Dynamic and Robust Allocation of On-Street Parking for Passenger and Delivery Vehicles

Alp Darendeliler* Barış Yıldız* Fikri Karaesmen*

Abstract

***Problem definition*:** Curb space has long been a scarce public resource in automobilized cities, serving competing uses for passenger parking and commercial activities. The rapid growth of e-commerce and home deliveries, combined with increasing urban density, has further intensified pressure on this already constrained resource, making effective curbspace management a critical policy challenge. Yet, in practice, curb access remains largely unmanaged, relying on first-come–first-served rules or, more recently, static loading-zone allocations that fail to exploit real-time sensing and control capabilities to improve utilization. In this study, we address this gap by examining the dynamic parking space allocation problem (DyPARK) from the perspective of public authorities that regulate curb access, with the explicit objective of enhancing urban sustainability and social welfare. ***Methodology/results*:** We formulate DyPARK as an admission-control problem for a stochastic service system with heterogeneous users. Our methodology accommodates general phase-type representations of parking durations, allowing empirically supported, non-exponential service-time behaviour to be incorporated. To address uncertainty in demand and parking characteristics, we develop a robust optimization framework that relies on a novel machine-learning-based performance approximation, enabling tractable evaluation of robust policies when classical dynamic programming approaches are infeasible. The framework is calibrated using detailed curbside parking data from Istanbul and validated with large-scale data from Melbourne. Numerical experiments show that dynamic curb management substantially reduces congestion- and emission-related externalities—by more than 40% relative to unmanaged access and by up to 21% compared to static allocation schemes—particularly under demand uncertainty. ***Managerial implications*:** The results provide critical insights for public agencies on how flexible, data-driven curbside governance can improve the societal value generated by scarce public space. We identify when investments in sensing, monitoring, and enforcement technologies are justified by reductions in traffic delays and emissions, and when simpler policies suffice.

Keywords: curb-side management, urban sustainability, parking optimization, dynamic allocation, robust optimization

*Department of Industrial Engineering, Koç University, İstanbul, Turkey. Emails: adarendeliler@ku.edu.tr, byildiz@ku.edu.tr, fkaraesmen@ku.edu.tr

1 Introduction

For decades, finding a parking spot has been a familiar, and frustrating ritual of city dwellers. While this chronic inconvenience is hardly novel to cities, the relentless march of urbanization has escalated it into an acute crisis. This challenge, traditionally viewed through the lens of passenger vehicles, has recently been compounded by the exponential growth of e-commerce, which has drastically inflated the demand for logistics parking (World Economic Forum, 2024). Yet, the specific needs of delivery vehicles are frequently overlooked in urban planning. Consequently, last-mile delivery operators are often forced into inadequate curbside access, resorting to practices like double parking, exacerbating the traffic congestion and increased pollution from idling vehicles. Reflecting the operational reality, UPS reportedly allocated $33.8 million in 2018 just for covering parking fines in New York City, instructing drivers to double-park when legal loading zones are unavailable (Baker, 2019). Such double-parking practices significantly exacerbate urban congestion, contributing to the disproportionate congestion impacts of freight trucks, which account for only about 7% of total traffic yet are associated with roughly 28% of congestion in the U.S. (Humes, 2019). The resulting congestion imposes substantial economic costs, estimated at approximately $74 billion in 2024 due to lost time alone (INRIX, 2024).

In response to these mounting pressures, transportation and city planning agencies are increasingly turning their attention to the complex challenge of curb space allocation between passenger use and logistics parking. While both passenger vehicles and delivery vehicles compete for the same limited resource, they exhibit markedly different behaviours and constraints in their curbside use. Passenger vehicles generally have greater flexibility, often able to park slightly farther from their destinations and stay for extended durations, such as for shopping or dining. In contrast, logistics vehicles operate under tight schedules, require immediate proximity to drop-off points, and typically occupy curb space for much shorter periods (often just a few minutes) (Kim and Wang, 2021). Effective policies must account for these distinct characteristics of passenger and logistics parking to mitigate the negative externalities associated with elevated emissions and increased traffic resulting from the lack of parking space.

Practically, two broad policy approaches are observed in curb space management. The first is the status quo, which remains largely ignorant of the growing demand for curbside use by logistics vehicles and the resulting externalities. This approach follows a first-come, first-served logic and does not involve any explicit allocation or prioritization decisions. It is by far the most widespread practice in cities today, as the predominant share of urban curb space is still managed (or rather, left unmanaged) in this way, without systematic differentiation between vehicle types or purposes. The second approach involves a static allocation of certain curb spaces to logistics vehicles, an approach that has become increasingly common in urban areas through the designation of loading zones or delivery bays. While such static assignments can help reduce double parking and improve delivery reliability, they frequently fail to accommodate the highly variable and stochastic nature of delivery operations. Logistics vehicles may arrive in short intervals, requiring more dedicated bays than available. Yet, unlike passenger vehicles, they typically occupy curb space only for a few minutes,

leading to rapid turnover. Consequently, static assignments can cause temporary congestion when demand spikes, but they also lead to underutilization at other times, with reserved bays remaining empty while nearby passenger parking demand goes unmet. In effect, static allocation introduces rigidity into a system that inherently requires flexibility, limiting the overall efficiency of curb space use.

In contrast, our study proposes a dynamic allocation policy, where curb spaces are not permanently dedicated to a specific vehicle type, but rather allocated in real time based on the prevailing demand and system conditions to better utilize the scarce curbside capacity. The central research question we address is therefore: under which conditions is such a dynamic policy beneficial, and to what extent does it outperform static or unmanaged alternatives? Answering these questions is critical for city planners and policymakers, as it provides the basis for determining when the additional investments required to enable dynamic control capabilities—such as sensing, monitoring, and enforcement technologies—are justified by the resulting efficiency and sustainability gains.

Given the stochastic and dynamic nature of the parking demand from passenger and logistics vehicles, optimizing curb space allocations can be conceptually framed as an admission control policy for a queueing system, where arriving vehicles compete for limited service capacity—in this case, curb access. However, translating this abstraction into practical policies presents significant methodological challenges. Classical queueing models often assume exponentially distributed service times, which rarely reflect the nature of real curbside occupancy. Moreover, even when general service time distributions are adopted, widespread practical implementation faces a critical challenge: the parameters governing these distributions are often uncertain or imprecisely estimated. As a result, any fixed parameterization of service times may lead to poor allocation strategies. An important contribution of our study is the development of a new methodological framework for real-time curb space allocation that can operate with empirically supported distributions of parking durations while incorporating a robustness perspective. Specifically, our methodology is designed to work with Coxian-distributed service times, a member of the phase-type distribution family that is substantially more flexible than the exponential distribution in capturing diverse stochastic processes. Built upon a Coxian distribution framework, our solution approach can be extended to any admission control system where the service times can be effectively represented by this flexible distribution family. This broadens its applicability to a wide range of settings where the exponential distribution assumption is unrealistic, yet the stochastic behavior of the system can be well captured through an appropriate parameterization of the Coxian distribution. Furthermore, we introduce a robustification approach that leverages machine learning to construct an accurate functional approximation of the system's performance, over which robustness against uncertainties in the governing stochastic parameters can be efficiently imposed.

Next, we present the definition of the problem examined in this paper and describe the Sarıyer (Istanbul) case study, which provides the empirical motivation and context for our analysis.

1.1 Problem definition

We study the dynamic parking space allocation problem (DyPARK) for an on-street parking lot that is accessible for both passenger and delivery vehicles. The arrival and parking duration of both vehicle types are stochastic and independent. When an incoming vehicle of any type does not find an available space, it generates a cost and leaves the system. The cost for the passenger cars is due to the cruising to find an available spot for parking, while the cost for the logistics vehicles is due to the double parking that blocks a lane. In both cases, the cost is calculated as the monetary value of the extra emissions and time losses due to waiting in traffic.

The parking lot is managed by the municipal authority, which aims to efficiently allocate scarce parking spaces and minimize negative externalities (emissions and traffic) resulting from uncontrolled parking rather than maximizing revenues collected for parking. For this purpose, the vehicles' admission is controlled, meaning that an arriving vehicle of any type can be rejected so as to allocate a parking spot for the other vehicle type whose rejection will be more costly to the system.

The objective of DyPARK is to find an optimal admission control policy that minimizes the expected long-run average cost. The admission policy determines whether to admit or reject an arriving vehicle based on the state of the system, consisting of the current number of passenger and delivery vehicles, as well as the vehicles' predicted parking duration and rejection costs.

1.2 Sariyer - Istanbul case study

To validate and test the proposed mathematical models, we collected empirical parking data through manual surveys conducted in an on-street parking area located in the Sarıyer district of Istanbul (Figure 1). The parking area is operated by İSPARK, a municipal affiliate company, and is bounded by three street segments: Şehit Mithat Yılmaz Street, Sarıyer Deresi Street, and Yeni Mahalle Street (marked blue in Figure 1), which provide 15, 15, and 10 parking spaces, respectively. Aligning closely with the practical motivation of our study, the surrounding area features a mix of residential and commercial buildings, and the streets experience high traffic volumes. Consequently, there is a consistently high demand for curbside parking from both passenger and freight vehicles throughout the day.

Collected over two weekdays and one weekend day between 10:00 and 18:00, the dataset includes detailed information on vehicle arrivals and departures, vehicle type (passenger or delivery), and parking behavior (legal or illegal) for each observed vehicle. The complete data set is publicly available. More information can be found in the Online Appendix.

The base scenario for the numerical analyses (Section 4) is based on data from two consecutive working days, since parking demand patterns are similar on working days, but differ on weekends. Table 1 reports the estimated arrival rates for Şehit Mithat Yılmaz (ŞMY), Sarıyer Deresi (SD), and Yeni Mahalle (YM) streets over two working days. To identify the underlying distributions of parking durations for passenger and delivery vehicles, we perform statistical analyses. The results of the Kolmogorov-Smirnov (KS) goodness-of-fit tests indicated that an exponential distribution passes the test for parking durations of delivery vehicles. On the other hand, an exponential

Figure 1: Sariyer study area

distribution fails the test for vehicle parking times. Although several flexible distributions pass the KS test for our data (shown with asterisks in Table 2), we choose the Coxian distribution because it is both highly flexible in fitting diverse data and straightforward to represent within an MDP formulation. In addition to providing a good fit for the empirical parking data used in our study, this choice has a broader benefit: the Coxian distribution can capture a wide range of stochastic processes, which paves the way for applying our methodology to other settings where this family is a good candidate for modelling the underlying stochastic behavior. For the base-case scenario, vehicle-specific arrival rates are estimated as pooled averages across all streets and time slots given in Table 1. To handle data uncertainty, we define interval uncertainty sets in the robust model (Section 3), with the observed overall minimum and maximum as the lower and upper bounds, respectively. Parking duration distributions in the base-case are taken as the statistically best-fitting Coxian and exponential distributions for passenger and delivery vehicles, respectively, and data uncertainty is modelled using moment-based uncertainty sets.

To further validate the above findings, we conducted KS tests using the extensive parking data collected by the city of Melbourne in 2018 (City of Melbourne (2020)). For a particular street (Bourke Street), the data were divided into four-hour clusters on working days between 6:00 and 18:00, corresponding to each street marker separating the street and each month. Analysis of $3,279$ data sets with sample sizes greater than 100 reveals that the null hypothesis, that parking times follow a Coxian distribution, cannot be rejected for 96% of instances, whereas an exponential distribution passes the test in only 16% of instances.

Estimating the Delay and Emission Costs: We use a commercial traffic microsimulation model (PTV Vissim) to estimate traffic delays and vehicle emissions resulting from insufficient parking availability. The simulation is calibrated using real-world traffic flow data for the considered street segments (the data and detailed implementation steps are provided in the Online Appendix). Based on background traffic conditions and the arrival rates of logistics and passenger vehicles searching for parking, the model captures additional waiting times and traffic delays arising from parking

Table 1: Arrival rates (per parking space) for passenger and delivery vehicles over two days

Vehicle Type	**Time Slot**	**ŞMY**		**SD**		**YM**	
		Day 1	**Day 2**	**Day 1**	**Day 2**	**Day 1**	**Day 2**
Passenger	10:30–12:00	0.37	0.48	0.56	0.80	1.13	1.13
Passenger	12:00–14:00	0.46	0.30	1.08	0.73	1.05	0.95
Passenger	14:30–18:00	0.33	0.55	0.72	0.49	0.69	1.46
Delivery	10:30–12:00	0.11	0.03	0.18	0.09	0.13	0.13
Delivery	12:00–14:00	0.02	0.04	0.03	0.13	0.05	0.10
Delivery	14:30–18:00	0.01	0.01	0.08	0.03	0.03	0.03

Note: ŞMY = Şehit Mehmet Yılmaz Street; SD = Sarıyer Deresi Street; YM = Yeni Mahalle Street.

Table 2: Kolmogorov–Smirnov goodness-of-fit test results

Distribution	KS statistic	p-value
Log-Normal*	0.0875	0.080
Weibull*	0.0901	0.066
Coxian*	0.0927	0.054
Gamma	0.0992	0.032
Erlang	0.0992	0.032
Generalized Exponential	0.1021	0.025
Exponential	0.1029	0.024
Normal	0.5035	$< 10^{-48}$

Note: The 5% critical value for $n = 206$ is $D_{0.05} = 0.0948$.

search behavior, curbside interference, and queue spillbacks. In our analysis, delay and emission outcomes are converted into a single aggregate monetary cost using the unit cost values proposed by the European Commission (2019). Table 3 reports the simulation results for average marginal externality costs (aggregated monetary cost of delays and emissions) induced by a cruising passenger vehicle and a double-parked delivery vehicle across different time periods.

Table 3: Marginal externality costs (€), including delay and emissions, caused by a single cruising passenger vehicle and a single double-parked delivery vehicle

Time Period	**Passenger Vehicle**	**Delivery Vehicle**
10:00–12:00	0.055	0.443
12:00–14:30	0.071	1.268
14:30–18:00	0.087	1.761

1.3 Contribution of this study

The main contributions of this study are as follows.

- We propose a solution approach for DyPARK that accommodates parking durations modelled by a Coxian distribution supported by empirical data, offering substantial flexibility in capturing parking duration data while preserving a Markovian structure, with each phase being memoryless. This leads to a better alignment for applying MDP-based models to a broader class of problems with non-exponential durations. Our numerical experiments show that the

Coxian-based MDP model, which more accurately represents parking durations, significantly outperforms the exponential benchmark.

- We propose a novel robust optimization framework to handle uncertainty in vehicle arrival rates and parking times. Since the robust counterpart of DyPARK cannot be solved using dynamic programming algorithms due to its structure, we develop an effective and tractable machine-learning based method. Our numerical experiments on empirical data show that the robust approach achieves substantial cost reductions compared to the nominal model that ignores parameter uncertainty.

- Our extensive numerical experiments based on real-world data reveal several critical insights for practice. First, the results indicate that, relative to the no-reservation benchmark, the suggested approach can reduce negative externalities by more than 40%. When compared with static allocation schemes, the reductions are up to 21%, depending on whether logistics vehicles are permitted to temporarily use regular parking spaces when available (soft allocation) or whether the allocation remains strictly enforced (hard allocation). By jointly considering the number of parking spots managed and the relative share of logistics versus private vehicles, we identify the operational regimes in which the proposed dynamic policy delivers substantial performance improvements and those where its marginal benefits may not justify the additional implementation effort.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 presents the model formulation and derives structural properties of the optimal control policy, as well as introduces a robust optimization approach. Section 4 reports numerical experiments on empirical data, and we conclude in Section 5 with some final remarks.

2 Literature Review

The problem studied in this paper falls under the broader class of parking management, an extensively studied problem in the operations literature. Beginning with early quantitative work in this area, most studies have focused on pricing and reservation decisions as the primary levers for improving parking efficiency (D'Acierno et al., 2006; Lin et al., 2017; Zhang et al., 2023). However, the majority of this literature has primarily considered passenger vehicles, and thus does not account for the specific characteristics of urban delivery fleets, whose parking behavior, turnover, and curbside utilization patterns differ substantially. Our work contributes to a more recent and relatively limited body of research that has begun to adopt a holistic view of parking management by explicitly incorporating freight and service vehicles. This perspective is largely motivated by the exponential growth of e-commerce and last-mile logistics, and micromobility technologies which has dramatically increased the pressure on curbside use (Lim and Masoud, 2024).

Research on combined passenger and logistics vehicle parking has examined equilibrium models for pricing and allocation (Amer and Chow, 2017), reservation-based systems to reduce congestion

and double parking (Burns et al., 2024), and slot assignment in loading areas (Roca-Riu et al., 2015). Parking has also been incorporated into vehicle routing optimization (Le Colleter et al., 2023; Jodiawan et al., 2025). Distinct from these studies, our work adopts the perspective of public agencies managing curbspace and develops a dynamic allocation framework that accounts for stochastic arrivals of both passenger and logistics vehicles, with the objective of mitigating the negative externalities due to double parking by freight vehicles and cruising by passenger vehicles.

Queueing models have been utilized to analyze parking systems. Xiao et al. (2018) and Caliskan et al. (2007) use a queueing model M/M/c/c, also known as the loss model, to describe the occupancy of a parking system and to predict future occupancy. Abdeen et al. (2021) also investigate a loss model to estimate parking availability. Dowling et al. (2020) model a curbside parking area as a network of interdependent queues that drivers join if there is a free parking spot; otherwise, they search for available spots in an adjacent queue. This model is used to estimate the proportion of cruising drivers in a given neighbourhood. Abhishek et al. (2021) study a parking system with both delivery bays and general on-street parking spaces, analysing it using a loss model. In their setting, delivery bays are reserved exclusively for freight vehicles, with general parking available to passenger vehicles and, when necessary, to freight vehicles once bays are full. By contrast, we propose a dynamic admission-control policy that flexibly allocates curbspace between passenger and logistics vehicles, and demonstrate the benefits of such an approach through numerical experiments.

The admission-control policy we develop essentially addresses a sequential decision-making problem. Markov Decision Process (MDP) models are widely used in such contexts, as decisions at one point affect both immediate and future outcomes. In parking applications, MDPs have supported reservation, routing, and admission decisions. For example, Bogoslavskyi et al. (2015) propose an MDP to guide drivers under uncertain occupancy, minimizing expected travel and walking times, while Levin and Boyles (2019) use an MDP to optimize parking reservations and routing to reduce search-induced congestion.

In queueing systems with multiple job types, admission control policies determine whether to accept or reject incoming jobs based on system state and job type, reallocating server capacity to less costly jobs when rejections occur. Such problems are typically formulated as MDPs and solved via dynamic programming. Loss systems, where arrivals are turned away if all servers are busy, have been studied extensively. For instance, Altman et al. (2001) and Ormeci et al. (2001) analyse two-class loss systems with exponential interarrival and service times, characterizing the structural properties of optimal admission policies. Closest to our work, Legros and Fransoo (2024) model a parking lot with passenger and delivery vehicles as a loss system. Their framework always admits delivery vehicles, while passenger arrivals are regulated through pricing and admission control to maximize revenue and maintain service levels for freight vehicles. They formulate the problem as an MDP, assuming exponential interarrival and parking times. By contrast, the real-world data we consider in our study required us to relax this exponential-time assumption, leading us to develop methods that account for more general parking-time distributions and uncertainty in parameter estimation.

Since our model relies on real-world data, transition probabilities in the MDP are estimated empirically and thus subject to estimation errors, which can significantly affect solution quality (Mannor et al., 2007; Nilim and El Ghaoui, 2005). To mitigate this risk, robust control of MDPs has been widely studied (Iyengar, 2005; Nilim and El Ghaoui, 2005; Wiesemann et al., 2013). Building on this line of work, we develop a robust counterpart of our nominal MDP to obtain admission-control policies that are immunized against parameter uncertainty. Unlike the common robust MDP framework—which assumes state-action–specific uncertainty sets that preserve the dynamic programming structure—we focus on uncertainty in parameters such as vehicle arrival rates and parking times. Since these parameters remain constant over the planning horizon, the induced transition probabilities are correlated across states. To address this, we design a novel machine learning–based approach to solve the resulting robust MDP.

3 Model Formulation

Consider an on-street parking lot with m parking spots. Passenger and delivery vehicles arrive according to an independent Poisson processes with rates λ_p and λ_d, respectively. As motivated by our case study, passenger parking times follow a two-phase Coxian distribution with independent exponential phases with rates μ_p^1 and μ_p^2. After the first phase, the process proceeds to the second phase with probability p; otherwise, it ends. Thus, with probability $1-p$, the parking time is exponential with rate μ_p^1, and with probability p, it is the sum of the two phases (see Appendix A for details). Delivery parking times are exponentially distributed with rate μ_d. Note that these assumptions were validated via Kolmogorov–Smirnov goodness-of-fit tests for the case study.

The state of the system is denoted as $x = (x_p^1, x_p^2, x_d) \in S = \left\{x : x_p^1 + x_p^2 + x_d \leq m\right\}$, where x_p^i denotes the number of passenger vehicles in phase $i = 1, 2$, and x_d denotes the number of delivery vehicles. State transitions occur when a vehicle arrives or leaves the parking lot. Upon arrival, a vehicle of any type can be accepted or rejected, where action $\mathbf{a}(\mathbf{x}) = (a_p(x), a_d(x))$ with $a_p(x), a_d(x) = 0$ denoting rejection and $a_p(x), a_d(x) = 1$ denoting acceptance of passenger and delivery vehicles, respectively. When the system is full, that is, $x_p^1 + x_p^2 + x_d = m$, or when a vehicle of any type is rejected, a penalty cost c_p or c_d is incurred for an arriving passenger or delivery vehicle. Let $R_p^\pi(t)$ and $R_d^\pi(t)$ denote the number of rejected passenger and delivery vehicles up to time t under policy π. The goal is to find a policy π to minimize the expected cost per unit time $g(\pi)$ where:

$$g(\pi) = \lim_{T\to\infty} \frac{\mathrm{E}^\pi \left[\int_0^T c_d dR_d^\pi(t) + \int_0^T c_p dR_p^\pi(t)\right]}{T}. \tag{1}$$

Since the system spends an exponential time in each state and the maximum transition rate out of any state is finite, the above problem can be modelled by a discrete-time Markov decision process (DTMDP) using the uniformization method. Without loss of generality, the uniformization rate can be set to $\lambda_p + \lambda_d + m\bar{\mu} = 1$, with $\bar{\mu} = \max\left\{\mu_p^1, \mu_p^2, \mu_d\right\}$, so that the arrival and departure probabilities are equal to their transition rates. Then, the optimality equation can be expressed as

$$
\begin{aligned}
& V\left(x_p^1, x_p^2, x_d\right) - g^\star \\
& \quad = \lambda_p \min\left\{c_p + V\left(x_p^1, x_p^2, x_d\right), V\left(x_p^1 + 1, x_p^2, x_d\right)\right\} \\
& \quad + \lambda_d \min\left\{c_d + V\left(x_p^1, x_p^2, x_d\right), V\left(x_p^1, x_p^2, x_d + 1\right)\right\} + x_p^1 p \mu_p^1 V\left(x_p^1 - 1, x_p^2 + 1, x_d\right) \\
& \quad + x_p^1 (1-p) \mu_p^1 V\left(x_p^1 - 1, x_p^2, x_d\right) + x_p^2 \mu_p^2 V\left(x_p^1, x_p^2 - 1, x_d\right) + x_d \mu_d V\left(x_p^1, x_p^2, x_d - 1\right) \\
& \quad + \left(1 - \lambda_p - \lambda_d - x_p^1 \mu_p^1 - x_p^2 \mu_p^2 - x_d \mu_d\right) \cdot V\left(x_p^1, x_p^2, x_d\right), \quad \text{for} \quad x_p^1 + x_p^2 + x_d \leq m,
\end{aligned} \tag{2}
$$

and $V\left(x_p^1, x_p^2, x_d\right) = \infty$, if $x_p^1 + x_p^2 + x_d = m + 1$, so that no vehicle is accepted when the parking lot is full, i.e., $x_p^1 + x_p^2 + x_d = m$. In the above equation, $V(x)$ is the optimal relative value function for state x, and g^* is the long-run minimum average cost. Any admission control policy that minimizes the right-hand side of (2) is average cost optimal. The optimal policy can be found by applying the value iteration algorithm which recursively solves the above equation for each time period until the value functions converge.

3.1 Structural Properties

In this section, we concentrate on the structural properties of the optimal policy. The following proposition states that delivery vehicles are always admitted to the parking lot as long as there is a free parking spot if a certain condition is satisfied.

Proposition 1. *If* $\frac{\lambda_p}{\mu_d + \lambda_p} \leq \frac{c_d}{c_p}$ *and* $\mu_d > \mu_p^1$*, delivery vehicles are always accepted whenever there is an available parking space.*

Proof. Let w be the upper bound of the difference $V_n(x + e_3) - V_n(x)$, such that $w \geq \max_{x \in S}\{V_n(x + e_3) - V_n(x)\}$, for all $x \in S' = \left\{x : x_p^1 + x_p^2 + x_d \leq m - 1\right\}$, and for all $n \in \mathbb{N}_0$, where $V_n(x)$ is the minimal expected cost with n periods remaining, and e_i is the i^{th} unit vector. Delivery vehicles are always admitted as long as there is a free parking spot if and only if $V_n(x + e_3) - V_n(x) \leq c_d$, $\forall x \in S'$, when n periods remaining.

For $n = 0$, $V_0(x + e_3) - V_0(x) \leq c_d$ is satisfied, since $V_0(x) = 0$, $\forall x \in S$. Assume that $V_n(x + e_3) - V_n(x) \leq c_d$ holds for some n. Then, consider a system that starts in state $x + e_3$ when $n + 1$ periods left, rejects both types of vehicles if the next event is an arrival and then follows the optimal policy. Let $T_{n+1}(x + e_3)$ is the value of following this policy starting from state x. Then,

the following inequality holds

$$
\begin{aligned}
&V_{n+1}(x+e_3) - V_{n+1}(x) \leq T_{n+1}(x+e_3) - V_{n+1}(x) \\
&\leq \lambda_p\big((c_p + V_n(x+e_3) - \min\{c_p + V_n(x), V_n(x+e_1)\}\big) + \lambda_d\big(c_d + V_n(x+e_3) \\
&- \min\{c_d + V_n(x), V_n(x+e_3)\}\big) + x_p^1\mu_p^1(1-p)\left(V_n(x-e_1+e_3) - V_n(x-e_1)\right) \\
&+ x_p^1\mu_p^1 p\left(V_n(x-e_1+e_2+e_3) - V_n(x-e_1+e_2)\right) + x_p^2\mu_p^2\left(V_n(x-e_2+e_3) - V_n(x-e_2)\right) \\
&+ x_d\mu_d\left(V_n(x) - V_n(x-e_3)\right) + \mu_d\left(V_n(x) - V_n(x+e_3)\right) \\
&+ \left(m\bar{\mu} - x_p^1\mu_p^1 - x_p^2\mu_p^2 - x_d\mu_d\right)\left(V_n(x+e_3) - V_n(x)\right) \\
&\leq \lambda_p \max\{c_p, w\} + \lambda_d \max\{c_d, w\} + (m\bar{\mu} - \mu_d)w.
\end{aligned} \tag{3}
$$

Then, to ensure that w is an upper bound, the following inequality can be established

$$\lambda_p \max\{c_p, w\} + \lambda_d \max\{c_d, w\} + (m\bar{\mu} - \mu_d)w \leq w.$$

Assuming that $w \leq c_d$, the above inequality becomes

$$\lambda_p \max\{c_p, w\} + \lambda_d c_d + (m\bar{\mu} - \mu_d)w \leq w. \tag{4}$$

First Case: $c_p \leq w \leq c_d$: Inequality (4) becomes

$$\lambda_p w + \lambda_d c_d + (m\bar{\mu} - \mu_d)w \leq w.$$

By uniformization, $m\bar{\mu} + \lambda_p + \lambda_d = 1$. Then

$$\frac{\lambda_d c_d}{\mu_d + \lambda_d} \leq w.$$

Let

$$\underline{w} = \frac{\lambda_d c_d}{\mu_d + \lambda_d}, (\underline{w} \leq c_d),$$

then

$$c_p \leq \underline{w}, \text{ must hold, implying that } \quad \frac{\mu_d + \lambda_d}{\lambda_d} \leq \frac{c_d}{c_p}.$$

Second Case: $w \leq c_d < c_p$: In this case, inequality (4) becomes

$$
\begin{aligned}
&\lambda_p c_p + \lambda_d c_d + (m\bar{\mu} - \mu_d)w \leq w, \\
&\text{implying} \quad \frac{\lambda_p c_p + \lambda_d c_d}{\mu_d + \lambda_p + \lambda_d} \leq w.
\end{aligned} \tag{5}
$$

Accordingly, by the same reasoning as in the previous case,

$$\frac{\lambda_p}{\mu_d + \lambda_p} \leq \frac{c_d}{c_p} < 1.$$

Third Case: $w < c_p \leq c_d$: Again inequality 5 holds, and accordingly,

$$\frac{\mu_d + \lambda_d}{\lambda_d} < c_p \;\rightarrow\; 1 \leq \frac{c_d}{c_p} < \frac{\mu_d + \lambda_d}{\lambda_d}.$$

The results of the three cases are summarized as follows:

- If $\frac{\mu_d+\lambda_d}{\lambda_d} \leq \frac{c_d}{c_p}$, then $V_{n+1}(x+e_3) - V_{n+1}(x) \leq \frac{\lambda_d c_d}{\mu_d+\lambda_d} \leq c_d$,
- If $\frac{\lambda_p}{\mu_d+\lambda_p} \leq \frac{c_d}{c_p} < 1$, then $V_{n+1}(x+e_3) - V_{n+1}(x) \leq \frac{\lambda_p c_p+\lambda_d c_d}{\mu_d+\lambda_p+\lambda_d} \leq c_d$,
- If $1 \leq \frac{c_d}{c_p} < \frac{\mu_d+\lambda_d}{\lambda_d}$, then $V_{n+1}(x+e_3) - V_{n+1}(x) \leq \frac{\lambda_p c_p+\lambda_d c_d}{\mu_d+\lambda_p+\lambda_d} \leq c_d$,

which imply that Proposition 1 is true. □

Given that Proposition 1 holds, the following proposition ensures that passenger vehicles are also admitted into the system as long as there is an available parking spot with the assumption $\mu_p^2 > \mu_p^1$.

Proposition 2. *If* $\frac{\lambda_p}{\mu_d+\lambda_p} \leq \frac{c_d}{c_p} \leq \frac{\lambda_d+\mu_p^1(p-1)}{\lambda_d}$ *and* $\mu_p^2 > \mu_p^1$*, then both vehicle types are always accepted whenever there is an available parking space.*

Proof. Let w be the upper bound of the difference $V_n(x+e_1) - V_n(x)$, such that $w \geq \max_{x \in S}\{V_n(x+e_1) - V_n(x)\}$, $\forall x \in S'$ and $\forall n \in \mathbb{N}_0$. Passenger vehicles are always admitted when n periods remaining if and only if $V_n(x+e_1) - V_n(x) \leq c_p$, $\forall x \in S'$. This inequality holds for $n = 0$. Assume that it also holds for a given n. Then, consider a system that starts in state $x+e_1$, rejects both types of vehicles if the next event is an arrival, and then follows the optimal policy. Let $T_{n+1}(x+e_1)$ be the value of following this policy starting from state x, with $n+1$ periods left. Then, the following inequality holds

$$\begin{aligned}
&V_{n+1}(x+e_1) - V_{n+1}(x) \leq T_{n+1}((x+e_1) - V_{n+1}(x) \\
&\leq\; \lambda_p\big((c_p + V_n(x+e_1) - \min\{c_p + V_n(x), V_n(x+e_1)\}\big) \\
&+ \lambda_d\big(c_d + V_n(x+e_1) - V_n(x+e_3)\big) + x_p^1\mu_p^1(1-p)\left(V_n(x) - V_n(x-e_1)\right) \\
&+ \mu_p^1(1-p)V_n(x) + x_p^1\mu_p^1 p\left(V_n(x+e_2) - V_n(x-e_1+e_2)\right) + \mu_p^1 p V_n(x+e_2) \\
&+ x_p^2\mu_p^2\left(V_n(x+e_1-e_2) - V_n(x-e_2)\right) + x_d\mu_d\left(V_n(x+e_1-e_3) - V_n(x-e_3)\right) \\
&+ \big(m\bar{\mu} - x_p^1\mu_p^1 - x_p^2\mu_p^2 - x_d\mu_d\big)\cdot\left(V_n(x+e_1) - V_n(x)\right) \\
&\leq \lambda_p \max\{c_p, w\} + \lambda_d c_d + (m\bar{\mu} + \mu_p^1 - p\mu_p^1)w.
\end{aligned} \tag{6}$$

□

The following inequality ensures that w is an upper bound

$$\lambda_p \max\{c_p, w\} + \lambda_d c_d + (m\bar{\mu} + \mu_p^1 - p\mu_p^1)w \leq w. \tag{7}$$

Assuming that $w \leq c_p$, the above inequality becomes

$$\lambda_p c_p + \lambda_d c_d + (m\bar{\mu} + \mu_p^1 - p\mu_p^1)w \leq w. \tag{8}$$

By uniformization, Equation (8) reduces to

$$\frac{c_d}{c_p} \leq \frac{\lambda_d + \mu_p^1(p-1)}{\lambda_d}. \tag{9}$$

If the value functions satisfy the supermodularity property in x_p^1 and x_d, then rejecting a vehicle type associated with the i^{th} component of state x implies that it is also rejected in state $(x + e_j)$ under the optimal policy for $i, j \in \{1, 3\}$ with $i \neq j$. This corresponds to the following inequality:

$$V_n(x + e_1 + e_3) - V_n(x + e_1) - V_n(x + e_3) + V_n(x) \geq 0, \quad \forall x \in S'', n \in \mathbb{N}_0,$$

with $S'' = \{x : x_p^1 + x_p^2 + x_d \leq m - 2\}$. Thus, in state (x_p^1, x_p^2, x_d), there exists a threshold level $S_i(x_i)$ such that a vehicle of type i is accepted if and only if $x_j < S_i(x_i)$, where $i, j \in \{1, 3\}$ and $i \neq j$. This condition is satisfied in a two-class loss system with exponentially distributed interarrival and service times (Altman et al. (2001)). Although we could not prove supermodularity in our setting—where the parking duration of passenger vehicles follows a nonexponential distribution—our extensive numerical experiments did not reveal any counterexamples.

3.2 Optimal Policy Illustration

In this part, we illustrate the optimal policy under the base-case scenario, with the following parameter values from the Sariyer case study: $\tilde{\lambda}_p = 0.74m$, $\tilde{\lambda}_d = 0.07m$, $\tilde{\mu}_p^1 = 0.82$, $\tilde{p} = 0.83$, $\tilde{\mu}_p^2 = 8.16$, $\tilde{\mu}_d = 1.97$, and $m = 10$. Figure 2 shows the optimal decisions for states with two and three passenger vehicles in phase 2. The crosses denote rejection of passenger vehicles, dots indicate acceptance, and the line marks the boundary.

Proposition 1 holds, implying that delivery vehicles are always admitted whenever a parking space is available, whereas access for passenger vehicles is restricted. Specifically, passenger vehicles are rejected only when nine spaces are already occupied. However, even when only a single space remains, admission can still be optimal in states with more delivery vehicles and fewer passenger vehicles, where the effective departure rate

$$x_p^1(1-p)\mu_p^1 + x_p^2\mu_p^2 + x_d\mu_d,$$

is sufficiently high to justify accepting an additional passenger vehicle. This property also holds when comparing the acceptance and rejection regions across the two cases with $x_p^2 = 2$ (left) and $x_p^2 = 3$ (Figure 2).

Let $p_b^p(\pi)$ and $p_b^d(\pi)$ be the blocking probabilities for passenger and delivery vehicles, denoting the probabilities that a vehicle of associated type cannot find a free parking space. Then, using

Figure 2: The optimal admission control decisions in states with $x_p^2 = 2$ (left) and $x_p^2 = 3$ (right)

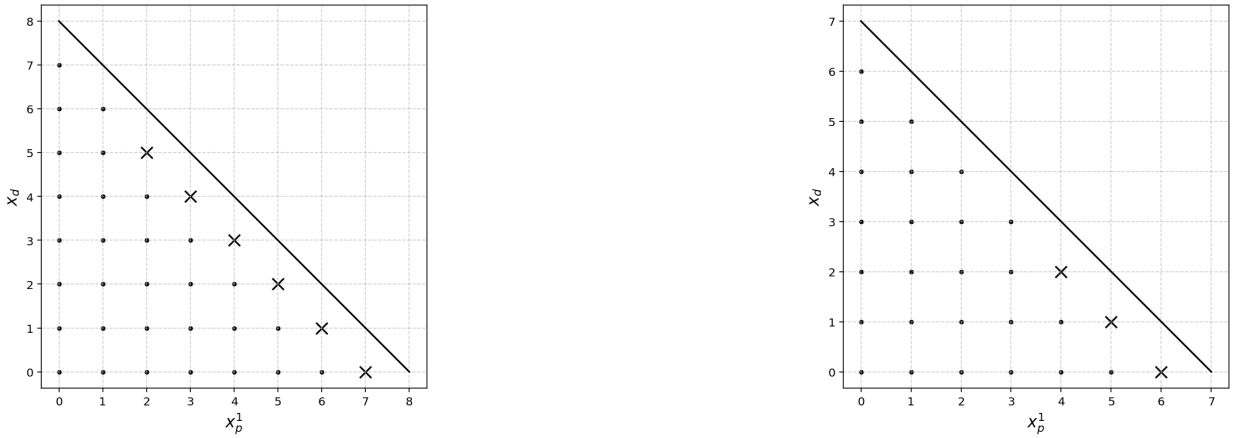

these, the long-run cost rate under a certain policy π can be expressed as

$$g(\pi) = c_p \lambda_p \, p_b^p(\pi) + c_d \lambda_d \, p_b^d(\pi).$$

Under the common practice of no admission control (called all-acceptance policy), the blocking probabilities for the passenger and delivery vehicles are equal, i.e, $p_b^p(\pi^a) = p_b^d(\pi^a) = p_b(\pi^a)$. If the optimal policy rejects vehicle type j in states with available capacity, then $p_b^j(\pi^*) > p_b(\pi^a)$, while $p_b^i(\pi^*) < p_b(\pi^a)$; and the cost difference between the optimal and all-acceptance policies is given by

$$g(\pi^a) - g(\pi^*) = c_p \lambda_p \left(p_b(\pi^a) - p_b^p(\pi^*) \right) + c_d \lambda_d \left(p_b(\pi^a) - p_b^d(\pi^*) \right).$$

In Section 4.1, to quantify the benefit of applying optimal allocation policy, we will analyse the cost and blocking probability differences between the optimal and all-acceptance policies across realistic problem instances.

3.3 Robust Formulation

In the nominal problem (1), the distributional parameters corresponding to the vehicle arrival rates and parking times are assumed to be exactly known and are collected in the vector $\mathbf{z} = \left(\lambda_p, \lambda_d, \mu_p^1, p, \mu_p^2, \mu_d\right)$. Since these parameters are estimated from limited data, they are inherently imprecise and subject to estimation errors. There are also other input parameters that affect the cost but are known with certainty (for instance the number of parking spaces m and the costs c_d and c_p). We denote the vector of input parameters that are not subject to uncertainty as $\mathbf{v}$.

Robust optimization considers formulations where the input parameters that may be subject to uncertainty are allowed to take values in an uncertainty set $\mathbf{Z}$. The robust formulation then seeks a solution to:

$$\max_{\mathbf{z} \in \mathbf{Z}} \min_{\pi \in \Pi} g\left(\pi, \mathbf{z}, \mathbf{v}\right), \tag{10}$$

In the context of our problem, there are many challenges to the above formulation. The standard robust dynamic program approach (Nilim and El Ghaoui (2005)) allows nature to take independent actions in time and states. This gives the adversary (nature) a lot of power and makes a continuous-time infinite horizon MDP problem extremely difficult. It is also not appropriate for our problem because the uncertainty is on the stationary rates of events that are observed and which should remain constant across states and time. We therefore focus on the stationary version of the problem where the adversary is allowed to choose $\mathbf{z} \in \mathbf{Z}$ to be constant for the entire horizon. We can then focus on a stationary version of the problem (10) where the adversary is able to offer the cost maximizing parameters from the uncertainty set against the long-run cost minimizing actions of the controller. Note that the adversary does not need to know the exact policy used by the controller but is aware that the controller will take the cost-minimizing actions.

The other challenge is that to have a tractable optimization formulation $\min_\pi g(\pi, \mathbf{z}, \mathbf{v})$ is needed. To have an exact value for the minimizer, the solution of a dynamic program is required and the result is not in closed form. We overcome this challenge by using tools from statistical learning to estimate $g(\pi^*, \mathbf{z}, \mathbf{v})$ as a function of $(\mathbf{z}, \mathbf{v})$ that can be expressed in closed form. To this end, we solve the dynamic program numerically for a large number of instances with varying parameter vectors $(\mathbf{z}, \mathbf{v})$ and construct a training data set for a learning algorithm. The relationship between the parameters of the problem and the corresponding optimal cost is non-linear since the optimal cost of an MDP is in general a non-linear function of the inputs. We therefore focus on statistical-learning approximations that allow us to estimate a closed form function from the training data: $g(\pi^*, \mathbf{z}, \mathbf{v}) \approx \tilde{g}(\mathbf{z}, \mathbf{v})$. This enables us to replace (10) with its approximate counterpart:

$$\max_{\mathbf{z} \in \mathbf{Z}} \tilde{g}(\mathbf{z}, \mathbf{v}) \tag{11}$$

3.3.1 Interval Uncertainty on Individual Parameters

Assume that we construct prediction intervals $[z_i^{\min}, z_i^{\max}]$ for each input parameter z_i. We can then construct the following uncertainty set:

$$\mathbf{Z}_0 = \left\{ \mathbf{z} \in \mathbb{R}_+^n : \ z_i \in [z_i^{\min}, z_i^{\max}], \ i = 1, 2, ...n \right\}. \tag{12}$$

Let us assume that statistical learning from the data is done using a Generalized Additive Model (James et al., 2023)

$$\tilde{g}(\mathbf{z}, \mathbf{v}) = \sum_i^n \tilde{g}_i(z_i) + \sum_j^l \tilde{g}_i(v_i) \tag{13}$$

where $\tilde{g}_i()$ is the basis function of parameters i, j.

Proposition 3. *Assume that $\tilde{g}_i(z_i)$ is a differentiable function, then:*

1. *If $\tilde{g}_i(z_i)$ is convex on $[z_i^{\min}, z_i^{\max}]$, then*

$$z_i^* = \left\{ \begin{array}{ll} z_i^{\max} & \text{if } \tilde{g}_i(z_i^{\min}) < \tilde{g}_i(z_i^{\max}) \\ z_i^{\min} & \text{otherwise} \end{array} \right\}.$$

2. *If $\tilde{g}_i(z_i)$ is concave on $[z_i^{\min}, z_i^{\max}]$, then*

   (a) *If the unconstrained global maximizer $z_i^{*(u)} \in [z_i^{\min}, z_i^{\max}]$, then $z_i^* = z_i^{*(u)}$.*

   (b) *If the unconstrained global maximizer $z_i^{*(u)} \notin [z_i^{\min}, z_i^{\max}]$, then $z_i^* = \underset{z_i \in \{z_i^{\min}, z_i^{\max}\}}{\arg\max} \tilde{g}_i(z_i)$.*

Proof. Since we have a separable optimization problem we obtain the solution by considering the individual problems for each z_i.

$$\max_{z_i \in [z_i^{\min}, z_i^{\max}]} \tilde{g}_i(z_i), \tag{14}$$

The proposition then follows based on well-known optimization properties of convex and concave functions in a finite interval.

□

In the case where the approximation is a simple linear function (i.e. a regression based on the individual parameters), the above proposition simplifies.

Corollary 1. *Let the estimated function be linear in the parameters:*

$$\tilde{g}(\pi^*, \mathbf{z}, \mathbf{y}) = \tilde{\beta}_0 + \sum_{i=1}^{n} \tilde{\beta}_i z_i + \sum_{j=1}^{l} \tilde{\kappa}_i v_i, \tag{15}$$

Then:

$$z_i^* = \left\{ \begin{array}{ll} z_i^{\max} & \text{if } \tilde{\beta}_i^* > 0 \\ z_i^{\min} & \text{otherwise} \end{array} \right.$$

and the corresponding cost is:

$$g_{\max}(\pi, \mathbf{z}, \mathbf{v}) = \tilde{\beta}_0 + \sum_{i=1}^{n} \tilde{\beta}_i z_i^* + \sum_{j=1}^{l} \tilde{\kappa}_i v_i^*$$

Proof. This follows from part 2 of Proposition 3. □

We can now discuss the following tradeoffs for estimation and optimization for interval uncertainty sets:

- If $\tilde{g}$ is obtained from a linear model as in (15), the optimization problem has a trivial solution from Corollary 1. In addition, the effects of varying individual parameters on the cost is clear. On the other hand, it should be expected that $g(\pi, \mathbf{z})$ is a non-linear function in general and the linear approximation cannot capture the non-linearities.

- If $\tilde{g}$ is obtained from a Generalized Additive Model as in (13), and if the individual basis functions are convex or concave then the optimization problem is still easy from Proposition 3. However, for $\tilde{g}_j(z_j)$ that are neither convex or concave some search algorithm would be required. This model captures some of the non-linear effects of the parameters and is still reasonably easy to interpret.

- A more accurate estimation might be possible if engineered features involving interactions of multiple input parameters are used. For instance, for queueing systems the offered load, measured as a ratio of the arrival rate to the service rate is known to have an important effect. Such engineered features make the estimation more accurate but would render the optimization problem more complicated since separability is lost.

Interval uncertainty formulations are attractive from an optimization perspective and in terms of explainability but are limited in their estimation capabilities when generalized additive models are used. In the next subsection, we propose strengthening the interval uncertainty sets with moment uncertainty constraints.

3.3.2 Improved Uncertainty Sets

In the collected data, there are some event times that are well-estimated by an exponential distribution. The remaining estimation problem is to estimate the mean (or rate) of this distribution. Let $\mathbf{q} = (\lambda_p, \lambda_d, \mu_d)$ denote the parameters of these exponentially distributed random variables. For these three event times, we maintain the interval uncertainty sets:

$$\mathbf{Q} = \left\{ \mathbf{q} \in \mathbb{R}^3_+ : \; q_i \in [q_i^{\min}, q_i^{\max}], \; i = 1, 2, 3 \right\}. \tag{16}$$

On the other hand, the parking times of passenger vehicles do not seem to follow an exponential distribution but we observe that a Coxian random variable appears to be a good fit. This random variable is a function of the three parameters: $\mathbf{y} = \left(\mu_p^1, p, \mu_p^2\right)$. Our data only records the total parking duration since the individual parameters are not observable. To have a better representation of the corresponding uncertainty, we assume that in the robust formulation, the adversary is able to choose the three parameters of the distribution but require that the choice is consistent with the observed first and second moments of the data. Let μ_0 be the sample mean from the data, we propose the following uncertainty set for the parking duration W:

$$Y = \left\{ y \in {\mathbb{R}_+}^3 \;\middle|\; \begin{array}{l} \left| \mathrm{E}\left[W\right] - \mu_0 \right| \leq \gamma_1 \\ \mathrm{E}\left[\left(W - \mu_0\right)^2 \right] \leq \gamma_2 \end{array} \right\}, \tag{17}$$

where μ_0 denotes the sample mean. The first inequality restricts the deviation of E[W] from μ_0 to at most γ_1, while the second bounds the second central moment around μ_0 by γ_2. The parameters γ_1 and γ_2 are derived from the confidence intervals of the mean and variance, and the overall uncertainty set is defined as $Z = Q \cup Y$.

3.4 Estimating the Optimal Cost using Machine Learning

As outlined in subsection 3.3, we aim to find a functional approximation to estimate the optimal cost: $\tilde{g}(\mathbf{z}, \mathbf{v}) \approx \min_\pi g(\mathbf{z}, \mathbf{v}, \pi)$. The functional form has an impact on the optimization problem. We therefore experiment with different models starting from linear regression on individual parameters to non-linear models and to engineered features.

To generate the data set for learning, we numerically solve the MDP problem for given input parameters and compute the corresponding optimal cost. The data set that was generated comprises 15000 observations of different sets of input parameters and the corresponding optimal cost. We split the generated data into training and test sets. We then estimate the models using the training set and use regularization with cross-validation to prevent overfitting. Finally, we report and compare results on the test set.

Table 4 presents the performance measures of four different classes of models: the first order regression model (LR) on the basic feature vector $(\mathbf{z}, \mathbf{v})$, the quadratic General Additive Model (GAM) on $(\mathbf{z}, \mathbf{v})$, the quadratic (QRE) and cubic (CRE) regression models with engineered features. We observe that all models demonstrate reasonable performance but the Root Mean Squared Error (RMSE) decreases significantly by using more complicated models.

Table 4: Performance comparison of the four regression models (R^2 : coefficient of determination, MAE: Mean Absolute Error, RMSE: Root Mean Squared Error, MAPE: Mean Absolute Percentage Error)

Model	$\mathbf{R^2}$	**MAE**	**RMSE**	**MAPE (%)**
LR	0.9798	0.4026	0.5252	2.5099
GAM	0.9849	0.3548	0.4540	2.2679
QRE	0.9993	0.0758	0.0992	0.2340
CRE	**0.9994**	**0.0681**	**0.0884**	**0.1619**

Due to its superior RMSE performance, we prefer to proceed with CRE and solve the robust optimization problem with the moment constraints using a solver. To elaborate, we train a third-order regression function with engineered features, $\tilde{g}(\mathbf{x})$, capturing the relationship between the problem parameters and the optimal value of the MDP (details are provided in Appendix B). This closed-form representation enables us to transform the original maxmin problem into a single-level tractable optimization problem.

The features for the third-order regression model are:

$$\mathbf{x} = \left(\lambda_p, \lambda_d, c_d, m, A_p, A_d, \mathrm{Var}[S_p], \mathrm{Var}[S_d], \mathrm{CV}^2_{S_p}, \mathrm{E}[S_p^2], \mathrm{E}[S_d^2], \frac{A}{m}, \frac{A_p}{m}, \frac{A_d}{m}, \frac{A_p}{A_d}, P_b\right)$$

where S_p and S_d are random variables that represent the parking times of passenger and delivery vehicles. The coefficient of variation of S_p is denoted as CV_{S_p}. $A_p = \lambda_p \mathrm{E}[S_p]$ and $A_d = \lambda_d \mathrm{E}[S_d]$ are the offered loads for the passenger and delivery vehicles, and the total offered load is A. P_b is the blocking probability under the policy that accepts both types of vehicles whenever the parking

lot is available and is identical for both types of vehicles. This blocking probability is given by the Erlang B formula

$$P_b = \frac{\frac{(A_p+A_d)^m}{m!}}{\sum_{k=0}^{m} \frac{(A_p+A_d)^k}{k!}}.$$

Note that only the rejection cost for the delivery vehicles, c_d, is considered, as the rejection cost for the passenger cars is fixed at 1, while c_d varies. The reason is that the optimal admission policy depends on their ratio rather than on their individual values. In addition, the coefficient of variation of S_d is excluded, as it is constant for the exponential distribution.

3.5 On the effects of the accuracy of the ML estimators for the optimization problem

To gain some insight into how ML estimation errors reflect onto robust optimization errors, let us assume that

$$\tilde{g}(\mathbf{z}, \mathbf{v}) = g(\mathbf{z}, \mathbf{v}) + \epsilon$$

where ϵ is a normally distributed error term with mean zero and variance σ_ϵ^2.

Let us now consider the worst-case parameters for the MDP problem $(\mathbf{z}^*, \mathbf{v}^*)$ with corresponding cost $g(\mathbf{z}^*, \mathbf{v}^*)$ and the worst case parameters for the estimated problem $(\mathbf{z}', \mathbf{v}')$ with the corresponding true (MDP) cost $g(\mathbf{z}', \mathbf{v}')$. Since $(\mathbf{z}^*, \mathbf{v}^*)$ are optimal we must have: $g(\mathbf{z}^*, \mathbf{v}^*) \geq g(\mathbf{z}', \mathbf{v}')$ at the same time we must have for some realizations of the error terms: $\hat{g}(\mathbf{z}^*, \mathbf{v}^*) \leq \hat{g}(\mathbf{z}', \mathbf{v}')$.

Let us define the following differences: $\Delta_{T,T} = g(\mathbf{z}^*, \mathbf{v}^*) - g(\mathbf{z}', \mathbf{v}')$, $\Delta_{T,ML} = g(\mathbf{z}^*, \mathbf{v}^*) - \hat{g}(\mathbf{z}', \mathbf{v}')$, and $\Delta_{ML,ML} = \hat{g}(\mathbf{z}^*, \mathbf{v}^*) - \hat{g}(\mathbf{z}', \mathbf{v}')$.

Proposition 4. *We have the following bound between the true robust solution and the ML-approximated robust solution $\Delta_{T,ML}$:*

For $u \geq \Delta_{T,T}$,

$$\mathbb{P}\big(|\Delta_{T,ML}| > u \mid \Delta_{ML,ML} < 0\big) \;\leq\; \frac{2\exp\Big(-\frac{(u-\Delta_{T,T})^2}{2\sigma_\epsilon^2}\Big)}{1-\Phi\Big(\frac{\Delta_{T,T}}{\sqrt{2}\,\sigma_\epsilon}\Big)}$$

where $\Phi()$ is the cdf of the standard normal random variable. The proof of the proposition can be found in Appendix C. By Proposition 4, we observe that the optimization error roughly decays exponentially with the variance of the estimation error σ_ϵ . If the error is estimated by the RMSE from the training data, we observe from Table 4 that there is a big gain going from the simplest model (LR) to the best performing one (CRE).

4 Numerical Experiments

In this section, we evaluate the performance of our data-driven and robust optimization framework through a series of numerical experiments using real-world parking data from the Sariyer case study (Section 1.2). Our analysis is structured around three key performance indicators (KPIs), each of which quantifies a distinct dimension of value added by our modeling framework.

KPI 1: Relative Value of Dynamic Admission Control. We measure the performance gains of the proposed dynamic policy by comparing it to the current practice of unrestricted access—where all vehicles are admitted whenever space is available—and to hard and soft static allocation policies that reserve a fixed number of spaces for delivery vehicles. For comprehensive analysis, we also report blocking probabilities, although they are not used as primary KPIs.

KPI 2: Relative Value of Using a Non-Exponential Model. We compare the data-driven MDP, which incorporates empirically estimated Coxian parking-time distributions, against a traditional exponential MDP. This experiment quantifies the benefit of using a more realistic, non-exponential representation of parking durations.

KPI 3: Relative Value of Robustification. Finally, we evaluate the robust model relative to the nominal model to assess the performance improvements achievable through robustness against parameter uncertainty.

Table 5 summarizes the problem parameters we use in our experiments. The underlying parking area includes three street sections with 10, 15, and 15 parking spaces; however, we also consider the cases with five parking spaces to observe the impact of the parking lot size on the performance of the solutions. The ranges of the distributional parameters are derived from their uncertainty sets, estimated from data. For the comparisons, the experiments are carried out on 100 problem instances with randomly generated distributional parameters from their uncertainty set (defined in Equations (16) and (17)). The parameters of the two-phase Coxian parking time distribution can take any values within the moment uncertainty set defined in Equation (17), bounded by γ_1 and γ_2, as reported in Table 5. As shown in Table 5, the arrival rates are specified such that the offered loads $\rho_p = \lambda_p \mathrm{E}[S_p]$ and $\rho_d = \lambda_d \mathrm{E}[S_d]$ scale proportionally with the size of the parking lot.

Table 5: Problem Parameters

Parameter	Value	Parameter	Value	Parameter	Value
λ_p	$[0.30m, 1.46m]$	λ_d	$[0.01m, 0.13m]$	γ_1	0.19
γ_2	2.436	μ_d	[1.57, 2.62]	m	{5, 10, 15}
c_d/c_p	15.41				

4.1 Comparison of Dynamic and All-Acceptance Policies

In this part, we compare the proposed dynamic optimal policy, π^*, with the all-acceptance policy (status quo), π^a, under which both vehicle types are admitted whenever a parking space is available and no spaces are reserved for delivery vehicles.

Table 6 displays the relative mean percentage differences in total cost for varying numbers of managed parking spaces. The results show that as the number of parking spaces increases, the dynamic admission control policy achieves increasingly higher cost savings over the all-acceptance policy.

Table 6: Relative mean cost differences between the optimal and all-acceptance policies for m = 5, 10, 15, including the base-case scenario

Parking Spaces	Base Case (%)	Mean (%)
5	36.49	34.21
10	41.16	38.36
15	43.19	40.31

Table 7 presents the mean differences of the blocking probabilities for both vehicle types under policies π^* and π^a. In all instances, Proposition 1 holds, implying that the optimal policy always admits delivery vehicles and restricts passenger cars when necessary. Hence, the blocking probabilities satisfy $p_b^d(\pi^*) \leq p_b(\pi^a)$ and $p_b^p(\pi^*) \geq p_b(\pi^a)$, as explained in the previous section. As the number of parking spaces increases, the optimal policy admits a relatively larger proportion of passenger vehicles. Consequently, substantial percentage cost savings can be achieved with only a small loss of passenger vehicle admissions when applying the dynamic optimal policy.

Table 7: Mean differences between the blocking probabilities under the optimal and all-acceptance policies for passenger and delivery vehicles, including the base-case scenario

	Passenger Vehicles		Delivery Vehicles	
Parking Spaces	Base Case (%)	Mean (%)	Base Case (%)	Mean (%)
5	11.74	10.57	26.13	29.43
10	6.61	7.10	19.93	27.45
15	4.68	5.29	16.84	25.66

4.2 Comparison of Dynamic and Static Policies

In this subsection, we evaluate the optimal dynamic policy in comparison to static allocation policies, which are the common strategy for allocating additional curbside spaces for freight vehicles (Nourinejad et al. (2014); Amer and Chow (2017); Abhishek et al. (2021)). Static policies designate specific curbside spaces for delivery vehicles, while the dynamic policy allows both types of vehicles to park in any available space.

First, we consider a hard allocation scheme that dedicates k spaces to delivery vehicles and $m-k$ to passenger cars, where k denotes the cost-minimizing number of delivery bays. Each vehicle type is permitted to park only in its designated spaces. The results are given in Tables 8 and Tables 9. The optimal policy yields significantly lower costs than the hard allocation policy, with an average difference of 17.77% across all problem instances. Under the optimal policy, a higher percentage of passenger cars are admitted to the parking lot, while slightly fewer delivery vehicles are accepted, which accounts for the observed cost differences.

Table 8: Relative mean cost differences between the optimal and hard allocation policies for m = 5, 10, and 15, including base-case scenario

Parking Spaces	Base Case (%)	Mean (%)
5	17.75	13.98
10	19.45	18.27
15	19.72	21.06

Table 9: Mean differences between the blocking probabilities under the optimal and hard allocation policies for passenger and delivery vehicles, including the base-case scenario

	Passenger Vehicles		Delivery Vehicles	
Parking Spaces	Base Case (%)	Mean (%)	Base Case (%)	Mean (%)
5	59.46	54.84	-2.67	-5.15
10	71.55	63.70	-1.92	-1.76
15	76.87	67.88	-1.62	-1.17

In the second case, we compare the optimal policy with the soft allocation scheme proposed by Abhishek et al. (2021), which allows delivery vehicles to park in common parking spaces when the delivery bays are full. The number of delivery bays are optimized as in the previous case. Tables 10 and 11 summarize the results. The optimal policy still yields considerably lower costs, but the difference is smaller than in the previous case with the hard allocation scheme. Although the soft allocation policy slightly increases the rejection rate for delivery vehicles, it allows more passenger cars to enter the system due to the flexible use of the shared parking spaces. Since passenger vehicles arrive more frequently, admitting more of them outweighs the small increase in the delivery vehicle rejections, resulting in lower overall rejection costs.

Table 10: Relative mean cost differences between the optimal and soft allocation policies for m = 5, 10, and 15, including base-case scenario

Parking Spaces	Base Case (%)	Mean (%)
5	6.45	5.38
10	10.38	7.67
15	10.02	9.04

To identify the operational regimes in which the dynamic policy yields substantial cost reductions and where its advantage is marginal, we conducted additional numerical experiments. Overall, the dynamic policy exhibits consistent improvements across a wide range of arrival rate ratios and parking capacities (Figure 3a). In particular, the highest performance occurs when λ_p/λ_d is between one and four and m exceeds eight.

When the cost ratio c_d/c_p equals 0.5 or 1, the dynamic and flexible allocation policies perform similarly. In these regimes, it is optimal to admit all arriving vehicles, so neither policy exclusively reserves parking spaces for delivery vehicles. However, when $c_d/c_p = 0.25$, the dynamic policy begins to reject some delivery vehicles, leading to observable cost improvements (Figure 3b). The most significant benefits occur when $c_d/c_p \geq 4$. As confirmed by our traffic microsimulation results, the rejection cost ratios exceed eight in practice, suggesting that these scenarios are likely to reflect

Table 11: Mean differences between the blocking probabilities under the optimal and soft allocation policies for passenger and delivery vehicles, including the base-case scenario

	Passenger Vehicles		Delivery Vehicles	
Parking Spaces	Base Case (%)	Mean (%)	Base Case (%)	Mean (%)
5	-1.50	2.79	3.23	0.07
10	4.29	1.37	-0.46	0.98
15	2.06	1.45	0.55	0.84

realistic operating conditions.

Figure 3: Heatmaps showing the relative percentage differences between optimal and flexible allocation policies (%) for varying numbers of parking spaces (x-axis) and (left) the ratio of arrival rates, and (right) rejection cost ratio levels (y-axis).

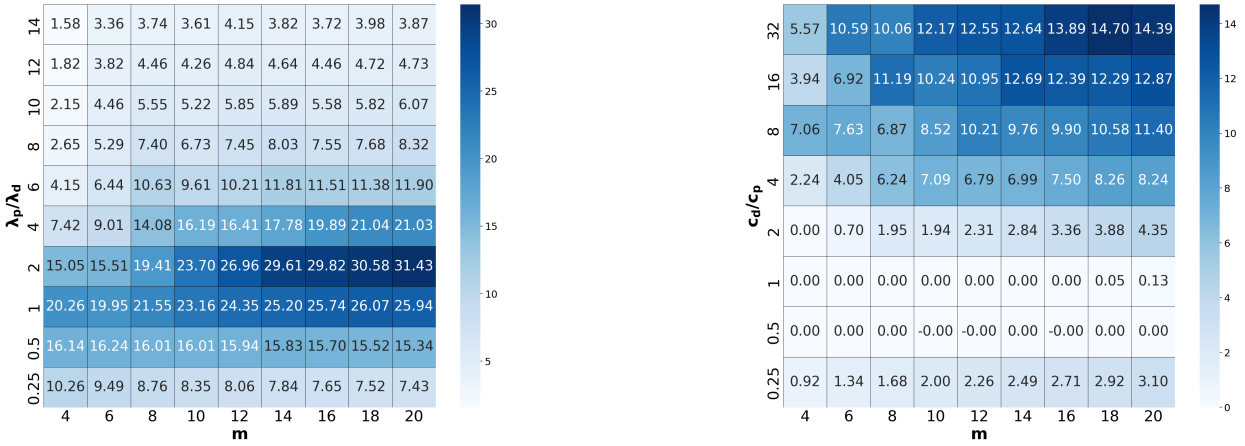

(a) Arrival rate ratio

(b) Rejection cost ratio

4.3 Comparison of Policies under Fitted Coxian and Exponential Parking Times

In this part, we compare the data-driven MDP model with fitted Coxian passenger vehicle parking times to the corresponding MDP under the assumption of exponentially distributed parking times. For each generated problem instance, the exponential distribution rate is set in a way that its mean matches that of the associated Coxian distribution, while all other problem parameters are kept identical.

Table 12: Relative mean cost differences between the data-driven MDP model and the corresponding MDP with the assumption of exponentially distributed parking times for m = 5, 10, and 15, including base-case scenario

Parking Spaces	Base Case (%)	Mean (%)
5	0.10	2.74
10	0.01	3.55
15	0.002	4.70

The comparison is made based on the average costs per time unit between the models. In the

MDP model with exponential parking times, the state is $s^e = (x_p, x_d)$, with the optimal policy π^e obtained by the value iteration algorithm. The average cost rate under π^e is then evaluated in the data-driven MDP, where $\pi^e(x_p, x_d) = \pi^*\left(x_p^1, x_p^2, x_d\right)$, for $x_p = x_p^1 + x_p^2$, and the resulting cost rate, $g(\pi^e)$ is compared to $g(\pi^*)$. Table 12 shows the percentage of the relative mean difference in average costs for 5, 10, and 15 parking spots. The overall mean difference is 3.67%, and the cost rate increases slightly with the number of parking spots.

4.4 Comparison of Robust and Nominal Policies

In this subsection, we compare the nominal model with its robust counterpart. The nominal model assumes that the parameters take their point estimate values and ignores parameter uncertainty. We additionally introduce a risk-averseness coefficient $\delta \in [0, 1]$, by which the reduced uncertainty sets Q_δ and Y_δ are defined as follows

$$Q_\delta = \left\{ q \in \mathbb{R}_+^3 : \; q_i \in \left[\tilde{q}_i - \delta(\tilde{q}_i - q_i^{\min}),\; \tilde{q}_i + \delta(q_i^{\max} - \tilde{q}_i)\right],\; i = 1, 2, 3 \right\},$$

$$Y_\delta = \left\{ y \in \mathbb{R}_+^3 \;\middle|\; \begin{array}{l} |\mathrm{E}[W] - \mu_0| \le \delta\gamma_1, \\ \mathrm{E}\left[(W - \mu_0)^2\right] \le \delta\gamma_2 \end{array} \right\},$$

where Q_δ is the uncertainty set of the exponentially distributed parameters with point estimates $\tilde{q}_i$ and Y_δ is the moment-based uncertainty set of the Coxian distributed parking time. Setting $\delta = 0$ corresponds to the nominal model, $\delta = 1$ corresponds to the fully robust model, and the intermediate values provide semi-robust policies with increasing risk-averseness.

Figure 4 presents box plots of nominal and robust costs with different risk-averseness coefficients. Robust policies achieve substantially lower costs in worst-case scenarios, while offering slightly lower median costs compared to nominal policies, with the lowest costs in the three cases observed at $\delta = 0.25$.

4.5 Main findings

The main findings from our numerical experiments are as follows.

- For KPI 1, we find that there is a significant improvement in the cost by the use of a fully dynamic policy. The relative percentage savings in cost is more than 34% with no allocation and more than 14% with hard allocation. Soft allocation closes the gap but the relative percentage cost may still be over 7% for our base scenario. We further observe that the relative percentage cost can be over 30% depending on the ratio of λ_p/λ_d. Moreover, dynamic allocation does not lead to major changes in rejection probabilities. This highlights the value of using fully dynamic policies and quantifies the potential cost savings to be weighed against the costs of implementing the system.

Figure 4: Box plots for costs under the nominal and robust policies with risk-awareness levels $\delta \in \{0.25, 0.50, 0.75, 1.00\}$ for 5, 10, and 15 parking spaces.

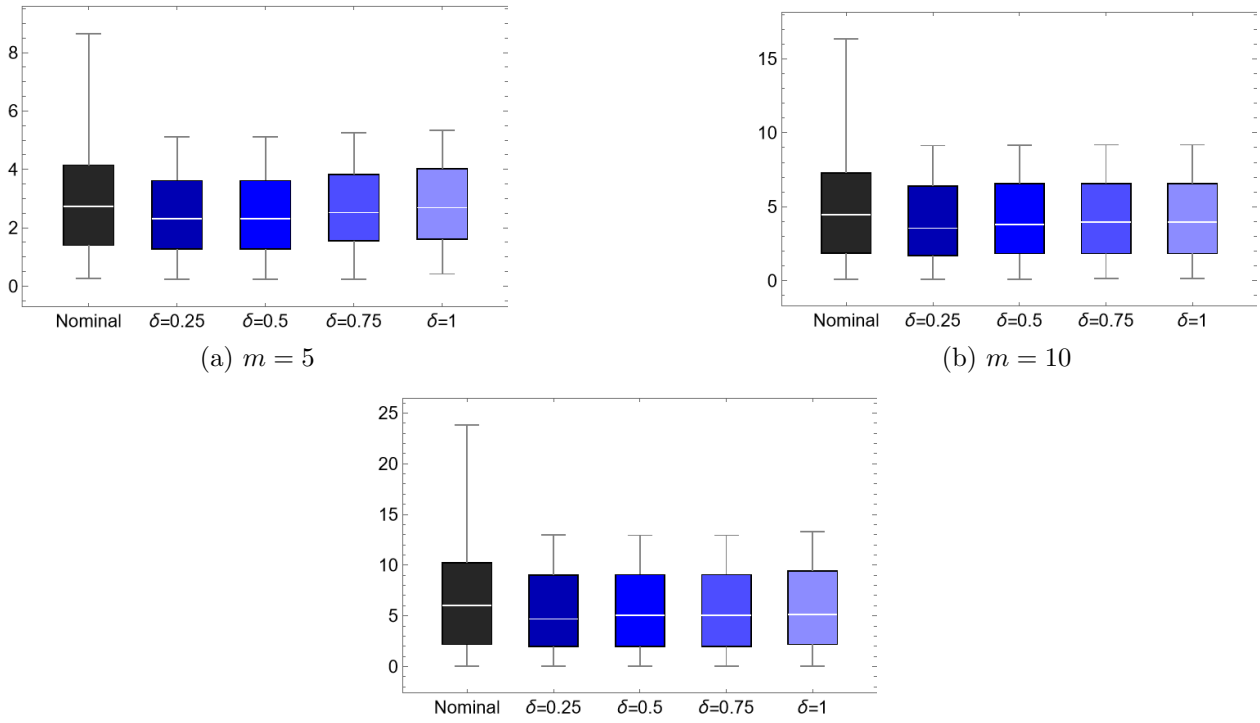

(a) $m = 5$ (b) $m = 10$

(c) $m = 15$

Note: On each box, the central line indicates the median, the upper and lower edges of the box indicate the 25% and 75% quartiles, and the whiskers show the maximum and minimum values.

- For KPI 2, we observe that there is some non-negligible value (up to 5%) in modelling non-exponential parking durations to determine the optimal dynamic allocation policy. This effect appears to be more pronounced as the number of parking spaces increases.
- For KPI 3, we observe that robustification adjustments to adapt the dynamic allocation policy for parameter estimation uncertainty have significant value for our scenario. We think that estimation uncertainty is a relevant issue, since driver preferences depend on a large number of factors, some of which are difficult to capture in a forecasting model.

5 Conclusions

This paper studied the dynamic parking space allocation problem for an on-street parking lot shared by passenger and delivery vehicles, each with distinct and stochastic parking demand behaviour. The objective was to minimize the externality costs (traffic delay and emission costs) due to double parking by delivery vehicles and cruising by passenger vehicles, caused by the scarcity of parking spaces. To this end, we developed a data-driven Markov decision process (MDP) model to construct an optimal admission control policy that decides whether to accept or reject an arriving vehicle

based on the current system dynamics. To evaluate the benefits of dynamic admission control, we benchmarked it against static allocation policies proposed in the literature, which reserve a predetermined number of parking spaces for delivery vehicles. Our numerical experiments across a wide range of data demonstrated that the dynamic policy can deliver significant performance gains.

Moreover, we proposed a novel robust optimization framework to develop robust policies against parameter uncertainties. Due to the intractability of the robust problem, we developed an effective machine-learning method. The numerical results on empirical data showed that the robust approach achieves significantly lower costs relative to the nominal model that does not take parameter uncertainty into account.

Our findings demonstrated that the proposed methodological framework can support policymakers in managing curbside parking more efficiently. By modelling parking durations with flexible Coxian distributions, the framework captures realistic parking-time behaviour. The generality of the Coxian family also makes our approach applicable to a broad set of admission-control systems beyond curbside parking where exponential assumptions are inadequate to represent service times. Combined with the machine-learning-based robustification scheme, the framework offers a reliable and adaptable decision-support tool that remains effective under uncertainty in the underlying stochastic parameters.

In this study, the parking lot is publicly managed, and the price is not a decision variable. In other contexts, pricing could be used as a control mechanism to manage parking demand and could be integrated with admission control policies. In addition, our focus in this study was on reducing negative externalities. In settings where the primary objective is profit maximization, pricing decisions could play a central role.

References

Abdeen, M. A. R., Nemer, I. A., and Sheltami, T. R. (2021). A balanced algorithm for in-city parking allocation: A case study of al madinah city. *Sensors*, 21(9):3148.

Abhishek, Legros, B., and Fransoo, J. C. (2021). Performance evaluation of stochastic systems with dedicated delivery bays and general on-street parking. *Transportation Science*, 55(5):1070–1087.

Altman, E., Jimenez, T., and Koole, G. (2001). On optimal call admission control in resource-sharing system. *IEEE Transactions on Communications*, 49(9):1659–1668.

Amer, A. and Chow, J. Y. (2017). A downtown on-street parking model with urban truck delivery behavior. *Transportation Research Part A: Policy and Practice*, 102:51–67.

Baker, L. (2019). Today's pickup: Ups hit with $33.8 million in nyc parking fines; fedex, $14.9 million. `https://www.freightwaves.com/news/todayspickup-ups-fedex-parking-fines?`

Bogoslavskyi, I., Spinello, L., Burgard, W., and Stachniss, C. (2015). Where to park? minimizing

the expected time to find a parking space. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.

Burns, A. J., Michalek, J. J., and Samaras, C. (2024). Estimating the potential for optimized curb management to reduce delivery vehicle double parking, traffic congestion and energy consumption. *Transportation Research Part E: Logistics and Transportation Review*, 187:103574.

Caliskan, M., Barthels, A., Scheuermann, B., and Mauve, M. (2007). Predicting parking lot occupancy in vehicular ad hoc networks. In *2007 IEEE 65th Vehicular Technology Conference - VTC2007-Spring*. IEEE.

City of Melbourne (2020). On-street car parking sensor data - 2018. `https://data.melbourne.vic.gov.au/explore/dataset/on-street-car-parking-sensor-data-2018/information/`. Accessed: 2025-06-20.

D'Acierno, L., Gallo, M., and Montella, B. (2006). Optimisation models for the urban parking pricing problem. *Transport Policy*, 13(1):34–48.

Dowling, C. P., Ratliff, L. J., and Zhang, B. (2020). Modeling curbside parking as a network of finite capacity queues. *IEEE Transactions on Intelligent Transportation Systems*, 21(3):1011–1022.

Humes, E. (2019). Online shopping was supposed to keep people out of traffic. it only made things worse. https://time.com/5481981/online-shopping-amazon-free-shipping-traffic-jams/.

INRIX, I. (2024). Inrix 2024 global traffic scorecard. Technical report, INRIX. Accessed 23 June 2025.

Iyengar, G. N. (2005). Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280.

James, G., Witten, D., Hastie, T., Tibshirani, R., and Taylor, J. (2023). *An Introduction to Statistical Learning: With Applications in Python*. Springer, 2nd edition.

Jodiawan, P., Côté, J.-F., and Coelho, L. C. (2025). The flexible park-and-loop routing problem. *Transportation Research Part C: Emerging Technologies*, 178:105173.

Kim, W. and Wang, X. (2021). Double parking in new york city: a comparison between commercial vehicles and passenger vehicles. *Transportation*, 49(5):1315–1337.

Le Colleter, T., Dumez, D., Lehuédé, F., and Péton, O. (2023). Small and large neighborhood search for the park-and-loop routing problem with parking selection. *European Journal of Operational Research*, 308(3):1233–1248.

Legros, B. and Fransoo, J. C. (2024). Admission and pricing optimization of on-street parking with delivery bays. *Eur. J. Oper. Res.*, 312(1):138–149.

Levin, M. W. and Boyles, S. D. (2019). Optimal guidance algorithms for parking search with reservations. *Networks and Spatial Economics*, 20(1):19–45.

Lim, J. and Masoud, N. (2024). Dynamic usage allocation and pricing for curb space operation. *Transportation Science*, 58(6):1252–1276.

Lin, T., Rivano, H., and Le Mouël, F. (2017). A survey of smart parking solutions. *IEEE Transactions on Intelligent Transportation Systems*, 18(12):3229–3253.

Mannor, S., Simester, D., Sun, P., and Tsitsiklis, J. N. (2007). Bias and variance approximation in value function estimates. *Management Science*, 53(2):308–322.

Nilim, A. and El Ghaoui, L. (2005). Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798.

Nourinejad, M., Wenneman, A., Habib, K. N., and Roorda, M. J. (2014). Truck parking in urban areas: Application of choice modelling within traffic microsimulation. *Transportation Research Part A: Policy and Practice*, 64:54–64.

Ormeci, E. L., Burnetas, A., and van der Wal, J. (2001). Admission policies for a two class loss system. *Stochastic Models*, 17(4):513–539.

Roca-Riu, M., Fernández, E., and Estrada, M. (2015). Parking slot assignment for urban distribution: Models and formulations. *Omega*, 57:157–175.

Wiesemann, W., Kuhn, D., and Rustem, B. (2013). Robust markov decision processes. *Mathematics of Operations Research*, 38(1):153–183.

World Economic Forum (2024). Transforming urban logisics: Sustainable and efficient last-mile delivery in cities. `https://reports.weforum.org/docs/WEF_Transforming_Urban_Logistics_2024.pdf`. White Paper.

Xiao, J., Lou, Y., and Frisby, J. (2018). How likely am i to find parking? – a practical model-based framework for predicting parking availability. *Transportation Research Part B: Methodological*, 112:19–39.

Zhang, X., Pitera, K., and Wang, Y. (2023). Parking reservation techniques: A review of research topics, considerations, and optimization methods. *Journal of Traffic and Transportation Engineering (English Edition)*, 10(6):1099–1117.

Appendix A Description of the Two-Phase Coxian Distribution

The two-phase Coxian distribution describes the time to absorption of a sequential two-stage process. The phases are independent and exponentially distributed with rates μ_1 and μ_2. After

completion of the first phase, the process proceeds to the next stage with probability p. Let T denote the time until absorption, then the cumulative distribution function is given by

$$F(t) = \frac{\mu_1 - \mu_2 + e^{-\mu_1 t}(\mu_2 + \mu_1(-1+p)) - e^{-\mu_2 t}\mu_1 p}{\mu_1 - \mu_2},$$

where $t > 0$. The first and second moments and the coefficient of variation are as follows:

$$\mathrm{E}[T] = \frac{1}{\mu_1} + \frac{p}{\mu_2}, \quad \mathrm{E}[T^2] = \frac{2\left(\mu_2^2 + \mu_1(\mu_1+\mu_2)p\right)}{\mu_1^2\mu_2^2}$$

$$\mathrm{Var}[T] = \frac{1}{(\mu_1)^2} + \frac{p(2-p)}{(\mu_2)^2}, \quad \mathrm{CV} = \sqrt{\frac{\mathrm{Var}[T]}{(\mathrm{E}[T])^2}}.$$

The exponential distribution (when $p = 0$) and the Erlang distribution (when $p = 1$ and $\mu_1 = \mu_2$) are special cases of the two-phase Coxian distribution. Depending on the choice of parameters, the coefficient of variation can be less than, equal to, or greater than 1, providing flexibility to fit a wide range of data.

Appendix B Machine Learning Model

In this section, we present the third-order regression model obtained after eliminating irrelevant features using lasso regression with cross-validation. The model estimates the minimum long-run expected cost rate $g\left(\pi^*, \mathbf{z}, \mathbf{v}\right)$, defined as:

$$\tilde{g}(\mathbf{x}) = \beta_0 + \sum_{1\le i\le n} \beta_i x_i + \sum_{1\le i\le j\le n} \beta_{ij} x_i x_j + \sum_{1\le i\le j\le k\le n} \beta_{ijk} x_i x_j x_k$$

We note that there are $\begin{pmatrix} n+3 \\ n \end{pmatrix}$ distinct terms in the cubic model. We have $n = 16$ which results in 969 terms. After cross-validation and model reduction, many of these terms are eliminated. The reduced model has 37 terms. The intercept and coefficients for the final reduced model, $\tilde{g}^r\left(\mathbf{x}\right)$, are given as follows:

Feature i	$\tilde{g}_i^r(\mathbf{x})$	β_i	Feature i	$\tilde{g}_i^r(\mathbf{x})$	β_i	Feature i	$\tilde{g}_i^r(\mathbf{x})$	β_i
0	1	-0.140	13	$\mathrm{Var}[S_p]\mathrm{Var}[S_d]$	-0.419	26	$\frac{c_d^2 A_p}{A_d}$	-0.0000031
1	λ_p	0.025	14	$\mathrm{Var}[S_d]\,\mathrm{E}[S_p^2]$	-0.027	27	$c_d\,\mathrm{CV}_{S_p}^2 P_b$	0.00125
2	c_d	0.006	15	$\mathrm{CV}_{S_p}^2 P_b$	0.697	28	m^3	-0.000081
3	P_b	0.283	16	$\mathrm{E}[S_p^2]\mathrm{E}[S_d^2]$	-0.055	29	$\frac{\mathrm{Var}[S_p]^2 A}{m}$	-0.288
4	$\lambda_p c_d$	0.002	17	$\frac{A_d}{m} P_b$	1.992	30	$\mathrm{Var}[S_p]\mathrm{E}[S_p^2]\mathrm{Var}[S_d]$	0.0129
5	$\frac{\lambda_p A}{m}$	0.041	18	$\frac{\lambda_p^2 A_p}{A_d}$	2.7×10^{-6}	31	$\mathrm{Var}[S_p]\mathrm{E}[S_p^2]\mathrm{E}[S_d^2]$	0.0258
6	$\frac{\lambda_p A_d}{m}$	0.230	19	$\lambda_p c_d\,\mathrm{CV}_{S_p}^2$	0.000074	32	$\frac{\mathrm{Var}[S_p]\mathrm{E}[S_p^2]A}{m}$	0.188
7	$\lambda_p P_b$	0.724	20	$\lambda_d \mathrm{CV}_{S_p}^2 P_b$	0.265	33	$\mathrm{E}[S_p^2]^2\mathrm{Var}[S_d]$	0.0557
8	$\lambda_d P_b$	0.553	21	$c_d^2 A_p$	-0.000018	34	$\frac{\mathrm{CV}_{S_p}^2 A_d}{m} P_b$	1.707
9	$c_d\,\mathrm{CV}_{S_p}^2$	0.004	22	$c_d^2 A_d$	-0.000255	35	$\frac{\mathrm{E}[S_p^2]^2 A}{m}$	-0.0217
10	$\frac{c_d A_d}{m}$	0.220	23	$c_d^2\mathrm{Var}[S_d]$	-0.000057	36	$\frac{\mathrm{E}[S_p^2]^2 A_p}{m}$	-0.0207
35	$\frac{\mathrm{E}[S_p^2]^2 A}{m}$	-0.0217	36	$\frac{\mathrm{E}[S_p^2]^2 A_p}{m}$	-0.0207	37	$\frac{\mathrm{E}[S_p^2]^2 A_p}{A_d}$	0.000010
12	mP_b	0.111	25	$\frac{c_d^2 A}{m}$	-0.000598			

Appendix C Proof of Proposition 4

Proof. We need to compute a conditional probability $P(A \mid B)$ where A corresponds to $\{|\Delta_{T,ML}| > u\}$ and B corresponds to $\{\Delta_{ML,ML} < 0\}$.

$$\mathbb{P}(A \mid B) = \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \leq \frac{\mathbb{P}(A)}{\mathbb{P}(B)}.$$

Then,

$$\mathbb{P}\big(|\Delta_{T,ML}| > u \mid \Delta_{ML,ML} > 0\big) \leq \frac{\mathbb{P}\big(|\Delta_{T,ML}| > u\big)}{\mathbb{P}(\Delta_{ML,ML} > 0)}.$$

For the numerator, $\mathbb{P}(B)$, we have

$$\Delta_{T,ML} = \Delta_{T,T} - \varepsilon_2.$$

For $u \geq \Delta_{T,T}$,

$$|\Delta_{T,ML}| = |\Delta_{T,T} - \varepsilon_2| \leq |\Delta_{T,T}| + |\varepsilon_2|,$$

so

$$\{|\Delta_{T,ML}| > u\} \subseteq \{|\varepsilon_2| > u - \Delta_{T,T}\}.$$

Thus, for $u \geq \Delta_{T,T}$,

$$\mathbb{P}(|\Delta_{T,ML}| > u) \leq \mathbb{P}\big(|\varepsilon_2| > u - \Delta_{T,T}\big) = 2\Big(1 - \Phi\Big(\frac{u - \Delta_{T,T}}{\sigma}\Big)\Big).$$

Using the standard Gaussian tail inequality

$$1 - \Phi(x) \leq e^{-x^2/2} \quad \text{for } x \geq 0,$$

and putting the numerator and denominator together, we obtain the result. □ □

Let us assume that $\Delta_{T,T} = k\sigma$ (i.e. within some confidence region), then we can get a sharper bound on $\Delta_{T,ML}$ being far away from $\Delta_{T,T}$ (i.e. the difference due to estimation error cannot be too different from Δ_{TT}.)

Corollary 2. *If $\Delta_{T,T} = k\sigma$ and we take $u = k\sigma + \Delta_u$ then*

$$\mathbb{P}\big(|\Delta_{T,ML}| > u \mid \Delta_{ML,ML} < 0\big) \;\leq\; \frac{2\exp\Big(-\frac{\Delta_u^2}{2\sigma^2}\Big)}{1 - \Phi\Big(\frac{k}{\sqrt{2}}\Big)}$$

Proof. This directly follows from Proposition 4. □ □

Online Appendix

OA1 Traffic Microsimulation

In what follows, we provide the operational details of the microsimulation and emission modeling framework employed in this study. All simulations were conducted in PTV VISSIM [Vissim(2022)], replicating the physical and control characteristics of the study corridor with high accuracy. The geometric layout, lane configurations, conflict areas, and priority rules were obtained from field observations and the planning documents provided by the local transportation authority, and were then directly transferred into the simulation environment. Two signalized intersections within the corridor were modeled using their exact field-controller settings, including phase sequences, intergreen periods, and cycle lengths, ensuring that the simulated control logic reflects the operational conditions observed in practice.

Time-dependent background traffic demand was constructed using the turning movement counts and temporal volume distributions supplied by the local authority. These empirical inputs were used to define vehicle entry volumes and routing percentages, allowing the simulation to capture the temporal variability inherent in urban traffic conditions.

Driver behavior was represented using the Wiedemann 74 psycho-physical car-following logic [Wiedemann(1974)]. The model distinguishes between free-driving, approaching, following, and braking regimes through threshold-based perception processes, and is particularly well-suited for dense and heterogeneous urban conditions. For completeness, we note that Wiedemann 99 [Vissim(2022)] introduces smoother spacing adaptation and extended perception thresholds tailored to higher-speed environments; however, given the urban characteristics of the study corridor, Wiedemann 74 provides a more realistic representation of local driving behavior. To account for temporal variability in demand and system loading, three simulation horizons were evaluated: (i) 2 hours, (ii) 2 hours 30 minutes, and (iii) 3 hours 30 minutes. Each scenario included a 30-minute warm-up period to allow initial transients, queue formation, and signal synchronization to stabilize before data

collection. All simulations were repeated using random seeds to mitigate stochastic variation and ensure statistically robust indicators.

Fuel consumption and emissions were computed by integrating the MOVESTAR model [Wang et al.(2020)Wang, Wu, Scora] directly into VISSIM through its C++ API. MOVESTAR provides an open-source framework for estimating second-by-second emissions based on instantaneous vehicle-specific power (VSP) and operating-mode classifications. During the simulation, the MOVESTAR module processes the instantaneous speed and acceleration data generated by VISSIM in real-time, enabling the synchronized and consistent estimation of fuel consumption and emissions across all scenarios.

OA2 Parking Data

Table OA1: Traffic inflow data of a weekday

Time Interval	Passenger Car	Heavy Vehicle	Bus/Minibus/Van	Total
10:00–11:00	672	192	96	960
11:00–12:00	742	149	99	990
12:00–13:00	1085	310	155	1550
13:00–14:00	980	280	140	1400
14:00–15:00	1169	387	114	1670
15:00–16:00	1272	335	107	1714
16:00–17:00	1336	367	93	1796
17:00–18:00	1349	341	126	1816

Table OA2: Observed parking times (hours) for two weekdays and a weekend day

Weekdays							
7.5	4.7	3.917	3.083	2.95	2.4	2.25	2.25
1.283	1.267	1.267	1.033	0.867	0.633	0.583	0.567
0.467	0.433	0.35	0.283	0.233	0.2	0.183	0.167
0.167	0.133	0.133	0.067	0.05	6.25	2.35	2.333
1.45	1.4	1.3	1.25	1.25	1.217	1.183	1.183
1.167	1.017	1.017	1	0.983	0.9	0.883	0.833
0.767	0.683	0.567	0.517	0.4	0.383	0.383	0.267
0.217	0.2	0.067	4.45	2.467	1.95	1.6	1.467
1.467	1.45	1.417	1.167	1	1	0.983	0.617
0.55	0.55	0.55	0.533	0.5	0.5	0.5	0.5
0.5	0.5	0.5	0.483	0.483	0.45	0.433	0.417
0.283	0.017	7.033	4.983	4.05	4.05	2.983	2.317
2.083	1.833	1.567	1.533	1.5	1.483	1.45	1.2
1.017	0.983	0.983	0.75	0.733	0.483	0.467	0.433
0.283	0.25	0.25	0.25	0.233	0.217	0.15	0.083
6.667	5.633	5.167	3.5	3.5	3.5	3.5	3.5
3.5	3.5	2.717	2.617	2.5	2.5	2.5	2.5
2.5	2.167	2.15	1.667	1.667	1.667	1.667	1.5
1.5	1.5	1	1	1	1	1	1
1	1	0.75	0.75	0.633	0.633	0.5	0.5
0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
0.367	0.367	0.333	0.217	0.183	0.167	0.017	5.983
3.633	3.117	1.867	1.467	1.2	1.2	1.067	0.983
0.983	0.967	0.967	0.933	0.883	0.65	0.633	0.55
0.517	0.517	0.5	0.5	0.5	0.5	0.483	0.483
0.467	0.467	0.45	0.433	0.433			
Weekend							
4.067	3.483	2.6	1.517	1.033	1	1	1
1	1	0.983	0.983	0.967	0.533	0.517	0.517
0.517	0.517	0.483	0.467	0.467	0.467	0.467	0.467
3.033	2.467	2.1	2.083	1.883	1.883	1.883	1.383
0.583	0.467	0.383	3.85	3.85	3.85	3.85	3.467
3.467	3.1	2.967	2.55	2.55	1.967	1.967	1.767
1.317	1.3	0.983	0.967	0.75	0.383	0.167	2.117
2.067	1.767	1.517	1.517	1.5	1.5	1.017	1
0.617	5.033	3.683	3.217	3.183	2.2	1.967	1.633
1.267	1.233	1.233	1	1	0.983	0.983	0.833
0.75	0.717	0.7	0.6	0.533	0.5	0.483	0.4
0.317	0.317	0.25	0.2	0.15	0.017	2.517	2.35
2.017	2	1.867	1.567	1.533	1.417	0.833	0.367
0.35	0.35						

References

[Vissim(2022)] Vissim P (2022) Ptv vissim. *Karlsruhe, Germany* .

[Wang et al.(2020)Wang, Wu, Scora] Wang Z, Wu G, Scora G (2020) Movestar: An open-source vehicle fuel and emission model based on usepa moves. *arXiv preprint arXiv:2008.04986* .

[Wiedemann(1974)] Wiedemann R (1974) *Simulation des Straßenverkehrsflusses*. Ph.D. thesis, Institut für Verkehrswesen, Universität Karlsruhe, Karlsruhe, Germany.