses a discrete-time simulation framework to evaluate the performance of SAEVs in urban environments. The framework models SAEV operations by simulating trips, vehicle movements, and charging processes using real-world data from New York City's Taxi and Limousine Commission (TLC). These include records from high-volume for-hire vehicle services (HVFHS, e.g. Uber, Lyft), general FHV bases, and yellow/green taxis. Each dataset provides trip-level information such as pickup and drop-off times, distances, locations, and fares, which are essential for reconstructing demand patterns and calibrating the simulation.

3.2 Key Variables

Table 1 summarizes the most important fields across the different TLC trip record datasets and their role in the simulation framework.

3.3 Example Record Sample

To illustrate the structure of the dataset, Table 2 shows a subset of HVFHS records used in the simulation. Each row corresponds to a single trip completed by one of the licensed high-volume for-hire vehicle services in NYC.

3.3.1 Data Quality Issues

Before running the simulation, several anomalies were identified and addressed:

- Trips with negative or zero duration, where `dropoff_datetime` < `pickup_datetime`.
- Implausible distances (e.g., >100 miles within NYC boundaries).
- Outlier fares, such as zero-fare rides or extremely high values inconsistent with pricing rules.
- Missing or inconsistent flags for shared rides in FHV and HVFHS records, due to differences in reporting across platforms.

These preprocessing steps ensured that the simulation relied on clean, consistent, and realistic input data.

Table 1: Overview of key fields in NYC TLC trip records

`Field`	**Description**	**Relevance for Simulation**
`pickup_datetime`	Timestamp when the passenger entered the vehicle	Defines trip demand arrival process and temporal patterns
`dropoff_datetime`	Timestamp when the passenger left the vehicle	Determines trip duration and vehicle availability windows
`request_datetime`	Time when passenger requested the ride (HVFHS only)	Used to approximate waiting time and measure service reliability
`on_scene_datetime`	Time when driver arrived at pickup (HVFHS only)	Enables calculation of pickup delays and dispatch efficiency
`PULocationID, DOLocationID`	TLC taxi zone IDs for pickup and drop-off (1-263 zones)	Encodes spatial distribution of demand and enables zone-based assignment strategies
`trip_miles / trip_distance`	Reported trip distance in miles	Determines battery consumption in SAEV scenario and trip-level cost
`trip_time`	Trip duration in seconds	Used for temporal vehicle availability and utilization calculation
`base_ passenger_fare, fare_amount`	Base fare before tips, taxes, and surcharges	Core component of revenue calculations
`tips`	Passenger tips (credit card only in yellow/green)	Used for realistic driver-pay benchmarks in baseline scenarios
`driver_pay (HVFHS)`	Net driver compensation, excluding tolls and tips	Important for comparison with SAEVs, where driver costs are eliminated
`tolls, airport_fee, congestion_surcharge, sales_tax, bcf`	Itemized surcharges and fees	Reflected in net platform revenue calculation
`shared _request_flag, shared_match_flag`	Indicates if passenger requested and was matched to a pooled ride (HVFHS only)	Relevant for studying carpooling and comparing with SAEVs

Table 2: HVFHS sample

	request$_{d}atetime$	on$_{s}cene_{d}atetime$	pickup$_{d}atetime$	dropoff$_{d}atetime$	PULocationID	DOLocationID	trip$_{m}iles$	trip$_{t}ime$	base$_{p}assenger_{f}are$	tolls	bcf	sales$_{t}ax$	congestion$_{s}urcharge$	airport$_{f}ee$	tips	driver$_{p}ay$
0	2024-01-01 00:21:47	2024-01-01 00:25:06	2024-01-01 00:28:08	2024-01-01 01:05:39	161	158	2.83	2251	45.61	0.0	1.25	4.05	2.75	0.0	0.0	40.18
1	2024-01-01 00:10:56	2024-01-01 00:11:08	2024-01-01 00:12:53	2024-01-01 00:20:05	137	79	1.57	432	10.05	0.0	0.28	0.89	2.75	0.0	0.0	6.12
2	2024-01-01 00:20:04	2024-01-01 00:21:51	2024-01-01 00:23:05	2024-01-01 00:35:16	79	186	1.98	731	18.07	0.0	0.5	1.6	2.75	0.0	0.0	9.47
3	2024-01-01 00:35:46	2024-01-01 00:39:59	2024-01-01 00:41:04	2024-01-01 00:56:34	234	148	1.99	930	17.17	0.0	0.47	1.52	2.75	0.0	0.0	11.35
4	2024-01-01 00:48:19	2024-01-01 00:56:23	2024-01-01 00:57:21	2024-01-01 01:10:02	148	97	2.65	761	38.67	0.0	1.06	3.43	2.75	0.0	0.0	28.63

The simulation utilizes three datasets: The *Trips Dataset*, sourced from NYC taxi records, contains historical ride-hailing data, including pickup and drop-off locations, trip distance and duration, as well as request timestamps; the *Vehicle Dataset*, which tracks car's locations, battery levels, and operational states (e.g., idle, en route, or charging); and the *Charging Station Dataset*, which contains information on station locations, occupancy, and queuing status.

3.4 Simulation algorithm

The simulation process begins with the random initialization of the SAEV fleet and charging stations, which are distributed across various urban zones to reflect heterogeneous access to infrastructure. Each vehicle is assigned a random starting location and a fully charged battery. Similarly charging stations are placed at randomly selected locations with specified capacities.

Next, trip requests from the FHV dataset are used to simulate SAEV behavior over historical completed rides. The simulation proceeds iteratively in one-minute time steps. At each step T, all trip requests with timestamps in the interval [T,T+1) minutes are selected for processing. Trip requests are processed in batches (see Fig. 1), mimicking real-world dispatch strategies used in industry to improve assignment efficiency and avoid short-sighted matching (wild goose chasing). By considering multiple requests at once, the algorithm can prioritize assigning a vehicle to a more suitable upcoming trip rather than the nearest immediate one. Ride assignment is performed using a greedy two-pass algorithm: in the first pass, the best available vehicle-trip pairs are selected based on the shortest estimated pickup time, ensuring that no vehicle or trip is assigned more than once. In the second pass, the remaining valid pairs, already sorted by estimated pickup time and filtered for availability, are processed iteratively, assigning the next-best pairs.

After each assignment round, vehicle states are updated to reflect progress along the trip, changes in location, battery usage, and availability for future assignments.

Battery consumption is modeled using a linear rule, where the battery level is reduced proportionally to the distance traveled, based on an energy-per-kilometer parameter defined at the start of the simulation. While regenerative braking and acceleration dynamics are not explicitly modeled, the average consumption rate is calibrated to reflect typical urban driving behavior.

Through this design, the simulation reproduces SAEV operations under empirically observed spatial and temporal demand patterns, allowing for realistic modeling of urban mobility flows, vehicle utilization, charging behavior,

```python
def greedy_two_pass_assignment(valid_combinations):
    valid_combinations = valid_combinations.sort_values('est_time_to_pickup').copy()
    first_pass = valid_combinations.drop_duplicates(subset='trip_id', keep='first').drop_duplicates(subset='car_id', keep='first')
    assigned_trips, assigned_cars = set(first_pass['trip_id']), set(first_pass['car_id'])

    remaining_combos = valid_combinations[~valid_combinations['trip_id'].isin(assigned_trips) & ~valid_combinations['car_id'].isin(assigned_cars)
        ].copy()

    second_pass = []

    if not remaining_combos.empty:
        remaining_combos = remaining_combos.sort_values('est_time_to_pickup').copy()

        while not remaining_combos.empty:
            mask = (~remaining_combos['trip_id'].isin(assigned_trips)) & (
                ~remaining_combos['car_id'].isin(assigned_cars))
            if not mask.any():
                break

            valid_rows = remaining_combos[mask]
            if valid_rows.empty:
                break

            best_row = valid_rows.iloc[0]
            second_pass.append(best_row)
            assigned_trips.add(best_row['trip_id'])
            assigned_cars.add(best_row['car_id'])
            remaining_combos = remaining_combos[
                (remaining_combos['trip_id'] != best_row['trip_id']) &
                (remaining_combos['car_id'] != best_row['car_id'])
                ]

    second_pass_df = pd.DataFrame(second_pass)
    best_assignments = pd.concat([first_pass, second_pass_df], ignore_index=True)
    return best_assignments
```

Figure 1: Trips assignment realization

and performance metrics.

A similar simulation structure is applied to two baseline configurations, which model traditional taxi and ride-hailing services operating without battery charging constraints of SAEV operations. The first baseline represents a conventional taxi fleet, where each vehicle operates on a fixed 15-hour shift, reflecting standard labor schedules. In this model, trips are assigned to available vehicles without accounting for battery levels or charging needs, and the cost structure includes driver commissions. The second baseline simulates a 24-hour ride-hailing platform (e.g., Uber or Lyft), where vehicles are continuously available. Like the taxi model, it excludes charging infrastructure and assumes conventional fuel-based vehicles. Both baselines are evaluated under varying fleet sizes, enabling a comparative analysis of SAEVs against existing service models in terms of efficiency, cost, and environmental impact.

3.5 Simulation metrics

To evaluate the effectiveness of Shared Autonomous Electric Vehicles (SAEVs) compared to conventional taxi and ride-hailing services, several performance metrics are defined. Each metric captures a different aspect of system behavior and allows benchmarking across scenarios. Unlike abstract optimization

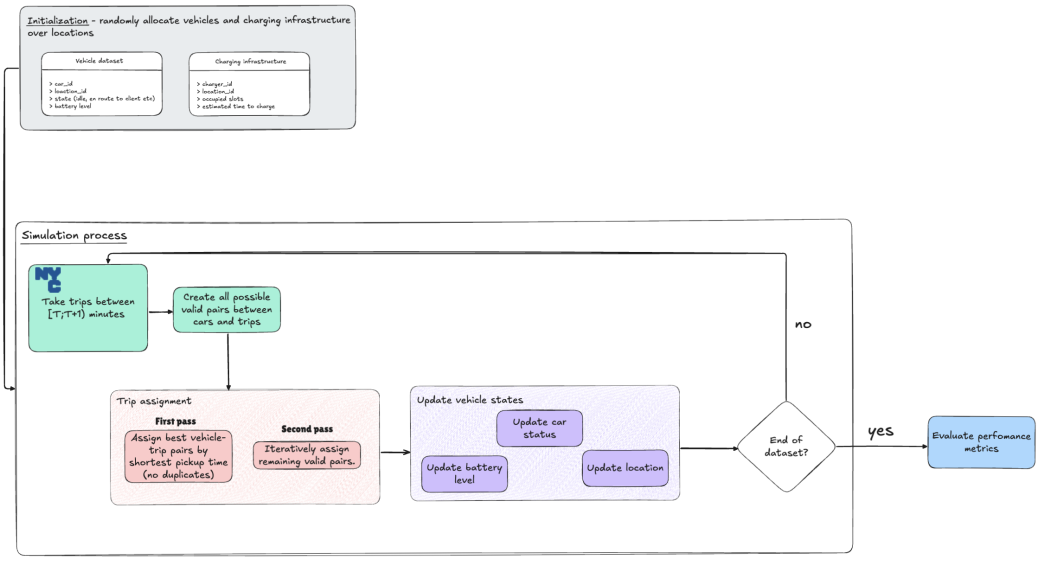

Figure 2: SAEV simulation process

criteria, these metrics reflect operational realities that matter both for passengers, operators, and policymakers.

3.5.1 Utilization Rate

The utilization rate measures the proportion of time vehicles spend with passengers compared to their total time in service. A higher utilization indicates that vehicles are serving riders more frequently and generating value instead of idling or waiting for assignments. In practice, this metric is closely related to the economic efficiency of the fleet: low utilization implies wasted capital investment and insufficient demand matching, while excessively high utilization may signal service saturation and long waiting times for customers. In previous research, utilization has often been used as the primary measure of efficiency for ride-hailing platforms, and the present study follows this tradition while extending it to the SAEV context.

3.5.2 Missed Trips Rate

This metric captures the share of requests that remain unserved because no vehicle was available within the acceptable waiting window. For passengers, it reflects service reliability and directly impacts customer satisfaction. For operators, high missed trip rates imply revenue losses and potential damage to the platform's reputation. Importantly, SAEV systems may exhibit different patterns of missed requests compared to conventional fleets, since charging downtime reduces the number of active vehicles. Thus, missed trips are not only a demand-supply balance indicator but also a proxy for how infrastructure constraints affect real service levels.

3.5.3 Revenue and Profitability

Revenue serves as the central economic metric, representing the income generated from completed trips. In conventional taxi and ride-hailing models, revenue must cover driver wages, fuel costs, and platform fees. By contrast, in SAEV scenarios driver compensation is eliminated, shifting the cost structure heavily toward electricity and vehicle maintenance. This fundamental change means that SAEVs may scale revenue more efficiently with demand, especially in high-volume markets. Profitability also depends on policy factors such as congestion surcharges or subsidies for electric fleets, making it a useful lens for economic comparison.

3.5.4 Environmental Impact

Reducing greenhouse gas emissions is one of the main motivations for electrifying and automating urban mobility. The environmental metric in this study is operational CO_2 emissions per trip, which directly depends on the energy source powering the grid. Although the simulation abstracts away from upstream factors such as battery production or recycling, the operational perspective already highlights the potential of SAEVs to drastically cut emissions compared to internal combustion engine taxis. This aligns with broader climate policies and offers a tangible way to compare different fleet configurations in terms of sustainability.

4 Analysis of the results

4.1 Exploratory Data Analysis

This section focuses on EDA to understand what the data looks like, examine its characteristics, and make sure it behaves as expected.

We observe that the average number of trips increases throughout the week, peaking on Saturday when most people are off work, and subsequently declining on Monday, as shown in Figure 3. Furthermore, there are noticeable peaks in trip volume around 8:00 AM and 6:00 PM, illustrated in Figure 4, which resemble a typical office-hours pattern. This pattern is particularly pronounced on weekdays, whereas on weekends (Saturday and Sunday), the number of trips tends to increase gradually throughout the day, reaching its

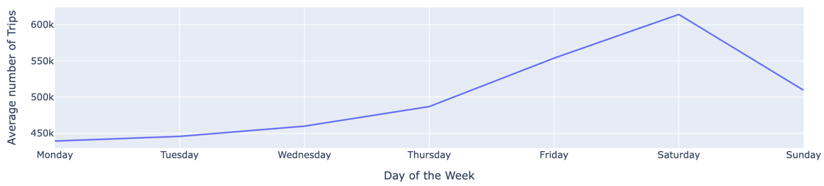

Figure 3: Average number of Trips per day of week

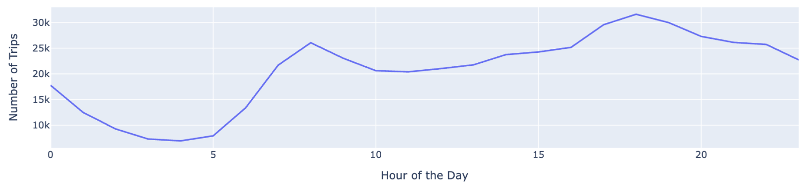

Figure 4: Average number of Trips per hour

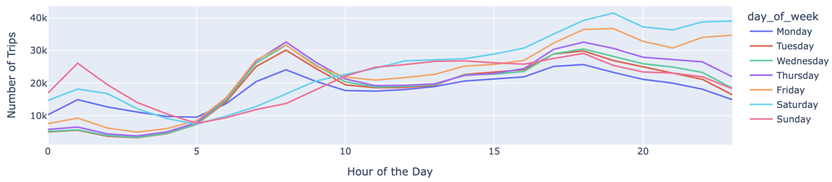

Figure 5: Average number of trips per hour and day of week

maximum around 7:00 PM according to Figure 5. Overall, these dynamics are consistent with expectations and reflect common urban mobility patterns associated with work schedules and leisure activities.

Both the trip distance in the Figure 7 and fare amount in the Figure 6 distributions exhibit strong positive skewness and visually resemble exponential distributions. In each case, the majority of observations are concentrated around 3 miles for distance and $18 for fares - followed by a sharp decline in frequency as values increase. This pattern indicates a clear tendency toward short, low-cost trips, typical for urban environment.

While the median trip distance stays about the same throughout the week, both the average and the range of distances increase slightly on weekends, as shown in Figure 8. This suggests that weekend trips are more varied and often longer, possibly due to travel to the suburbs or across boroughs. In contrast, fare amounts are slightly higher on weekdays and drop a bit on weekends (Figure 9). This difference may be due to heavier weekday traffic, especially in business areas, where delays can make trips more expensive. On weekends, lighter traffic and more spread-out travel likely lead to lower overall fares.

These observations are further supported by spatial trip patterns. As

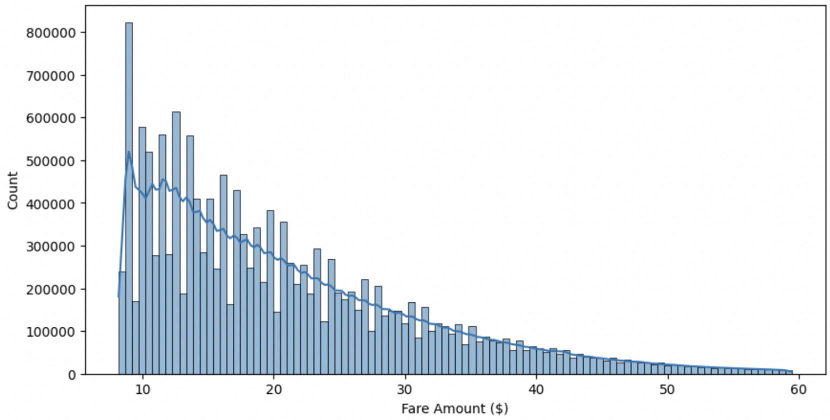

Figure 6: Fares distribution

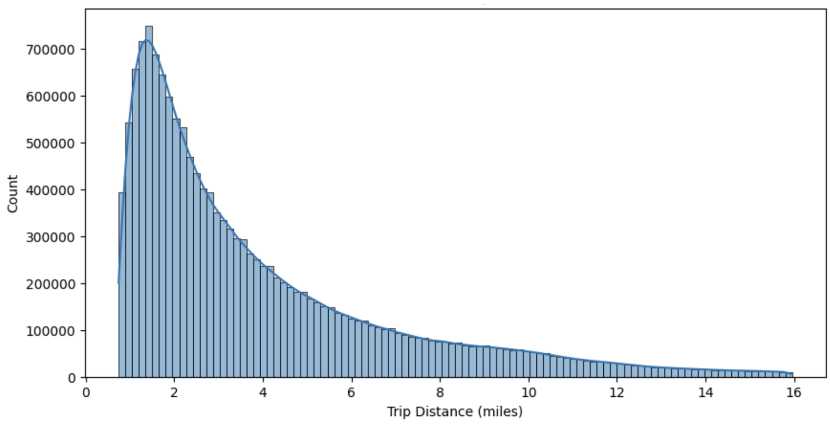

Figure 7: Trips distance distribution

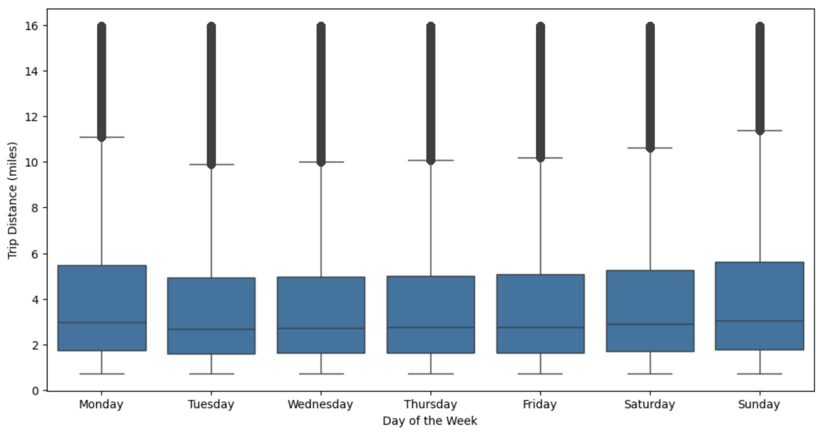

Figure 8: Trips distance boxplots over day of week

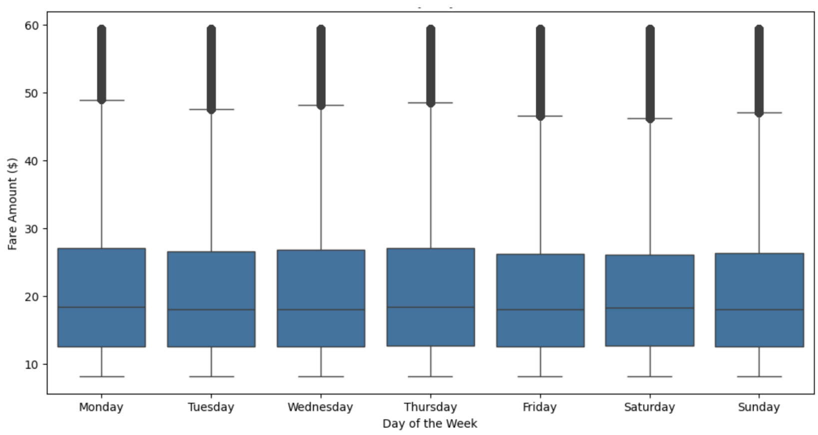

Figure 9: Trips fares boxplots over day of week

Figure 10: Trip Volume by Drop-Off Zone across days of the week

shown in the Figure 10, the volume of drop-offs *Outside of NYC* increases notably on weekends, aligning with the interpretation that longer, inter-borough or out-of-city trips are more frequent during this period. Also, drop-offs at *LaGuardia Airport* gradually increase over the course of the week, peaking on Friday, which likely reflects outbound business or leisure travel behavior. Additionally, *Midtown Center* emerges as one of the most frequent zones for both pick-ups and drop-offs throughout the weekdays, consistent with its role as a major office and commercial district. This concentration of weekday trips toward central business zones aligns with the observed fares and trip distance dynamics.

The dataset shows clear and expected patterns of urban mobility, supporting its appropriateness for further analysis. Weekday and weekend travel behavior, as well as differences in trip distances and fares, reflect typical commuting and leisure trends in a large city. Patterns by time of day and day of the week are especially important and should be carefully considered in future analysis.

4.2 Hypothesis testing

The performance of SAEVs will be statistically compared to that of traditional taxis and hailing platforms in all defined metrics. In other words, a formal hypothesis testing procedure will be conducted. Some of the selected metrics are proportional (for example, utilization efficiency), which means that standard t-test cannot always be directly applied without adjustments. This is due to the fact that when two independent variables are combined into a single metric, a joint distribution arises. Consequently, the assumption of independence, which is fundamental to the Student's t-test, is violated and this formally invalidates the standard variance estimation. Nevertheless, it remains reasonable to rely on the Student's t-test, provided that appropriate adjustments are made. Such adjustments can include the Delta Method [9] or Linearization [3]. In practice, there is no meaningful difference between the two approaches for the purposes of this research [24].

Another issue which arises is the choice of aggregation time when computing the metrics. If the aggregation window is too large (for instance, a week), the number of independent observations becomes small, leading to a loss of statistical power. On the other hand, if the aggregation window is too small (for example, 5 seconds), the data becomes noisier and more susceptible to outliers and extreme fluctuations, which can distort the analysis. Therefore, an appropriate balance must be found to maintain both sufficient sample size and stability of the measured metrics.

An empirical validation will also be provided to assess the performance of the methods and to validate the chosen aggregation window.

4.2.1 Type I Error Validation

When applying statistical tests, an implicit assumption is made: the false positive rate is controlled by the p-value threshold [24]. If this assumption is violated, the p-value can no longer serve as a valid indicator of False Positive Rate (FPR). To ensure validity, it is necessary that the p-values follow a uniform distribution under the null hypothesis. One of the most widely used methods to verify this assumption is A/A testing, which involves the following steps:

- On a dataset structurally similar to that of the planned experiment, randomly split observations into "test" and "control" groups such that no systematic difference should exist between them.
- Repeatedly sample paired subgroups (e.g., 1,000 iterations), applying the statistical test to each pair.
- Analyze the resulting distribution of p-values to verify whether it follows uniform distribution. Additionally, calculate the share of wrong rejection of the null hypothesis - FPR.

To justify the validity of the proposed test under the outlined criteria, the A/A test was conducted (see Fig. 11), mirroring the structure of the actual experimental setup.

For each metric and each aggregation level, the procedure was conducted. The resulting distributions of p-values were examined and found to be approximately uniform. An example of a p-value distribution is presented above in Figure 13.

4.2.2 Minimum detectable effect

Before conducting the tests, it is also important to assess the minimum detectable effect (MDE). However, it should not be used directly; instead, hypotheses can be reformulated based on the MDE. Specifically, when a p-value exceeds the threshold, it is not enough to simply state that there is no effect or that the evidence was insufficient. Rather, it should be interpreted as: if an effect exists, it is smaller than the minimum detectable effect at the chosen power level.

As discussed above, for our metrics we cannot use the raw standard deviation directly (the metrics are ratios/normalized aggregates and may involve

```
def check_aa(
    data: pd.DataFrame,
    nominator: str,
    denominator: Optional[str] = None,
    percentage_groups: Optional[Iterable[float]] = None,
    method: Optional[callable] = DEFAULT_METHOD,
    split_type: Optional[str] = "random",
    split_count: Optional[int] = 2,
    effects: Optional[list] = None,
    N: int = 1000,
    **kwargs,
) -> np.array:
    p_values = []

    df = data.copy()

    if effects is None:
        effects = [0 for _ in range(split_count)]

    combination = list(combinations([i for i in range(split_count)], 2))

    for _ in tqdm(range(N)):
        split_data(df, split_count, split_type, percentage_groups=percentage_groups, **kwargs)

        p_values_per_combo = []

        for pair in combination:
            values_a_nom, values_b_nom = (
                df.loc[df["split"] == pair[0], nominator],
                df.loc[df["split"] == pair[1], nominator],
            )

            if denominator:
                values_a_den, values_b_den = (
                    df.loc[df["split"] == pair[0], denominator],
                    df.loc[df["split"] == pair[1], denominator],
                )

                # подумать над структурой как здесь расчитывать эффекты лучше
                p_values_per_combo.append(
                    method(
                        apply_effect(values_a_nom, effects[pair[0]], denominator=values_a_den),
                        values_a_den,
                        apply_effect(values_b_nom, effects[pair[1]], denominator=values_b_den),
                        values_b_den,
                        **kwargs.get("criteria_args", DEFAULT_ARGS),
                    ).pvalue
                )

            else:
                p_values_per_combo.append(
                    method(
                        apply_effect(values_a_nom, effects[pair[0]]),
                        apply_effect(values_b_nom, effects[pair[1]]),
                        **kwargs.get("criteria_args", DEFAULT_ARGS),
                    ).pvalue
                )

        p_values.append(p_values_per_combo)

    return np.array(p_values)
```

Figure 11: Realization Type I error Validation in Python

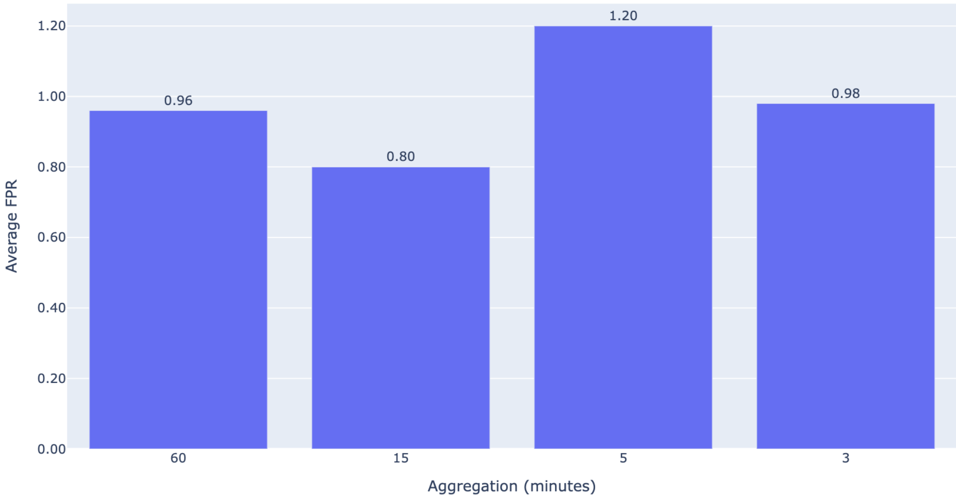

Figure 12: Average FPR

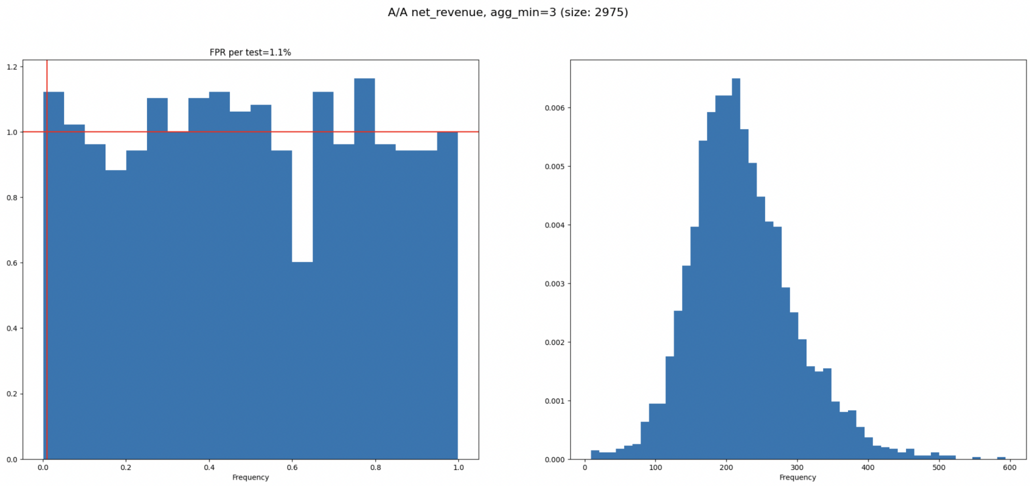

Figure 13: Net revenue p-value distribution under H0

```python
def get_metric_stats(x_0: np.array, y_0: Optional[np.array] = None):

    if y_0 is None:
        y_0 = np.ones(len(x_0))

    mean_nom, var_nom = np.mean(x_0), np.var(x_0)
    mean_den, var_den = np.mean(y_0), np.var(y_0)

    cov = np.mean((x_0 - mean_nom) * (y_0 - mean_den))

    std = np.sqrt(
        var_nom / mean_den ** 2 + var_den * mean_nom ** 2 / mean_den ** 4 - 2 * mean_nom / mean_den ** 3 * cov
    )
    mean = np.sum(x_0) / np.sum(y_0)

    return mean, std
```

Figure 14: Realization of standard deviation through Delta-Method in Python

dependence). Instead, we approximate the sampling variance via the Delta Method; for a ratio $\hat{\theta} = \bar{X}/\bar{Y}$ this yields

$$\mathrm{Var}(\hat{\theta}) \approx \frac{\sigma_X^2}{\mu_Y^2} + \frac{\mu_X^2}{\mu_Y^4}\sigma_Y^2 - \frac{2\,\mu_X}{\mu_Y^3}\mathrm{Cov}(X,Y),$$

which we then use to compute the effective standard deviation and thus the MDE (e.g., MDE $= (z_{1-\alpha/2} + z_{1-\beta})\,\sigma_{\text{eff}}/\sqrt{n}$ for a two-sided test). This approximation is implemented in Figure 14.

As expected, a smaller aggregation window leads to a lower MDE. On average, the MDE is approximately 3%, implying that if no statistically significant difference is observed, any potential effect is likely smaller than this threshold.

4.3 Analysis of Results

4.3.1 Economic and Operational performance analysis

As illustrated in the figure 16, the hailing platform achieves a consistently higher rate of completed trips compared to the SAEV model. This outcome is expected, given that SAEVs must periodically exit service for charging, reducing their availability.

Yet, when examining the net revenue (see Fig. 17), a different trend emerges: the net revenue generated by SAEVs increases more steeply with fleet size than for either traditional taxis or hailing services. This suggests that SAEVs benefit from stronger economies of scale. As the fleet size grows, the marginal gain in revenue per additional vehicle is significantly greater for

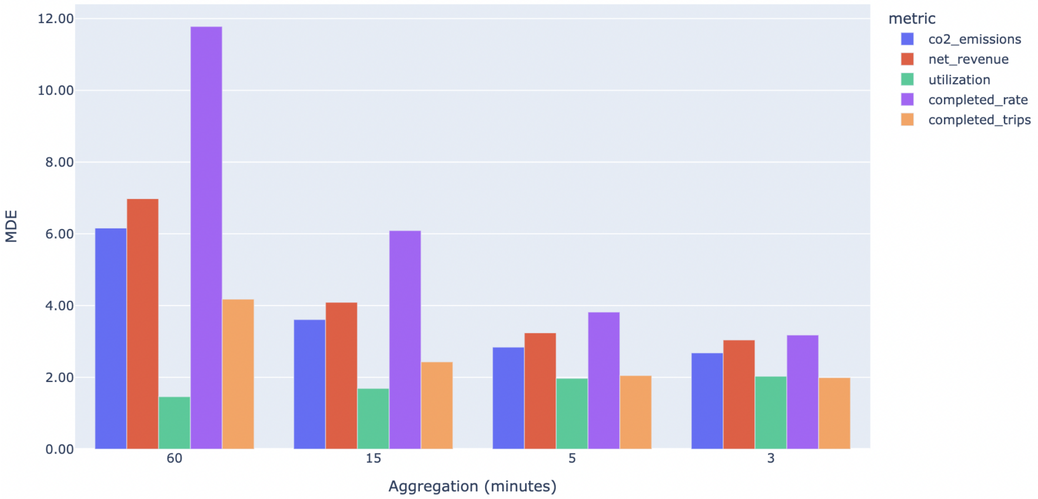

Figure 15: MDE comparison along aggregation windows

SAEVs than for the other two models. This nonlinear scaling effect may be attributed to their ability to operate without labor costs.

SAEVs have slightly higher utilization than the hailing service (see Fig. 18). The lift heavily depends on fleet size and varies from 1% to 11%. In contrast, traditional taxis show very low and almost constant utilization regardless of fleet size.

4.3.2 Environmental impact

As expected, emissions from the SAEV fleet are significantly lower than those from both hailing and traditional services (see Fig. 19). In contrast, hailing services-largely composed of internal combustion engine vehicles-exhibit the highest emissions, increasing steeply with fleet size. Traditional taxis fall in between but also show a steady rise as the fleet expands. These results highlight the strong environmental benefits of electrified fleets, particularly at scale.

It is also worth noting that the environmental performance of SAEVs is closely tied to the energy mix of the electricity grid. In regions with a high share of renewable energy, the shift from internal combustion engines to SAEVs can deliver near-zero operational emissions. By contrast, in areas where coal or natural gas remain dominant in power generation, the environmental benefits are reduced, as emissions are effectively transferred from tailpipes to power plants. Bauer et al. (2021) emphasize that transport electrification policies must therefore be coordinated with broader decarbonization of the electricity sector [2].

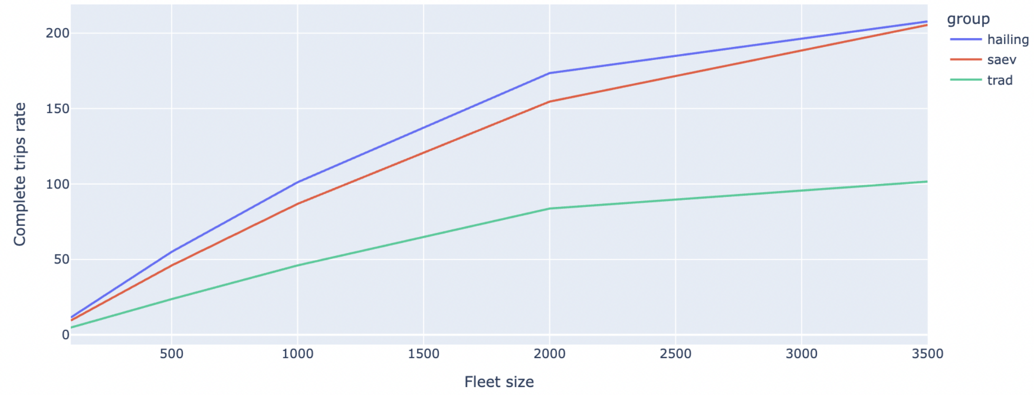

Figure 16: Average completed trips over various fleet sizes

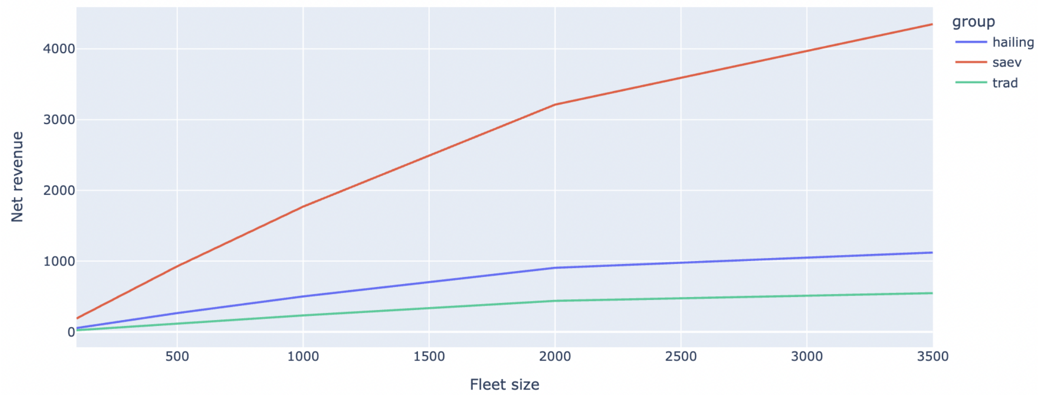

Figure 17: Average net revenue over various fleet sizes

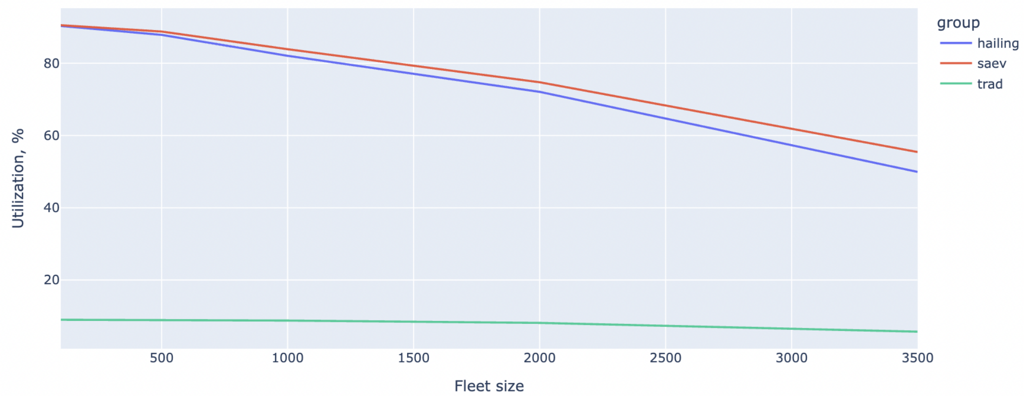

Figure 18: Utilization rate per fleet size

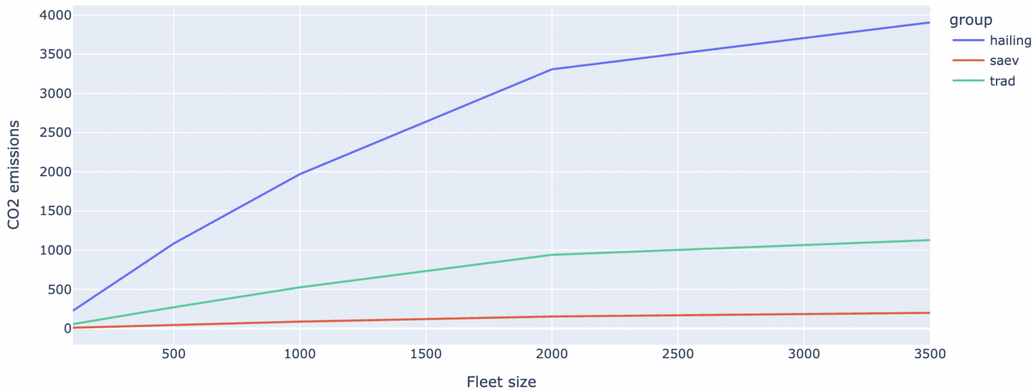

Figure 19: CO_2 emissions per fleet size

Another relevant consideration is the potential rebound effect. As the marginal cost of each trip declines, SAEVs could encourage higher travel demand, partially offsetting emissions reductions. This phenomenon has been discussed in sustainability transition studies (Silva et al., 2022), which warn that induced demand may arise if lower costs stimulate additional trips that would otherwise not have been taken [20]. While the present simulation holds demand constant, future extensions could explicitly model elasticity of demand to assess whether increased accessibility leads to higher aggregate vehicle kilometers traveled.

Finally, beyond operational emissions, the full life-cycle environmental impact of SAEVs must account for battery production, recycling, and disposal. Although outside the scope of this study, such factors are significant: manufacturing electric vehicle batteries is resource-intensive, involving materials such as lithium, cobalt, and nickel. Studies such as Sumitkumar et al. (2024) highlight that the environmental payback period for electric vehicles depends on both battery life and the carbon intensity of electricity used during operation [23]. Integrating such life-cycle perspectives would provide a more holistic assessment of the sustainability of SAEVs.

4.3.3 Hypothesis results

In summary, our analysis shows that the use of SAEVs leads to major improvements compared to both hailing and traditional taxi services, especially in terms of environmental and financial results (see Fig. 20). Across all fleet sizes, SAEVs consistently outperformed hailing and traditional taxi services in terms of both environmental and financial outcomes. CO2 emissions were significantly lower (as illustrated in Fig. 19), and net revenue was substan-

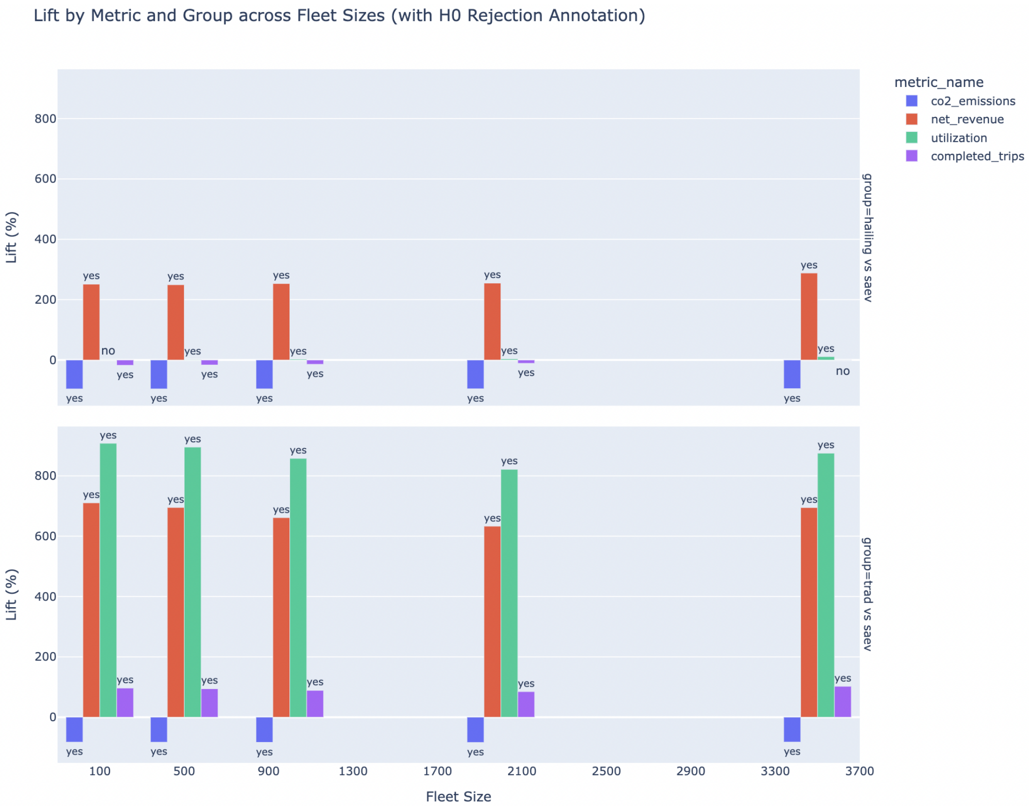

Figure 20: Overall comparison results of simulation

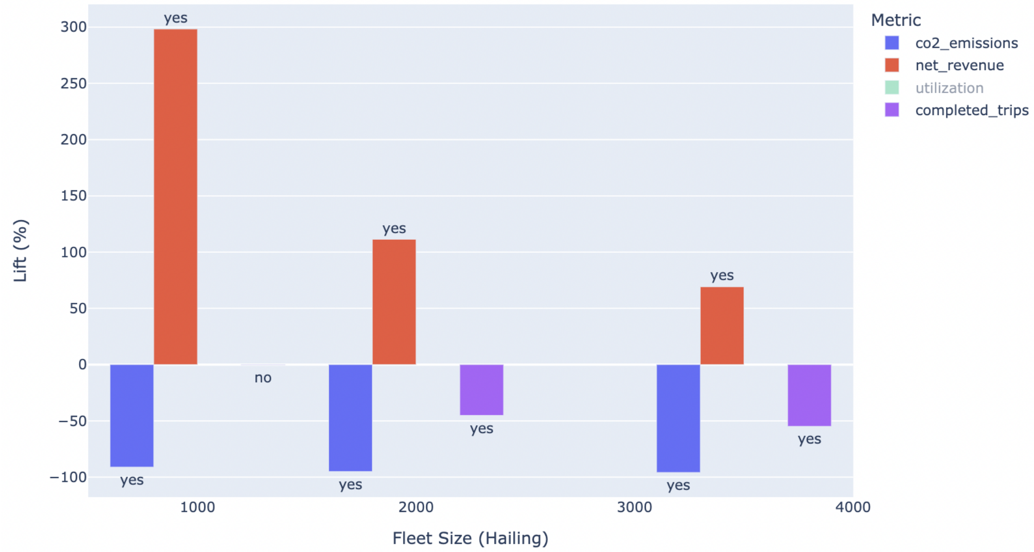

Figure 21: Lift by Metric: Hailing Fleet Sizes vs SAEV Fleet Size = 500

tially higher, regardless of whether there is a shortage of vehicles or enough supply. On average, net revenue increased by more than 250%, while CO2 emissions were reduced by over 100% relative to the benchmark services. Additionally, SAEVs demonstrated a statistically significant improvement in utilization rate, with the effect reaching 11% at a fleet size of 3,500 vehicles. However, not all operational metrics improve. The number of completed trips dropped by about 20% on average. Yet, this gap decreases as more SAEVs are introduced, and it can be largely closed with a larger fleet size.

In addition to within-fleet-size comparisons, we also conducted a cross-fleet-size evaluation only with hailing 24-hour services. Even when the fleet size of the hailing service is up to seven times larger, SAEVs still outperform in terms of net revenue 21. Specifically, despite this substantial disparity in fleet scale, SAEVs deliver a revenue lift of approximately 70%. This finding suggests that SAEVs maintain strong economic advantages even under conditions of supply scarcity, while hailing services require significantly more vehicles to reach comparable financial outcomes.

Overall, the results suggest that SAEV perform better than hailing and traditional taxi services in terms of revenue and environmental impact. Nevertheless, there is a trade-off: some decrease in operational efficiency might occur unless the SAEV fleet is scaled appropriately.

4.3.4 Statistical Robustness of Results

While the descriptive analysis already indicates clear performance differences between SAEVs and conventional fleets, it is important to ensure that these results are not artifacts of sampling variability or model noise. For this reason, the study incorporated several robustness checks.

First, Type I error validation was conducted to confirm that the false positive rate of the statistical tests remained close to the nominal significance level. This involved repeatedly resampling synthetic control and treatment groups from the same dataset. The resulting distribution of p-values was approximately uniform, confirming that the hypothesis testing framework is well calibrated and does not overstate significance. Such checks are critical in transport simulation contexts, where large datasets can create spurious statistical confidence if not carefully validated (Deng et al., 2018) [9].

Second, Type II error validation was performed by injecting synthetic treatment effects into one group and testing whether the statistical procedure was able to detect them. This approach provides an empirical measure of test sensitivity. The results showed that effects of the magnitude observed in the simulation (e.g., revenue increases above 50% or utilization lifts of more than 5%) were consistently detected at high statistical power. Conversely, smaller effects below the minimum detectable effect (MDE) threshold were rarely identified, ensuring that the conclusions reported here focus only on practically meaningful differences.

Finally, sensitivity tests on aggregation windows demonstrated that statistical reliability is affected by temporal granularity. At very small aggregation levels (e.g., five-second windows), noise dominates, while at very large levels (e.g., weekly averages), the number of independent observations declines. The chosen 3-minute window represented a balanced trade-off between sample size and variance stability, consistent with best practices in transportation data analysis (VK Team, 2020) [24].

Together, these robustness checks provide confidence that the observed performance advantages of SAEVs are not merely statistical artifacts but reflect systematic differences rooted in fleet design and cost structure.

5 Conclusion

This thesis evaluated the integration of Shared Autonomous Electric Vehicles (SAEVs) into urban mobility systems using a simulation-based framework. The analysis compared SAEVs against traditional taxi services and 24-hour ride-hailing platforms, focusing on operational efficiency, financial outcomes, and environmental impact.

The results are clear. SAEVs substantially reduce CO_2 emissions, with near-complete elimination of operational greenhouse gases relative to combustion-based fleets. Net revenue was consistently higher, on average more than 250% above benchmark services, even when compared against much larger ride-hailing fleets. Utilization efficiency also improved, reaching up to 11% at a fleet size of 3,500 vehicles. These advantages confirm that SAEVs can outperform conventional mobility services across both environmental and economic dimensions. At the same time, not all operational outcomes improved: the number of completed trips fell by roughly 20% on average due to charging downtime. However, this gap narrows as fleet size increases and can be largely closed with sufficient scaling of the system.

A key contribution of this research is the development of a flexible, open-source simulation framework. It enables systematic testing of fleet sizes, charging infrastructure, and demand conditions while ensuring statistical validity through Type I and Type II error checks. This framework provides a foundation for robust evaluation of SAEV strategies and can be directly reused in other urban contexts.

The findings also highlight important implications. For policymakers and city planners, SAEVs demonstrate strong economies of scale, meaning that early support for deployment may accelerate the point at which the technology becomes both environmentally and financially sustainable. For operators, results suggest that maximizing the benefits of SAEVs requires not only fleet expansion but also careful management of charging infrastructure to avoid service shortages.

Beyond their theoretical and methodological contribution, the results of this thesis have direct practical relevance. Policymakers can use these findings to design regulatory frameworks that prioritize investment in charging infrastructure and provide targeted incentives for early SAEV deployment. For operators, the framework illustrates how revenue potential can be maximized through optimized fleet sizing and proactive repositioning strategies, while also warning of potential service gaps if charging constraints are ignored. Finally, for researchers and technology developers, the open-source simulation tool offers a transparent environment in which new dispatching algorithms, battery technologies, or pricing strategies can be tested under realistic demand conditions.

In conclusion, the evidence shows that SAEVs offer substantial benefits compared to both taxis and ride-hailing services: higher revenues, lower emissions, and better utilization, balanced against manageable operational trade-offs. By making the framework publicly available, this thesis aims to support further academic inquiry and provide a practical tool for decision-makers in shaping the next generation of urban mobility.

6 Limitations and Future Work

Although the simulation provides valuable insights into the relative performance of SAEVs and conventional fleets, several limitations must be acknowledged.

6.1 Methodological Reflection

The methodological design of this study relied on a discrete-time simulation of Shared Autonomous Electric Vehicle (SAEV) operations. While this approach was chosen for reasons of transparency, computational efficiency, and data compatibility, it is important to critically reflect on the assumptions embedded in this choice and its implications for interpretation.

A key decision was to model the system in discrete one-minute intervals rather than in continuous time. This choice enabled straightforward integration with trip data, which is timestamped at minute-level granularity. It also simplified event scheduling, since arrivals, assignments, and charging completions could all be aligned with a common clock. However, this abstraction implies that sub-minute dynamics, such as dispatching delays or traffic-light interruptions, are not represented. In contexts where second-level precision matters, such as highly congested intersections-continuous-time or event-driven simulations might yield more accurate results, though at a higher computational cost.

The simulation is agent-based at the level of individual vehicles and trips, but aggregate in its treatment of travel times between zones. This hybrid choice balanced computational feasibility with behavioral richness. An alternative would have been a fully agent-based traffic microsimulation, where every vehicle movement is modeled in continuous space. While such models provide fine-grained insights into congestion patterns, they require immense data inputs and calibration, often exceeding the scope of thesis-level research. The present approach thus represents a pragmatic trade-off: capturing individual decision processes while abstracting from detailed traffic flows.

Vehicle assignment relied on a greedy matching algorithm, which minimizes pickup times within each discrete step. This method was selected for its simplicity and robustness, ensuring that all feasible matches are resolved quickly. Yet the algorithm is suboptimal compared to advanced approaches such as mixed-integer programming, reinforcement learning, or predictive demand models. Literature on ride-hailing platforms shows that anticipatory repositioning can significantly improve efficiency (Vazifeh et al., 2018) [26]. By not including such strategies, this simulation may underestimate the potential of SAEVs under sophisticated dispatching.

The study relies on New York City's TLC datasets, which are among the most comprehensive mobility records available globally. Their scale, detail, and public accessibility make them uniquely suited to simulation studies. However, the reliance on a single metropolitan area raises questions of external validity. Urban form, energy infrastructure, and cultural attitudes toward automation differ substantially across cities. While NYC offers an extreme case of high-density demand, the conclusions may not transfer directly to mid-sized or car-dependent cities. Future research could replicate the methodology with datasets from other contexts to explore robustness.

Battery dynamics were modeled using constant consumption rates and linear charging functions. This reflects average-case behavior but omits nonlinear charging curves, battery degradation, and variability in driving conditions. Similarly, charging station capacity was modeled as a fixed-slot queue, abstracting from grid-level constraints. While these simplifications enable tractable simulation, they may understate real-world bottlenecks, especially under rapid fleet scaling. A more detailed energy model would require coupling transport simulations with power-system models, as explored in Bauer et al. (2021) [2].

Overall, the methodological choices reflect a balance between realism, transparency, and feasibility. More complex models could capture additional dynamics but risk obscuring causal relationships behind layers of assumptions. The advantage of the present framework is that its mechanics remain interpretable and reproducible, aligning with the principles of open science. At the same time, this comes at the cost of certain simplifications that must be acknowledged when interpreting results.

6.2 Behavioral Assumptions

The analysis assumes identical user demand behavior across service types. In practice, SAEVs may face different demand elasticities due to passenger concerns about safety, comfort, or trust in automation. These behavioral responses could affect adoption rates, waiting-time tolerance, and willingness to share rides. Modeling such heterogeneity requires integrating demand elasticity functions and survey-based behavioral models into the simulation framework.

6.3 Market Competition

The current simulation evaluates a single operator in isolation. In real-world conditions, passengers may choose between multiple platforms (e.g., Uber, Lyft, taxis), creating competition for demand and altering pricing dynamics. Market fragmentation can lower utilization rates and lead to uneven spatial coverage. Future work should include multi-agent simulations where several competing operators interact, possibly under regulatory oversight.

6.4 Infrastructure Simplifications

This study focused primarily on fleet size while holding other system parameters constant. Charging stations were randomly distributed and assumed to provide homogeneous service capacity. In reality, charging infrastructure is unevenly distributed across cities, and congestion at stations can create bottlenecks. Future research could explore optimization of station placement, the impact of fast-charging technology, and integration with vehicle-to-grid (V2G) services.

6.5 Data Limitations

The TLC trip datasets provide extensive coverage of completed rides but do not capture requests that were rejected or abandoned. This limits the ability to model latent demand and true service quality. Moreover, data quality issues such as missing fields, implausible fares, and inconsistent shared-ride flags may introduce noise. Extending the framework with synthetic demand models or calibrated demand curves could improve robustness.

6.6 Environmental and Technical Factors

The environmental analysis only accounts for operational CO_2 emissions during vehicle use. It does not include upstream emissions from electricity generation, battery manufacturing, or end-of-life recycling. Additionally, the simulation does not consider battery degradation, which affects range and lifecycle costs. Future work should adopt life-cycle assessment (LCA) methods and incorporate degradation dynamics to better estimate long-term environmental impacts.

6.7 Governance, Ethics, and Societal Implications

The deployment of SAEVs cannot be evaluated solely from a technological or economic perspective. Ethical, governance, and societal dimensions play an equally important role in determining whether such systems can be adopted successfully and sustainably.

From an ethical and governance standpoint, several challenges arise. The large-scale use of SAEVs depends on extensive collection of geolocation and behavioral data, raising significant privacy concerns. Ensuring secure data handling and maintaining user trust will be critical. Algorithmic dispatching may also lead to fairness issues: if platforms optimize exclusively for efficiency, some neighborhoods may experience systematically lower availability, exacerbating spatial inequalities. Furthermore, SAEVs disrupt established labor structures by removing driver compensation, thereby generating economic efficiency gains while simultaneously displacing traditional driving jobs. Broader questions of liability and accountability also remain unresolved, especially regarding responsibility in cases of accidents or service failures.

At the same time, the empirical results of this thesis highlight clear opportunities for policy action. Demonstrated economies of scale suggest that early public support through subsidies, pilot projects, or favorable regulation could accelerate SAEV deployment until operations become financially self-sustaining. Near complete elimination of CO_2 emissions underscores the relevance of SAEVs for urban climate policy, making investments in charging infrastructure not an optional add-on but a central prerequisite for system reliability. Finally, the social consequences of job displacement call for anticipatory measures such as re-skilling programs and the creation of new employment opportunities in fleet maintenance, charging management, and system supervision.

Beyond direct policy measures, SAEVs also entail broader societal implications. The removal of human drivers not only reshapes the labor market but also changes the social fabric of urban mobility. Adoption will depend on whether passengers feel safe, perceive the technology as reliable, and view the service as accessible across different neighborhoods. Public trust will therefore require transparent communication about safety, clear liability frameworks, and inclusive service coverage.

In the long term, SAEVs should be understood not merely as a transport innovation but as part of a larger transition toward sustainable, technology-driven urban living. Their success will hinge on the integration of transport planning with climate goals, digital governance, and equity considerations. This highlights that SAEV deployment is not just a technical intervention but a catalyst for systemic change in how cities are organized and experienced.

Taken together, these considerations illustrate that SAEVs represent both an opportunity and a responsibility: technological gains are achievable, but long-term success will depend on how effectively ethical risks are mitigated, governance frameworks are established, and supportive public policies are implemented in ways that align efficiency with equity and sustainability.

6.8 Future Research Directions

While this thesis has provided empirical insights into the performance of Shared Autonomous Electric Vehicles (SAEVs) under different fleet sizes and infrastructure constraints, several avenues remain open for further investigation. Extending the simulation framework and empirical scope would allow for a deeper understanding of the systemic role SAEVs can play in future urban mobility.

One promising direction is to study SAEVs not as a stand-alone system, but as part of a multimodal ecosystem. In many metropolitan areas, ride-hailing demand is highly correlated with public transport usage, particularly for first- and last-mile connections. Integrating SAEVs with metro, bus, and commuter rail systems could therefore yield different results than those observed in a purely competitive setting. Future research could incorporate intermodal transfers and evaluate whether SAEVs complement or cannibalize existing public transport services. This line of inquiry would also have important implications for urban planning and transport equity, especially in areas where access to high-frequency public transport is limited.

Another avenue relates to the interaction between SAEVs and the broader energy system. This thesis measured environmental impact in terms of operational CO_2 emissions, but did not model the dynamic interaction between fleets and electricity grids. Future work could integrate SAEV charging with smart-grid management, accounting for renewable energy availability and time-of-use electricity pricing. Advanced scenarios such as Vehicle-to-Grid (V2G) could be tested, where SAEV fleets not only consume but also temporarily supply energy back to the grid during peak demand. Such an extension would connect the debate on sustainable mobility with that on energy transition and climate resilience.

Service equity also deserves closer examination. In practice, mobility platforms tend to concentrate supply in high-demand, high-income districts, which can exacerbate spatial inequality. Future SAEV simulations could explicitly model service provision across different neighborhoods and evaluate how dispatch algorithms, fleet sizes, or pricing policies affect underserved communities. This perspective links technological optimization with ethical and governance considerations and is especially important if SAEVs are to be integrated into publicly regulated mobility frameworks.

Finally, while this study focused on New York City, replication across different urban contexts would add external validity. European cities, for example, feature denser public transport networks, narrower streets, and stricter regulatory environments. Smaller or less dense urban areas may exhibit different adoption patterns, infrastructure bottlenecks, or cost structures. Extending the framework to multiple cities would allow researchers to disentangle context-specific outcomes from generalizable trends, and thereby improve both scientific understanding and policy relevance.

In summary, future research should not only refine technical assumptions within the simulation framework, but also broaden the analytical perspective to include multimodal integration, energy-mobility interactions, equity concerns, and comparative studies across diverse urban environments. Addressing these dimensions would provide a more comprehensive picture of the opportunities and challenges associated with SAEV deployment, and support the development of evidence-based policies for sustainable and inclusive mobility.

Appendix

The simulation framework developed in this thesis, including code and configuration scripts, is openly available on GitHub for further inspection and replication of results:

https://github.com/enbelodedova/simulation_nyc

The framework is designed to be modular and extensible, allowing researchers and practitioners to vary parameters such as fleet size, charging infrastructure, and vehicle battery capacity. It can therefore be directly reused for testing alternative deployment strategies in different urban contexts.

The empirical analysis in this study relies on publicly available data from the New York City Taxi & Limousine Commission (TLC), covering more than one billion trips across for-hire vehicle platforms such as Uber, Lyft, Via, and Juno. The raw data are accessible at:

https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page

Simulation Parameters

For reproducibility, the key parameters used in the simulation framework are summarized in Table 3. These values reflect both the technical assumptions (e.g., charging times and consumption rates) and the operational constraints imposed in the experiments.

Parameter	**Value / Range**
Fleet size	500 - 3,500 vehicles
Battery capacity	300 km range
Charging time	3,600 seconds (1 hour)
Electricity consumption	1 unit per km
Charging stations	10 (capacity: 5 vehicles each)
Maximum wait time	900 seconds
CO_2 emission factor (ICE taxis)	0.2 kg/km
Simulation time step	1 minute

Table 3: Main parameters used in the SAEV simulation framework.

References

[1] Deema Almaskati, Sharareh Kermanshachi, and Apurva Pamidimukkala. The convergence of emerging transportation trends: A comprehensive review of shared autonomous vehicles. *Journal of Intelligent and Connected Vehicles*, 2024. Forthcoming.

[2] Gordon S. Bauer, Jeffery B. Greenblatt, and Brian F. Gerke. Cost, energy and environmental impact of automated electric taxi fleets in manhattan. *Supporting Information for Manuscript*, 2021. Available as supplementary material; 27 pages, 26 figures and tables.

[3] Roman Budylin, Alexey Drutsa, Ilya Katsev, and Valeriya Tsoy. Consistent transformation of ratio metrics for efficient online controlled experiments. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 55–63, 2018.

[4] Lawrence D. Burns, William C. Jordan, and Bonnie Scarborough. Transforming personal mobility. *The Earth Institute, Columbia University*, 2012.

[5] Juan Camilo Castillo, Dan Knoepfle, and Glen Weyl. Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 241–242, 2017.

[6] T. Donna Chen, Kara M. Kockelman, and Jason P. Hanna. Operations of a shared, autonomous, electric vehicle fleet: Implications of vehicle

and charging infrastructure decisions. *Transportation Research Part A: Policy and Practice*, 94:243–254, 2016.

[7] Boyd Cohen and Jan Kietzmann. Ride on! mobility business models for the sharing economy. *Organization & environment*, 27(3):279–296, 2014.

[8] Judd Cramer and Alan B Krueger. Disruptive change in the taxi business: The case of uber. *American Economic Review*, 106(5):177–182, 2016.

[9] Alex Deng, Ulf Knoblich, and Jiannan Lu. Applying the delta method in metric analytics: A practical guide with novel ideas. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 233–242, 2018.

[10] Daniel J Fagnant and Kara Kockelman. Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. *Transportation Research Part A: Policy and Practice*, 77:167–181, 2015.

[11] Javed Farhan and T Donna Chen. Impact of ridesharing on operational efficiency of shared autonomous electric vehicle fleet. *Transportation Research Part C: Emerging Technologies*, 93:310–321, 2018.

[12] Michael Hyland and Hani S. Mahmassani. Dynamic autonomous vehicle fleet operations: Optimization-based strategies to assign and reposition autonomous vehicle fleets with electric charging constraints. *Transportation Research Part C: Emerging Technologies*, 114:597–617, 2020.

[13] Mohamed Khaleel, Abdussalam Ali Ahmed, Abdulagader Alsharif, and Mohamed Milad Beiek. Technology challenges and trends of electric motor and drive in electric vehicle. *International Journal of Electrical Engineering and Sustainability (IJEES)*, 1(1):41–48, 2023.

[14] Jinan Lin, Tingting Nian, Vidyanand Choudhary, Bo Tan, and Cheng Gong. Lower-tier products: Friends or foes? the impact of carpool on ride-hailing platforms. *The Impact of Carpool on Ride-hailing Platforms (July 26, 2022)*, 2022.

[15] Lucy Maybury, Padraig Corcoran, and Liana Cipcigan. Mathematical modelling of electric vehicle adoption: A systematic literature review. *Transportation Research Part D*, 107, 2022.

[16] Dimitris Milakis, Bart van Arem, and Bert van Wee. Policy and society related implications of automated driving: A review of literature and directions for future research. *Journal of Intelligent Transportation Systems*, 21(4):324–348, 2017.

[17] Onde App. Women in ride-hailing: Common issues and how companies can tackle them. https://medium.com/@ondeapp/women-in-ride-hailing-common-issues-and-how-companies-can-tackle-them-44d569163635, 2023. Accessed: 2025-08-22.

[18] Marc Rysman. The economics of two-sided markets. *Journal of economic perspectives*, 23(3):125–143, 2009.

[19] Susan Shaheen and Adam Cohen. Mobility on demand (mod) and mobility as a service (maas): Early understanding of shared mobility impacts and public transit partnerships. In *Demand for emerging transportation systems*, pages 37–59. Elsevier, 2020.

[20] Óscar Silva, Rubén Cordera, Esther González-González, and Soledad Nogués. Environmental impacts of autonomous vehicles: A review of the scientific literature. *Science of The Total Environment*, 830:154615, 2022.

[21] Kevin Spieser, Kyle Treleaven, Rick Zhang, Emilio Frazzoli, David Morton, and Marco Pavone. Toward a systematic approach to the design and evaluation of automated mobility-on-demand systems: A case study in singapore. In *Road Vehicle Automation*, pages 229–245. Springer, 2014.

[22] Statista. Ride-hailing - worldwide | market forecast. https://www.statista.com/outlook/mmo/shared-mobility/ride-hailing/worldwide, 2024. Accessed: 2025-04-01.

[23] Rathor Sumitkumar and Ameena Saad Al-Sumaiti. Shared autonomous electric vehicle: Towards social economy of energy and mobility from power-transportation nexus perspective. *Renewable and Sustainable Energy Reviews*, 197:114381, 2024.

[24] VK Team. Practitioner's guide to statistical tests. https://vkteam.medium.com/practitioners-guide-to-statistical-tests-ed2d580ef04f6f38, 2020. Accessed: 2025-04-06.

[25] Alejandro Tirachini, Emmanouil Chaniotakis, Mohamed Abouelela, and Constantinos Antoniou. The sustainability of shared mobility: Can a

platform for shared rides reduce motorized traffic in cities? *Transportation Research Part C: Emerging Technologies*, 117:102707, 2020.

[26] M.M. Vazifeh, P. Santi, S. Resta, S.H. Strogatz, and C. Ratti. Addressing the minimum fleet problem in on-demand urban mobility. *Nature*, (557):534–538, 2018.

[27] Chiwei Yan, Helin Zhu, Nikita Korolko, and Dawn Woodard. Dynamic pricing and matching in ride-hailing platforms. *SSRN Electronic Journal*, 2020. Available at SSRN: https://ssrn.com/abstract=3258234.

[28] Lihui Zhang, Sijie Wu, Kaiqiang Leng, and Jinrong Zhu. Carbon emission reduction pathways for taxis based on the whole life cycle. *Transportation Research Part D: Transport and Environment*, 132:104236, 2024.