Titre: Title:	Optimization Models for Parking Enforcement Planning: Integrating Driver Behavior, Resource Allocation, and Routing Strategies
Auteur: Author:	Mohsen Yahyaei
Date:	2025
Type:	Mémoire ou thèse / Dissertation or Thesis
Référence: Citation:	Yahyaei, M. (2025). Optimization Models for Parking Enforcement Planning: Integrating Driver Behavior, Resource Allocation, and Routing Strategies [Thèse de doctorat, Polytechnique Montréal]. PolyPublie. https://publications.polymtl.ca/71314/

URL de PolyPublie: PolyPublie URL:	https://publications.polymtl.ca/71314/
Directeurs de recherche: Advisors:	Michel Gendreau, & Nikolaj Van Omme
Programme: Program:	Doctorat en génie industriel

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Optimization Models for Parking Enforcement Planning: Integrating Driver Behavior, Resource Allocation, and Routing Strategies

MOHSEN YAHYAEI
Département de mathématiques et de génie industriel

Thèse présentée en vue de l'obtention du diplôme de *Philosophiæ Doctor*
Génie industriel

Décembre 2025

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Cette thèse intitulée :

Optimization Models for Parking Enforcement Planning: Integrating Driver Behavior, Resource Allocation, and Routing Strategies

présentée par **Mohsen YAHYAEI**
en vue de l'obtention du diplôme de *Philosophiæ Doctor*
a été dûment acceptée par le jury d'examen constitué de :

Louis-Martin ROUSSEAU, président
Michel GENDREAU, membre et directeur de recherche
Nikolaj VAN OMME, membre et codirecteur de recherche
Nadia LAHRICHI, membre
Andrés L. MEDAGLIA, membre externe

DEDICATION

To Roya,
For your patience, your perseverance
and your tireless support.

Through every difficulty, every uncertainty
and every endless night, you believed in me
more than I believed in myself.

This thesis is not only the result
of my work, but also of your strength,
your sacrifices, and your unwavering
faith in our shared path.

Everything written here
belongs to you as much as it does to me.

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my supervisors, Professor Michel Gendreau and Dr. Nikolaj Van Omme. Their guidance, support, and high standards have shaped every stage of this work.

Michel, your professionalism, guidance, and unwavering support, especially during a period of transition when you worked tirelessly to secure appropriate projects for your students, including me, shaped the course of this dissertation. Nikolaj, your openness, availability, and willingness to engage in discussion at any moment were invaluable to the progress of this dissertation. I am sincerely grateful for the trust you both placed in me.

I also acknowledge Tony (Muhammad) Yousefnezhad, a colleague and friend, whose support during the most challenging periods of this PhD was both meaningful and essential. Your presence and encouragement helped me remain steady when the process felt overwhelming.

My heartfelt thanks go to my family, whose support has been a constant source of strength, comfort, and love throughout this journey. To Roya, whose patience, diligence, and unwavering belief in this path have carried me further than I could have ever gone alone, this achievement is deeply shared with you.

I am also grateful to my friends, both within and outside the university, who provided encouragement, perspective, and balance along the way. There are many names that could be listed here, and I hope each of you recognizes the impact of your presence in my life during these years.

RÉSUMÉ

L'application du stationnement constitue un levier essentiel pour assurer la rotation des véhicules, l'accessibilité et une utilisation équitable de l'espace public. Toutefois, le comportement de stationnement illégal est adaptatif : les conducteurs ajustent leurs décisions en fonction de la probabilité perçue de recevoir une contravention. Cette thèse développe un cadre d'optimisation intégré qui relie explicitement les opérations d'application des règles au comportement des conducteurs, et ce, sur plusieurs horizons de planification.

La première contribution propose un modèle de planification opérationnelle dans lequel les décisions de routage des patrouilles influencent, et sont influencées par, la conformité. Un modèle de parcours périodique sensible au comportement est formulé, où les revenus proviennent conjointement des paiements de stationnement légal et des contraventions émises aux contrevenants, ces deux composantes étant modulées par l'intensité perçue de l'application. L'apport central réside dans une traduction d'échelle entre les décisions discrètes de patrouille et une intensité d'application continue, permettant l'intégration d'un équilibre de rencontre (meeting-rate) dans les modèles de décision opérationnels et tactiques. En raison de ce couplage comportemental et de la complexité combinatoire, le modèle est résolu à l'aide d'une matheuristique de type *Kernel Search* identifiant des itinéraires de patrouille de haute qualité. Une étude de cas réelle portant sur 32 installations de stationnement et 5 agents à Montréal illustre l'applicabilité du modèle dans un contexte municipal.

La deuxième contribution étend la planification opérationnelle à des situations où plusieurs inspections sont nécessaires. Nous formulons un problème d'orienteering en équipe avec visites multiples et contraintes de temps de récupération minimal, et développons une heuristique *Variable Neighborhood Descent* incluant un mécanisme de réparation de faisabilité de manière adaptative. Les résultats montrent que le moment choisi pour revisiter un site influence fortement la conformité et les revenus, et que la version adaptative de l'algorithme améliore la qualité des solutions par rapport à la version classique, tout en n'entraînant qu'un surcoût computationnel marginal.

La troisième contribution aborde l'allocation tactique des ressources entre différentes zones urbaines. En s'appuyant sur un modèle d'équilibre comportemental, un seuil critique d'effectif est obtenu sous forme fermée à l'aide de la fonction de Lambert–W, identifiant le niveau d'application à partir duquel les conducteurs basculent de comportements illégaux vers des comportements conformes. Ce mécanisme de dissuasion est intégré dans un modèle de programmation dynamique qui répartit les agents tout en garantissant une probabilité minimale uniforme de détection afin d'assurer l'équité territoriale. Les expériences numériques et une étude de cas à l'échelle d'un district montréalais montrent qu'un déploiement modéré et mieux réparti peut offrir une conformité plus élevée et des résultats plus équitables qu'une application fortement concentrée.

Dans l'ensemble, cette thèse démontre que l'efficacité de l'application du stationnement dépend non seulement de l'intensité déployée, mais également de sa *distribution spatiale*, de sa *coordination temporelle* et de sa *perception comportementale*. Les modèles et algorithmes proposés fournissent des outils concrets pour soutenir des stratégies d'application à la fois efficientes, équitables et crédibles du point de vue des usagers.

ABSTRACT

Parking enforcement is a critical mechanism for ensuring curbside turnover, accessibility, and the equitable use of public space. However, illegal parking behavior is adaptive: drivers adjust their decisions in response to the perceived probability of receiving a citation. This thesis develops an integrated optimization framework that explicitly links enforcement operations with driver behavioral response across multiple planning horizons.

The first contribution formulates an operational planning model in which patrol routing decisions influence, and are influenced by, compliance. A behavior-aware periodic routing model is developed, where revenue arises jointly from legal parking payments and citations issued to illegal parkers, with both components mediated by perceived enforcement intensity. The central insight is a translation of scales between discrete patrol decisions and continuous enforcement intensity, enabling the incorporation of a meeting-rate equilibrium into operational and tactical decision models. Due to the behavioral coupling and combinatorial complexity, the model is solved using a Kernel Search matheuristic that identifies high-quality patrol patterns. A real case study with 32 parking facilities and 5 enforcement officers in Montreal illustrates the applicability of the behavior-aware routing model to practical planning contexts.

The second contribution extends operational planning to settings where multiple inspections are required. We formulate a Multi-Visit Team Orienteering Problem with minimum recovery-time constraints and develop a Variable Neighborhood Descent heuristic with feasibility repair in adaptive fashion. Results show that strategically timed revisits significantly increase compliance and revenue, with the adaptive search improving solution quality compared to a classical VND at only marginal additional computational cost.

The third contribution addresses tactical resource allocation across multiple urban regions. Presenting a new behavioral equilibrium foundation, a closed-form critical staffing threshold is derived via the Lambert–W function, characterizing the point at which enforcement becomes sufficiently credible to shift drivers from illegal to legal parking. This deterrence mechanism is embedded in a dynamic programming model that allocates officers across regions while ensuring a uniform minimum average detection probability as an equity requirement. Computational experiments and a district-scale Montreal case study illustrate how moderate and well-distributed enforcement can achieve higher compliance and fairer outcomes than high-intensity deployment in limited zones.

Overall, this thesis demonstrates that effective parking enforcement depends not only on how

much enforcement is applied, but on how it is spatially distributed, temporally coordinated, and behaviorally perceived. The proposed models and algorithms provide actionable planning tools for municipalities seeking enforcement strategies that are efficient, equitable, and behaviorally credible.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS AND ACRONYMS

AI	Artificial Intelligence
API	Application Programming Interface
ALPR	Automatic License Plate Recognition
CPP	Chinese Postman Problem
CV	Commercial Vehicle
IoT	Internet of Things
LSTM	Long Short-Term Memory
MTZ	Miller–Tucker–Zemlin
MILP	Mixed-Integer Linear Programming
MIP	Mixed-Integer Programming
NP-hard	Nondeterministic Polynomial-time hard
RMIP	Relaxed Mixed-Integer Problem
SB	Symmetry-Breaking
DP	Dynamic Programming
ECV	Expected Cost of Violation
KS	Kernel Search
KUB	Knapsack-Based Upper Bound
LP	Linear Programming
MVRP	Multi-Visit Routing Problem
MVTOP	Multi-Visit Team Orienteering Problem
PVRP	Periodic Vehicle Routing Problem
PMR	Plateau–Mont-Royal
RRE	Rerouting and Enhancement
TOP	Traveling Officer Problem
VND	Variable Neighborhood Descent
VNS	Variable Neighborhood Search
VND+S	VND with Shaking

CHAPTER 1 INTRODUCTION

Urban parking management plays a central role in shaping mobility conditions, curbside access, commercial vitality, and the equitable use of public space. Parking enforcement, in particular, ensures adherence to regulations that govern turnover, prioritize essential activities, and mitigate externalities such as congestion, blocked access points, and unsafe circulation around high-demand street segments.

However, designing effective enforcement strategies is challenging. Illegal parking behavior does not remain the same over time: drivers adjust their decisions according to the perceived likelihood of receiving a citation. When enforcement presence appears inconsistent or weak, the expected cost of violation decreases, encouraging non-compliance. Conversely, credible and appropriately distributed enforcement promotes legal parking and sustainable curb usage. This behavioral feedback mechanism implies that enforcement planning must account for the strategic responses of drivers rather than be treated purely as a logistical routing problem.

At the same time, enforcement agencies operate under practical limitations. The number of available officers is fixed in the short run, their patrol time is finite, and the geography of parking demand is heterogeneous. Some neighborhoods experience high turnover and repeated violations throughout the day, whereas others may require only occasional monitoring. Thus, enforcement planning involves decisions across multiple temporal and spatial scales, including:

- Operational planning of daily patrol routes and inspection patterns.
- Tactical planning of how enforcement resources are allocated across different regions over longer horizons.

These two layers are interconnected. Operational routing shapes the effective enforcement intensity experienced by drivers, while tactical allocation influences the baseline level of deterrence that different neighborhoods receive. Planning that ignores either dimension risks inefficiencies, inequities, or unintended behavioral consequences.

1.1 Research Objectives

The overall goal of this thesis is to develop a coherent planning framework for parking enforcement that integrates behavioral responses, spatial heterogeneity, and operational feasibility. To achieve this, the research is structured around three high-level objectives:

1. **To model the interaction between enforcement presence and driver compliance.** This involves capturing how the perceived risk of citation shapes parking decisions, and how enforcement activity influences that perception. The aim is to represent the feedback loop between enforcement and behavior in a tractable yet realistic manner.

2. **To design operational planning methods for enforcement routing under resource constraints.** This includes developing optimization models and solution procedures that generate feasible patrol patterns while respecting limits on officer availability, travel times, and the effectiveness of repeated inspections.

3. **To establish a tactical resource allocation framework that balances efficiency and spatial fairness.** The objective is to determine how enforcement resources should be distributed across different areas of a city over a longer planning horizon, ensuring deterrence while supporting policy goals such as equity across neighborhoods.

These objectives collectively address both the behavioral and operational dimensions of parking enforcement, while recognizing that strategic allocation and day-to-day routing decisions must reinforce one another.

1.2 Thesis Outline

The remainder of this thesis is organized as follows. Chapter 2 reviews the relevant literature on parking policy and governance, behavioral modeling of compliance, and optimization approaches for enforcement planning. Chapter 3, Synthesis, situates the three methodological studies within a coherent multi-horizon research program. Chapters 4–6 present the core developments: behavior-aware multi period routing, multi-visit routing and scheduling with recovery-time constraints, and tactical resource allocation under equity requirements. Chapter 7, summary of work, synthesizes the findings, emphasizing how behavioral insights and operational strategies interact in practice and distilling managerial implications. Chapter 8 concludes with overarching contributions, limitations, and directions for future research.

CHAPTER 2 LITERATURE REVIEW

This chapter surveys interdisciplinary work spanning parking policy, behavioral modeling, smart enforcement technologies, and operational optimization, with the aim of revealing the theoretical and methodological gaps that motivate the dissertation. Parking sits at the nexus of land-use and transport policy, yet the evidence base guiding many parking decisions has historically been thin and politically constrained by tensions between regeneration, restraint, and revenue objectives [1]. Classic reviews further underline both the centrality of parking policy and persistent uncertainties around its impacts on access, the local economy, and travel behavior, underscoring the need for more rigorous analytical foundations [2, 3].

We therefore begin broadly, with urban parking governance and its shift from "predict-and-provide" supply expansion to demand management and sustainability, before narrowing to the behavioral mechanisms that shape compliance, the data/technology stack that measures curb activity, and the operational optimization frameworks for allocating and routing enforcement resources, highlighting the unaddressed integration of these fields. European neighborhood reviews describe this policy transition toward transportation demand management and the growing competition for scarce curb space, reinforcing why enforcement is a critical implementation lever [2]. Complementing the policy lens, the economics literature highlights welfare losses from mispriced or weakly enforced curb space and calls for integrating behavioral response, microsimulation, and empirically grounded elasticities (precisely the bridge this dissertation builds by embedding driver response within enforcement allocation and routing models [3]).

2.1 Urban Parking Policy and Management

This section provides an overview of how cities govern and manage curb space, tracing the evolution of parking policies from supply-oriented planning to contemporary demand management approaches. It highlights how pricing, spatial regulation, and neighborhood-level governance shape parking availability and use, while also noting the challenges created by limited curb capacity and competing stakeholder interests. By outlining both the policy foundations and the operational gaps in enforcement practices, this section sets the stage for understanding why regulation alone does not ensure compliance. However, the effectiveness of these tools depends on credible on-street enforcement.

2.1.1 Policy Evolution and Governance

Parking policy has evolved significantly from traditional zoning-based mandates to contemporary tools emphasizing transportation demand management and sustainability. Early frameworks often imposed minimum parking requirements tied to land use (typically designed to meet peak demand for free parking). This conventional “predict-and-provide” logic resulted in overbuilt infrastructure, distorted travel behavior, and entrenched car dependency [4].

Historical reviews underscore how planners relied on legacy handbooks and peer-city surveys to define supply thresholds, with little empirical grounding [1, 5]. In North America and Australia, for example, minimum off-street parking provisions were widely adopted without critical evaluation of urban context or transport alternatives [4]. Recent critiques challenge these norms, advocating for flexible standards or parking maximums in dense, transit-rich zones [6].

This historical trajectory aligns with broader goals of urban governance. Early work by McShane and Meyer emphasized that parking policy is inseparable from larger municipal objectives such as economic vitality, land conservation, and accessibility [7].

A shift is now underway: cities have begun deploying more adaptive management tools, metered pricing, dynamic time limits, permit zones, and geofenced areas, to regulate curb space and promote turnover [2, 8, 9].

However, governance literature continues to downplay the operational and logistical realities of enforcement. The relationship between regulation and compliance is not straightforward: chronic illegal parking often arises where enforcement capacity is limited, penalties are low, or monitoring strategies are poorly aligned with policy intent [10].

More recent studies similarly emphasize that enforcement is a necessary enabling condition of any curb management intervention. Without credible monitoring, even well-designed pricing policies fail to influence behavior or achieve turnover [11].

Furthermore, parking governance increasingly involves multiple institutional actors, municipal agencies, private operators, and residents, whose incentives do not always align. Game-theoretic analyses of shared governance arrangements show that pricing and access policies may be undermined when responsibilities for enforcement and revenue distribution are fragmented [12].

This disconnect is also visible in jurisdictions confronting illegal parking as a systemic governance issue rather than merely a behavioral problem. Recent municipal case studies demonstrate that addressing illegal parking requires not only regulation but institutional coordination and infrastructure investment [13].

Yet few policy models incorporate the planning, routing, or resource allocation strategies needed to operationalize dynamic policy tools. This omission is significant, as the effectiveness of pricing or regulatory instruments hinges on the presence of credible enforcement [1,3].

Taken together, the governance trends and policy instruments reviewed above show that curb rules only matter insofar as they are made credible on the street. This naturally points from policy design to the practical question of how presence is organized in space and time, setting the stage for the next subsections where human behavior, information, and logistics begin to interact at the curb. Because curb pricing and access controls operate through drivers' expectations, we next examine how individual parking decisions respond to enforcement and perceived risk.

2.1.2 Economic and Environmental Implications

The economic consequences of underpriced or unmanaged curb space are substantial. When parking is offered free or below cost, the resulting externalities include increased congestion, circulation delays, pollution, and inefficient use of valuable urban land [5, 9]. Shoup's seminal critique highlights that minimum parking requirements and free parking embed a large implicit subsidy for automobile use, shifting the cost of parking provision onto society and reinforcing car-dependent urban form [14].

A major driver of these externalities is *cruising*: motorists circulating in search of low-cost curb spaces. Empirical studies estimate that cruising can account for up to 30% of traffic in dense urban cores, substantially increasing vehicle-kilometers traveled and elevating emissions [8, 9]. Economic models show that when curb prices are set too low relative to demand, drivers rationally choose to continue searching, creating pure welfare loss; optimal, demand-responsive pricing can reduce cruising and congestion [15]. Simulation results show that non-myopic, location-dependent dynamic pricing with reservations, solved via approximate dynamic programming, consistently outperforms myopic rules and reliably handles spatial and temporal demand variation; moreover, infrequent price updates (e.g., every few weeks) limit effectiveness in practice [16]. Dynamic pricing approaches, which adjust curb rates according to occupancy, have been shown to reduce search time, emissions, and turnover inefficiencies [17].

Environmental burdens associated with cruising and idling are not spatially neutral. Freight and commercial parking shortages, in particular, lead to localized emissions and air quality impacts that disproportionately affect lower-income or environmentally burdened neighborhoods [18]. Thus, curb management is also an environmental justice concern.

However, even well-designed pricing or access rules depend critically on credibility and enforcement. Without visible and consistent enforcement, pricing signals weaken and compliance erodes [8, 19]. Studies show that insufficient fine levels or inconsistent enforcement can undermine demand moderation goals and even increase car use [20]. Similarly, price-based regulation may fail if behavioral responses and dwell-time elasticities are not considered, potentially increasing circulation rather than reducing it [21]. Integrated frameworks therefore emphasize that pricing and enforcement must be co-designed: policy establishes the incentive, but enforcement determines whether the incentive is believed and acted upon [11, 22].

In this sense, enforcement planning is not a secondary operational detail but a core determinant of policy effectiveness. Integrating enforcement into curb pricing and access strategies is essential to align behavioral response with intended economic and environmental outcomes [3, 23]. The stakes of misaligned policy and enforcement manifest directly on the street—as search behavior, congestion, and spatially uneven burdens. The next section therefore examines how drivers adjust their parking choices in response to perceived enforcement presence, establishing the behavioral foundations required to design effective, compliance-supportive enforcement operations.

2.2 Behavioral Foundations of Parking and Compliance

This section examines the behavioral mechanisms driving parking decisions and compliance with regulations. It emphasizes how drivers form perceptions of convenience, cost, and risk, and how these perceptions are shaped by personal habits, social norms, and past interactions with enforcement. The discussion underscores that illegal parking is not solely the result of low penalties or insufficient enforcement, but of adaptive decision-making in a dynamic environment where expectations and behaviors co-evolve.

2.2.1 Driver Decision-Making and Risk Perception

Parking decisions involve a continuous evaluation of trade-offs among monetary cost, convenience, enforcement risk, and trip purpose. In standard economic terms, the choice between legal and illegal parking can be framed through the *expected cost of violation*, defined as the fine multiplied by the perceived probability of detection [3]. However, drivers rarely know true enforcement intensities and do not compute probabilities explicitly. Instead, they form and update beliefs from past experience, visible patrols, and social cues. Enforcement therefore acts as a learning signal: personalized feedback on past violations has been shown to reduce subsequent illegal parking by approximately 14%, whereas generic reminders do not change behavior [24]. More broadly, compliance is embedded in a shared "culture of driving," where behaviors adjust in response to the observed norms and sanctions in a given area [25].

Beyond fully rational calculation, drivers frequently rely on boundedly rational heuristics under time pressure. Simple lexicographic rules (e.g., "take the closest acceptable spot") often guide parking search decisions and can even outperform fully optimized strategies in realistic settings [26]. Empirical models of search behavior show that motorists evaluate options sequentially rather than solving the problem in advance, leading to search lengths that can exceed rational optima [27]. Importantly, heterogeneity in schedules, time sensitivity, and curb access needs results in distinct parking segments. A recent synthesis of 34 studies finds that commercial vehicle (CV) drivers, in particular, exhibit different parking strategies: when curb space is scarce, they frequently accept the risk of temporary illegal stopping rather than cruise for parking, given tight delivery schedules and inelastic service demand [28, 29]. Structural equation and stated-preference studies further show that parking cost and walking distance consistently dominate decision-making, while contextual street conditions play a secondary role [30, 31].

Taken together, the evidence indicates that risk perception is dynamic, heterogeneous, and shaped by recent enforcement visibility. The same street segment may experience substantially different compliance behaviors across the day depending on purpose, time pressure, and perceived patrol presence. This behavioral granularity motivates operational strategies that vary patrol cadence and location to influence expectations—shifting the focus from static pricing or signage toward the spatial–temporal pattern of enforcement activity itself.

2.2.2 Compliance, Social Norms, and Displacement

Classical deterrence theory holds that stronger and less predictable enforcement increases compliance by raising the expected cost of violation [32]. Yet compliance is not determined solely by fines and detection probability. Cullinane and Polak emphasize that adherence also depends on the perceived legitimacy and clarity of street-use priorities [10]. When regulations are viewed as fair and enforcement is seen as consistent, voluntary compliance strengthens. For example, residential permit programs that align with neighborhood expectations tend to generate higher acceptance and lower conflict [33]. Conversely, when enforcement is perceived as arbitrary or punitive, even well-designed rules may fail to achieve behavioral change.

Public attitudes further illustrate this divergence. A case study in Edinburgh found that, despite high violation rates, residents overwhelmingly viewed enforcement as insufficient rather than excessive, suggesting that *consistency* rather than strictness is the key driver of perceived fairness [34]. Moreover, external policy shifts, such as tariff changes or transit investments, restructure demand spatially as much as they reduce it. A large-scale time-series analysis of Amsterdam parking transactions showed that price increases and metro expansion reduced central parking volumes while displacing demand to peripheral zones and park-and-ride facilities [35]. Similar spatial redistribution patterns emerge when built environments facilitate or constrain stopping opportunities: illegal parking clusters intensify around schools, hospitals, and institutional districts when legal curb supply is insufficient [36].

These displacement dynamics extend to commercial vehicles. The CV parking literature shows that expanding loading capacity, improving curb information, or tightening enforcement alters the balance between legal and illegal stops, stop durations, and cruising time [28]. Social norms also shape compliance. A recent field experiment found that enforcement-framed notices reliably increased fine repayment across user groups, whereas norm-framed appeals had more limited and segmented effects [37]. Likewise, evidence suggests that drivers behave more cautiously when operating outside their home jurisdiction, where social identity makes reputational consequences more salient [38].

Taken together, these findings show that compliance emerges from an interaction of deterrence, legitimacy, and behavioral expectations—not cost signals alone. Because drivers update beliefs from observed patrol rhythms, the *spatial and temporal pattern* of enforcement activity matters: it influences not only where violations occur but when they reappear and how they migrate across the network. This highlights the need for enforcement strategies that manage compliance dynamically, using patrol cadence and visibility to shape expectations. The next section therefore considers sensing and monitoring systems that support such behavior-aware operational planning.

2.3 Smart and Data-Driven Parking Systems

This section reviews technological and data-driven approaches that increasingly support parking management and enforcement. It introduces sensing infrastructures, digital monitoring tools, and analytics platforms that enable real-time visibility of curb activity and violation patterns. The section also highlights how these systems can uncover spatial-temporal regularities and inequities, providing the informational basis for more targeted, proactive, and adaptive enforcement strategies.

2.3.1 Technology Landscape and Architectures

Recent advances in sensing, connectivity, and distributed computing have transformed parking management from a largely reactive activity into a data-rich, networked service within broader smart city infrastructures. Modern smart parking systems typically consist of three integrated layers: a *sensing layer* responsible for occupancy detection, a *communication layer* that transmits data, and a *data-processing layer* that supports decision-making and service delivery [39].

Sensing technologies vary widely in their form and precision. They include inductive loops, ultrasonic detectors, magnetic field sensors, and camera-based systems, each offering distinct trade-offs between installation complexity, scalability, and accuracy [40].

Meanwhile, automatic license-plate recognition (ALPR) platforms enable continuous monitoring of curb spaces and structured parking facilities through edge-based computer vision [41].

In practice, ALPR-equipped patrol vehicles and handheld devices have become central enforcement tools in many cities, enabling continuous scanning of plates and time-stamped parking records. A field evaluation of license-plate–based surveying demonstrated that optimizing vehicle scanning paths can significantly improve coverage efficiency while reducing enforcement labor hours [42].

To reduce cloud latency and bandwidth pressure, recent architectures increasingly adopt *edge* and *fog computing* designs, processing visual and sensor data near the source rather than centrally [43]. This approach lowers transmission delays, improves responsiveness, and enables real-time control, especially critical in dense urban conditions.

However, the integration of sensing and data platforms has not automatically produced proactive enforcement strategies. Shao et al. propose a data-driven patrolling framework in which large-scale in-ground sensor logs are combined with machine learning models (LSTM auto-encoders) to predict violation persistence and guide officer routing; this connects smart sensing directly to operational optimization rather than passive data display [44].

Despite these technological advances, deployment strategies in many cities remain primarily operational rather than optimization-driven. Systems tend to provide real-time occupancy information without integrating predictive or prescriptive analytics for proactive enforcement or dynamic curb allocation [41]. As a result, the potential of these systems to reshape enforcement scheduling, violation deterrence, and long-run compliance behavior remains underutilized.

Modern sensing and data platforms convert curb activity into timely signals, but signals alone do not set priorities. Their value is realized when they inform where and when officers appear next time, so the discussion turns from architectures to how analytics highlight the specific blocks and windows where presence changes outcomes.

2.3.2 Data Analytics for Hotspots and Equity

Smart parking systems and curbside sensing technologies generate high-resolution spatial and temporal data that can be used to identify where and when parking violations recur. Hotspot detection methods, including kernel density estimation and statistically robust spatio-temporal clustering, help agencies detect persistent violation clusters and peak-demand periods [45]. Time-series forecasting models likewise support anticipatory patrol planning by predicting occupancy surges, event-driven pressures, and seasonal variation. Techniques originally developed for criminology and public safety hotspot detection have also been adapted to parking contexts to improve proactive patrol targeting [46].

Recent spatial econometric analyses further confirm that parking violations are strongly shaped by built environment features and land-use intensity. A geographically weighted Poisson regression of Nanjing shows that illegal parking clusters are concentrated around schools, medical facilities, and institutional corridors with constrained curb supply [36]. These findings indicate that hotspots are not simply "busy areas," but arise from structural access needs and curb scarcity—suggesting that uniform enforcement is unlikely to resolve concentrated non-compliance.

Beyond efficiency, data analytics reveal important equity dimensions in parking enforcement. Studies show that fines and street-management burdens can fall disproportionately on low-income and minority neighborhoods, raising concerns over fairness and distributional impact [47]. Similarly, research on freight and commercial-vehicle parking highlights that prolonged search time and idling frequently occur in environmental justice communities along freight corridors, concentrating emissions and exposure in vulnerable areas [18]. Hotspot analysis of double-parking further shows that even localized non-compliance can trigger broader network effects: eliminating double-parking in Athens increased average travel speeds by up to 44% and reduced delay by up to 33% [48]. Thus, hotspots carry both congestion and pollution externalities, not merely localized inconvenience.

Extending these insights, recent transportation infrastructure equity studies demonstrate that resource allocation and maintenance often correlate with socioeconomic characteristics, indicating systemic unevenness in municipal prioritization [49, 50]. Applying similar benchmarking to parking enforcement data can support equity-aware patrol allocation strategies that account not only for violation intensity, but also for neighborhood vulnerability.

Taken together, hotspot and equity analyses show that violations cluster with land use and that their impacts are uneven across neighborhoods. However, descriptive analytics alone do not determine how limited enforcement units should be deployed. The critical step is converting predictive insights into *prescriptive* patrol strategies that schedule repeat visits, vary presence over time, and align enforcement effort with both efficiency and fairness goals. The next section examines how such operational planning and routing decisions are structured.

2.4 Operational and Enforcement Optimization

This section focuses on the operational dimension of parking enforcement, where limited enforcement resources must be scheduled, allocated, and routed across complex urban networks. It traces the development of models that treat enforcement as a routing and patrol planning problem, ranging from simple beat assignments to advanced formulations that account for time-varying violation opportunities. By connecting enforcement logistics with spatial and temporal patterns of illegal parking, this section shows how operational design influences on-the-ground outcomes.

2.4.1 From Heuristics to Formal Models

Parking enforcement has historically evolved from manual and locally improvised patrolling practices to systematically optimized inspection schedules grounded in algorithmic routing models. Early parking enforcement was primarily labor-intensive, with officers assigned to fixed beats or cyclic patrols based on local knowledge and rules of thumb. These practices emphasized spatial coverage but rarely considered dynamic patterns of violation behavior, officer travel times, or the probabilistic nature of detection. Over time, advances in the fields of vehicle routing, graph theory, and inspection games enabled the development of formal models that treat parking enforcement as a specialized routing problem with time-dependent rewards and strategic user response.

Early Inspection and Staffing Approaches. The first analytical models did not treat patrolling as a routing problem but rather as a *staffing and coverage* problem. In a seminal case study from Singapore, the inspection of coupon-based parking facilities was modeled to determine optimal staffing levels and patrol schedules for wardens deployed across multiple lots [51]. These models emphasized the allocation of personnel across space and time, but patrol movement itself was treated in simplified terms, assuming either circular or grid-based patterns. The underlying logic reflected a *coverage paradigm*: violations would be minimized if wardens appeared regularly enough to increase the perceived likelihood of inspection.

A similar logic appears in early parking enforcement policy studies, which often assumed that compliance was primarily a function of fine levels and nominal enforcement intensity [32]. However, these frameworks did not optimize routing nor account for dynamic violation occurrence or strategic driver adaptation.

Emergence of Routing Formulations. By the late 1980s and 1990s, researchers began to cast patrol problem as an *arc-routing problem*, recognizing that enforcement officers may traverse street networks rather than nodes [52]. The Chinese Postman Problem (CPP) and its variants became foundational [53,54]. In the context of enforcement routing, the objective was to minimize traversal distance while ensuring required-edge traversal for inspection coverage. This formulation shifted focus from staffing to *efficient spatial patrol coverage*, which allowed the use of exact solvers and heuristics common in transportation network optimization.

The development of the "parking warden tour problem" explicitly recognized parking enforcement as a structured variant of arc routing, where each street segment contains a set of curbs requiring inspection rather than classical service points [55]. These models improved route efficiency but still assumed deterministic violation payoff: any inspection yields either zero or one detected violation if present.

Incorporating Violation Probabilities and Time-Sensitive Payoffs. As sensor technologies and violation timestamping systems matured, researchers began acknowledging that *violation opportunities are time-dependent.* A vehicle may become ticketable at a specific moment, and may also leave before an officer arrives. Thus, the enforcement problem became one of *maximizing expected detections*, requiring probabilistic modeling.

One of the earliest explicit formulations of the detection–reward relationship appeared in the economics literature, where the optimal fine and enforcement intensity were derived as joint determinants of compliance [56]. However, the full operational implications for routing were not yet addressed.

A breakthrough occurred with probability-based inspection frameworks, where each potential inspection location yields a time-decaying reward. The payoff of inspecting a vehicle diminishes as the delay increases, capturing the likelihood that the violator has already departed. This shift reframed enforcement as an analog to *deadline-constrained or stochastic prize-collecting routing.*

The Traveling Officer Problem (TOP) formalized this perspective, defining the enforcement task as choosing officer paths to maximize expected violation detections given time-based decay [57]. TOP marked a conceptual shift: enforcement was no longer about coverage, but *catching violations before they expire.*

Integration of Sensor Data and Machine Learning. The proliferation of in-ground parking sensors, ALPR systems, and camera-based curb monitoring enabled the estimation of violation duration distributions and spatial–temporal patterns. This made it possible to learn the probability that a violation persists over time and incorporate those estimates into routing optimization.

For instance, Shao et al. developed a machine learning–enhanced TOP solution that uses spatio-temporal sensor logs and LSTM auto-encoders to estimate violation persistence and guide real-time route selection [44]. These models move enforcement beyond static schedules toward adaptive routing systems that respond dynamically to changing curb conditions.

Exact and Heuristic Optimization Approaches. With increased problem complexity, researchers explored both exact and heuristic solution methods. Exact methods (e.g., integer programming, branch-and-cut) can optimally solve medium-sized arc-routing formulations with time constraints, but scale poorly in dense street networks.

Heuristic frameworks—such as greedy time-aware selection, simulated annealing, tabu search, and ant colony optimization—emerged as practical alternatives for real-world deployment. For example, ant colony optimization has been shown to outperform both naive patrol patterns and simple greedy approaches in maximizing detections under time limits [58].

Unifying View: Parking Enforcement as Stochastic Arc Routing with Strategic Feedback. Modern models increasingly adopt a unifying interpretation: parking enforcement is a *stochastic, time-sensitive arc-routing problem* embedded in a behavioral feedback loop. Patrol presence influences driver expectations; expectations influence parking decisions; parking decisions alter violation opportunities; and violation opportunities shape optimal patrol routing. Thus, objective functions now reflect not only expected detection yield but also *long-run compliance effects*, where consistent spatial presence reduces violations even if immediate citations fall.

Summary. The field has progressed from:

1. *Manual, rule-based patrol scheduling* (coverage focus),
2. to *Deterministic arc-routing models* (efficiency focus),
3. to *Stochastic, payoff-sensitive routing* (expected detection focus),
4. to *Machine learning–assisted dynamic routing* (predictive enforcement).

This progression lays the foundation for the next subsection, where enforcement is embedded in bi-level and equilibrium frameworks that explicitly model how drivers adapt to enforcement intensity over time.

2.4.2 Patrolling Taxonomy and Transferable Insights

Parking enforcement shares deep conceptual foundations with the broader police patrolling literature, where authorities must allocate limited mobile resources across space and time to discourage undesirable behavior. A comprehensive review of police patrol models identifies three interdependent decision layers: *strategic* (district design), *tactical* (resource allocation), and *operational* (route and schedule planning) [59]. This layered perspective is echoed in recent police planning work, in which districting, patrol assignment, and routing are treated as coordinated components of the same enforcement system [60].

Strategic Level: District Partitioning.

Strategic districting divides the urban area into zones that balance workload, maintain territorial coherence, and support predictable operational coverage [61]. Recent approaches use spatial optimization and simulation to redesign districts in ways that improve workload equity and service access [62]. Network-based districting further aligns zone boundaries with street connectivity, enabling more realistic patrol mobility [63]. Although parking enforcement does not require emergency response coverage, the principles of *equity and exposure balance* are directly relevant. If some neighborhoods consistently receive fewer inspections, patterns of illegal parking will concentrate there. Strategic districting thus provides a structured way to prevent geographic enforcement inequity.

Tactical Level: Resource Allocation.

At the tactical scale, patrol units are distributed across zones based on expected violation or incident levels. Early allocation models maximize the probability of detecting offenses given limited officer capacity [64], while stochastic patrol allocation frameworks determine revisit frequencies needed to sustain deterrence [65]. More recent multi-period location-allocation models adjust deployment across time to enhance visibility and balance workloads [66, 67]. Multi-coverage extensions account for hotspots that require repeated attention [68]. In parking enforcement, these allocation principles translate directly: expected compliance improvement replaces expected crime interception, but the mechanics of distributing limited patrol attention remain the same.

Operational Level: Routing and Scheduling.

Operational planning determines how officers traverse the network during a patrol shift. Patrol routing has evolved from early coverage heuristics to formulations based on orienteering with time, priority, and revisit constraints. Maximum-covering patrol routing models jointly optimize route feasibility and enforcement coverage [69]. Time-dependent security-level decay models similarly represent how the deterrent effect of a patrol visit fades and must be renewed through appropriately timed revisits [70]. These concepts directly parallel parking enforcement, where violations re-emerge over time and the benefit of inspection is inherently time-sensitive.

Transferable Insights to Parking Enforcement.

1. **Randomization Sustains Deterrence.** Predictable patrol loops lower perceived detection risk; strategic variability improves compliance [70].
2. **Equity Must Be Explicitly Modeled.** Districting work shows that fairness does not emerge automatically and requires deliberate constraints [62].
3. **Deterrence is Time-Dependent.** Violations regenerate, so routing must incorporate revisit intervals rather than single-pass coverage.
4. **Compliance is Adaptive.** Drivers respond to perceived enforcement patterns; state-dependent or stochastic violation models capture this feedback more realistically [69].

Implications for This Research.

Parking enforcement planning therefore benefits from the same layered reasoning that guides police patrol design:

$$\begin{aligned}\text{Strategic Districting} \;+\; \text{Tactical Resource Allocation} \;+\; \text{Operational Routing}\\ +\,\text{Deterrence Feedback}\end{aligned}$$

The optimization frameworks developed in later chapters build directly on this structure. They model enforcement as a dynamic interaction in which patrol presence shapes expectations, expectations shape compliance, and compliance patterns evolve over time. This perspective motivates the behavior-aware routing and allocation methods introduced in the following sections.

2.5 Integrating Behavioral and Optimization Perspectives

This section synthesizes the behavioral and operational perspectives introduced earlier, framing parking enforcement as an interactive system in which driver decisions and enforcement strategies influence one another over time. It introduces equilibrium and game-theoretic approaches that capture this feedback loop, as well as routing models that incorporate behavioral adaptation and deterrence dynamics. This integrated perspective provides the conceptual link between enforcement planning and compliance outcomes.

2.5.1 Bi-level/Equilibrium and Inspection-Game Models

Parking enforcement is best understood as a system of strategic interaction in which drivers and enforcement authorities continuously respond to one another's decisions. The central mechanism through which this interaction unfolds is the *perceived probability of detection*, which links patrol allocation decisions to driver compliance behaviors. As these perceptions evolve through repeated observation of enforcement patterns, compliance outcomes form an endogenous equilibrium rather than a static or predetermined state. This perspective motivates the study of parking enforcement using *bi-level* optimization and *inspection-game* formulations, where enforcement decisions represent upper-level policies, and driver reactions represent lower-level behavioral equilibrium.

Early research recognized that parking behavior cannot be adequately represented through comparative-static models because enforcement and compliance adjust over time in response to each other. Cullinane [71] argues that parking systems exhibit path dependence and memory effects, such that compliance in one period is shaped by past inspection visibility. This insight underpins the shift from descriptive parking studies toward models that explicitly incorporate temporal feedback.

A major advancement in this direction is introduced in Nourinejad and Roorda's work on commercial vehicle enforcement equilibria [72]. Their model treats illegal parking as the outcome of a bilateral meeting process between parked vehicles and enforcement officers, producing an endogenous citation probability. This formulation implies that increasing patrol density has diminishing returns when illegal parking volume is either too low (officers cannot find violators) or too high (violators cycle quickly). A similar meeting-based equilibrium approach appears in related work on steady-state illegal parking in downtown delivery and freight contexts [73].

The behavioral structure underlying such equilibria is reinforced by Guo et al. [74], who demonstrate that parking choice reflects *bounded rationality* under uncertainty rather than perfectly anticipatory optimization.

Further extensions recognize that equilibrium enforcement interacts with broader urban travel demand. Petiot [75] shows that parking enforcement and pricing can unintentionally increase car travel when fines are used as substitutes for congestion management rather than complements.

Similarly, Gur and Biemborn [76] develop one of the earliest equilibrium frameworks linking parking supply, travel demand, and spatial assignment patterns. These foundational models emphasize that enforcement cannot be separated from broader demand and occupancy dynamics—a key principle retained in this dissertation.

Bi-level programming formalizes this interdependence. In this structure, the *upper level* represents the enforcement authority choosing patrol intensity, fine levels, or zone allocation strategies, while the *lower level* models drivers optimizing their payoffs given enforcement conditions. Duan et al. [77] demonstrate such a bi-level model in shared parking allocation, in which operators allocate capacity and drivers optimize cost and access. Although their context concerns facility allocation rather than enforcement, the structure parallels the patrol–compliance feedback modeled here.

A further evolution appears in governance models involving multiple stakeholders. Hu et al. [12] develop a tripartite game among municipal authorities, operators, and drivers, showing that equilibrium outcomes depend on whether pricing incentives are aligned across actors.

Inspection Games and Randomized Deterrence.

While bi-level models explain equilibrium outcomes under deterministic patrols, enforcement is rarely deterministic in practice. When drivers learn patrol routines, they adapt—reducing deterrence. This motivates *inspection games* in which the enforcement authority randomizes its patrol strategy.

Brotcorne et al. [78] formulate fare inspection routing in transit systems as a Stackelberg game, where inspectors commit to mixed strategies to avoid being anticipated.

Trejo et al. [79] use extraproximal methods to compute *strong* Stackelberg equilibria, demonstrating how mixed strategies can be computed efficiently even under incomplete information.

Garrec [80] develops continuous-time patrolling and hiding games, emphasizing how mobility, location uncertainty, and attacker timing jointly shape optimal routing.

Recent work shows that randomized enforcement is not only theoretically desirable but also necessary in real-world public safety domains. Naja et al. [81] introduce the STOP model for speed trap deployment, demonstrating that predictable enforcement loses deterrent value.

Rosenfeld et al. [82] implement a deployed traffic enforcement allocation system used by the Israeli Police, showing measurable reduction in crashes when enforcement is spatially and temporally randomized.

Wang and Cui [83] extend inspection games to include coordination between mobile patrols and fixed monitoring devices, demonstrating that deterrence depends on the *joint* allocation of sensors and patrol units.

Finally, Rezazadeh et al. [84] show how periodic inspection of hazardous pipeline networks can be modeled analogously to patrol routing with risk accumulation, reinforcing the generality of inspection-game logic in infrastructure systems.

Positioning Within This Dissertation.

The insights from these models converge on three principles that shape the behavioral–operational framework developed in this dissertation:

1. *Compliance is endogenous* and depends on perceived, not nominal, enforcement intensity.

2. *Equilibrium outcomes emerge from repeated observation and adaptation*, requiring dynamic modeling.

3. *Deterrence is sustained by strategic unpredictability*, motivating randomized or varied patrol schedules.

These principles directly motivate the transition to the behavior-aware routing models developed in the following subsection, where compliance and violation potential evolve as a function of *time since last visit* to each curb segment.

2.5.2 Behavior-Aware Routing and Scheduling

While the previous subsection emphasized equilibrium relationships between enforcement intensity and driver compliance, the practical enforcement challenge centers on *how* officers should move through the network to sustain deterrence over time. A key insight emerging from the parking enforcement and patrolling literature is that violation incentives are inherently *time-dependent*. The likelihood of illegal parking increases with the *time since last patrol*, meaning that deterrence decays as the memory of enforcement fades. Therefore, routing and scheduling decisions are not merely logistical problems of minimizing travel time—they are strategic decisions that shape behavioral outcomes.

This dynamic structure aligns with classical stochastic enforcement theory. Han and Levis [85] demonstrate that parking violations are best deterred not through fixed or deterministic patrols, but through *probabilistic and temporally varied enforcement*. This principle remains foundational: predictable patrols invite temporal displacement, while unpredictability sustains perceived detection risk. In this sense, routing is both an operational and behavioral control mechanism.

Recent studies explicitly incorporate behavioral response within patrol optimization models. Lei, Zhang, and Ouyang [86] develop a parking enforcement patrol model in which drivers decide whether to pay or evade based on anticipated patrol arrival patterns. Their study shows that optimizing routes without accounting for behavioral response can lead to enforcement patterns that maximize coverage but unintentionally lower compliance by becoming predictable. This work provides direct evidence that enforcement effectiveness depends on patrol *temporal spacing*, not just total time in a zone.

This behavioral dimension is also reflected in routing formulations derived from the orienteering family of problems. Bruglieri [55] introduces the Parking Warden Tour Problem, where the benefit of visiting a curb segment depends on how long it has been since the last inspection. Similarly, Corrêa et al. [87] extend this logic to a multi-period setting, showing that deterrence effects accumulate and decay across days and shifts. These models mirror the decay of perceived detection probability in parking enforcement: if an officer revisits too soon, there is little additional deterrent effect; if too late, violations have already accumulated. Thus, the central routing challenge is to allocate officers across space and time to balance deterrence decay rates.

A complementary line of work integrates predictive modeling into patrol optimization. Shao et al. [44] use LSTM auto-encoders to forecast violation patterns and embed these predictions within a routing heuristic to target areas with the highest imminent violation risk. This approach demonstrates the value of data-driven models capable of learning local behavioral rhythms, such as delivery windows, lunch peaks, and school dismissal times. Incorporating such temporal patterns transforms routing from static coverage to proactive behavioral anticipation.

Taken together, three mechanisms underpin behavior-aware enforcement routing:

1. **Deterrence decays with time since the last patrol**, so the marginal value of revisiting a segment is time dependent.

2. **Compliance responds to anticipated (not merely observed) presence**, which calls for randomized or strategically varied patrol patterns.

3. **Violations follow predictable temporal rhythms**, which can be learned and exploited through forecasting.

These principles motivate the modeling and solution approach adopted in this dissertation. We develop routing models in which each location has a *time-dependent violation yield*, capturing both diminishing returns from rapid repeat inspections and the regrowth of illegal parking when presence is absent. Embedding this dynamic into a multi-visit orienteering framework enables patrol schedules that sustain deterrence more efficiently than static or distance-only routing. The temporal structure also lets schedules adapt to context-specific rhythms (e.g., delivery windows, school dismissal, nightlife peaks). In short, behavior-aware routing shifts the focus from *covering space* to *managing compliance over time*, forming the foundation for the multi-period routing and Variable Neighborhood Descent methods developed later.

2.6 Discussion and Research Gap

The literature reviewed in this chapter reveals a persistent disconnect: enforcement is often modeled either as a fixed backdrop to policy or as a purely logistical routing exercise, rather than as a strategic lever whose effectiveness depends on how drivers perceive and adapt to it. Pricing reforms and access rules only work when supported by credible, sustained presence; yet most optimization models do not endogenize compliance, and most behavioral or economic studies do not specify how enforcement should be scheduled or allocated. Consequently, the causal link from patrol decisions (where, when, how often) to outcomes (violations, turnover, spatial fairness) remains under-specified.

A second gap concerns the separation of decision horizons. Operational routing studies optimize short-run inspection efficiency but commonly ignore adaptation to predictable patrols or uneven neighborhood coverage. Tactical allocation studies aim for long-run deterrence or equity across zones, but rarely incorporate time-dependent re-emergence of violations or the feasibility of multi-visit patrols within shifts. Addressing these gaps requires models that integrate driver response with revisit-sensitive routing at the operational horizon and equity-aware resource distribution at the tactical horizon. The following chapters develop precisely this synthesis: behavior-aware operational routing (Chapters 4–5) and equitable, detection-threshold–based allocation across regions (Chapter 6).

CHAPTER 3 SYNTHESIS

This dissertation is organized around three methodological developments that address parking enforcement planning at different decision horizons. The overall research program progresses from daily operational routing with behaviorally-informed enforcement effects, to enhanced multi-visit operational planning with explicit temporal dynamics, and finally to tactical allocation of enforcement resources at the city scale. Together, the three studies address the research objectives introduced in Chapter 1 by integrating behavioral response, operational constraints, and policy considerations into coherent models and solution approaches.

3.1 Operational Enforcement Planning Framework

The first research contribution, presented in Chapter 4, develops an operational planning framework that incorporates the behavioral responses of drivers to varying levels of enforcement presence. The core problem concerns determining feasible patrol routes for enforcement officers over a short-term horizon, such as a day or week, under the premise that enforcement frequency influences the perceived likelihood of receiving a citation.

In this study, the probability of illegal parking and the likelihood of detection are modeled as functions of officer visit frequency at each location. The resulting formulation extends a periodic routing framework in which the objective is to schedule enforcement activities to enhance deterrence and improve overall compliance. Because direct optimization of this formulation can be computationally demanding for realistic urban networks, the study introduces a matheuristic solution strategy that incrementally constructs, evaluates, and refines candidate patrol patterns.

Computational analysis is conducted using both synthetic datasets and a real case study involving off-street parking facilities in Montreal. The results show that incorporating behavioral effects yields patrol schedules that differ meaningfully from those generated under purely structural routing assumptions. This study directly contributes to the first research objective by providing a modeling framework that explicitly links enforcement presence to compliance behavior within an operational routing environment.

3.2 Multi-Visit Scheduling and Temporal Compliance Effects

The second research contribution, detailed in Chapter 5, extends operational planning to account for the re-emergence of illegal parking after inspection. In many urban zones, a single visit during a patrol period is insufficient to maintain compliance. However, revisiting a location too soon offers little enforcement value. To capture this, the study formulates a single-period planning problem in which locations may be visited multiple times, subject to a minimum recovery time between successive inspections.

This leads to a multi-visit routing formulation that includes temporal feasibility constraints and limits on total patrol duration. A Variable Neighborhood Descent metaheuristic is developed to explore routing alternatives while ensuring recovery time and coordination constraints remain satisfied.

Empirical results from synthetic data and the Montreal use case show that optimal and near-optimal patrol schedules typically involve repeated but temporally separated visits to high-violation-prone locations. This work supports the second research objective by developing operational scheduling tools that explicitly account for time-dependent enforcement effectiveness and limited officer availability.

3.3 Tactical Resource Allocation Under Equity Constraints

The third contribution, presented in Chapter 6, shifts from short-run routing to medium-term allocation of enforcement resources across urban regions. Here, the goal is to determine how a fixed number of officers should be distributed across different neighborhoods to produce sustained compliance, preserve fairness, and maintain revenue stability over longer horizons.

The study develops a behavioral deterrence mechanism that characterizes how enforcement intensity influences the expected cost of violating parking rules. This mechanism yields a closed-form threshold describing when enforcement transitions from weak to strong deterrence. The deterrence structure is incorporated into a tactical allocation model that assigns enforcement resources across regions while ensuring a minimum detection-probability requirement, thereby supporting equity among neighborhoods.

This contribution addresses the third research objective by providing a resource allocation framework that accounts for behavioral adaptation, spatial heterogeneity, and policy considerations such as fairness. Results demonstrate how moderate enforcement can maintain compliance efficiently while avoiding diminishing returns from excessive deployment.

3.4 Synthesis and Concluding Remarks

Chapter 7 synthesizes the insights gained from the operational and tactical studies, emphasizing how short-run routing strategies and medium-run allocation decisions interact with behavioral responses. Chapter 8 concludes the dissertation by summarizing key findings and identifying directions for future research, including adaptive enforcement scheduling, integration with real-time data collection, and coordination with broader urban mobility policies.

CHAPTER 4 ARTICLE 1: PARKING ENFORCEMENT OPTIMIZATION VIA KERNEL SEARCH AND BEHAVIORAL–OPERATIONAL COUPLING

Authors: Mohsen Yahyaei, Michel Gendreau and Nikolaj Van Omme, Submitted on November 2, 2025, Under review at European Journal of Operational Research

Abstract: This paper addresses illegal parking problems in urban environments by developing a novel decision framework for optimizing parking enforcement strategies. Recognizing the limitations of existing approaches, we integrate driver behavior and enforcement dynamics into a comprehensive mathematical model. The model captures the interplay between patrol routes, citation probabilities, and driver decisions, allowing for a more realistic representation of the enforcement process. Due to the model's computational complexity, commercial solvers often struggle with medium- and large-sized instances, motivating the development of math-heuristic algorithms. To address this, we propose two Kernel Search (KS) algorithms that use iterative refinement to efficiently produce high-quality solutions in reduced computational time. Extensive testing was performed on randomly generated problem instances, and the results demonstrated the superior performance of the proposed algorithms in solving larger problems compared to Gurobi. In addition to synthetic tests, the algorithms were applied to a real-world case study involving the inspection of parking lots in Montreal. The results from this case study confirmed the efficiency of the KS algorithms in practical applications.

Keywords: Parking Enforcement Optimization, Periodic Routing Problem, Kernel Search Heuristic, Driver Behavior Equilibrium, Matheuristic Algorithms

4.1 Introduction

The growing demand for parking spaces in Montreal underscores the need for effective management and strategic urban planning. In 2024, the Agence de Mobilité Durable [88] increased the number of paid parking spaces by 10%, adding 1,735 meters, bringing the total to over 19,000, and projected revenues of $73 million, up from $63 million in 2023. Acknowledging that merely increasing parking capacity is not enough, there is also a need to prioritize stricter enforcement policies. Such enforcement plays a vital role in preserving traffic flow, enhancing safety, and bolstering economic sustainability [89].

Research highlights that cities with proper parking enforcement see marked improvements in compliance rates. For example, localized enforcement in European cities increased compliance from 20% to 70-80% after adopting digital tools and modern systems [90]. Neglecting enforcement can lead to severe consequences, including accessibility issues and revenue loss for businesses. For instance, improperly parked vehicles can obstruct emergency services and deter customers from local businesses [91].

A way to control illegal parking is through effective enforcement, which relies on three core components: detection technology, citation fines, and the level of enforcement [72]. Detection technology, such as human patrolling, surveillance systems, and video detection, provides the foundational capability to identify violations, though it often involves long-term investments and technological upgrades. Citation fines act as financial deterrents, but their implementation is usually governed by higher authorities and may not be adjustable by local agencies or private entities managing parking operations. In settings where detection technology and fines are constrained, the level of enforcement is the pivotal lever. It can be expressed at the tactical scale as the spatio-temporal density of officer presence in a region, or equivalently at the operational scale as the expected visit frequency per blockface (the number of passes an officer can make along a segment within the operational horizon). A higher enforcement level not only deters violations by increasing the perceived risk of detection but also enhances compliance rates, particularly when combined with significant fines.

Illegal parking can be examined from multiple perspectives, including psychological, operational, and spatial factors. Song et al. [92] explored drivers' psychological responses and found that perceived behavioral control strongly predicts willingness to avoid violations, even more than moral attitudes against illegal parking. Publicizing monetary damages and implementing regular enforcement can mitigate these behaviors. Meanwhile, Gao et al. [93] identified spatial dependencies and selection biases in parking citation data, showing that enforcement coverage and local conditions like land use and traffic affect violation patterns. The prevalence of commercial vehicles in double parking violations highlights operational pressures, such as limited parking near delivery points, further complicating enforcement efforts. These findings underscore the complexity of addressing illegal parking, which requires integrating psychological insights, spatial analysis, and targeted policy design.

Enforcing parking regulations often involves patrolling areas to issue citations for violations. These areas typically include multiple parking lots or streets, and planning an efficient route for patrols is important, especially given the limited human resources (officers) available to most agencies. This problem can be approached through the lens of the routing problem, a class of optimization problems that focuses on servicing a set of points (customers, demand locations, or parking lots in this case) efficiently.

In this context, the routing problem is adapted to scenarios where parking enforcement officers are required to visit multiple locations to conduct inspections and issue citations systematically [86, 94]. Similar to vehicle routing problems used in logistics for deliveries or service visits, parking enforcement routing requires optimizing an objective while managing limited available resources. By applying optimization techniques to determine the most efficient routes, agencies can maximize revenue, ultimately improving compliance with parking regulations.

By considering the effects of patrol planning on illegal parking, this paper develops a comprehensive decision framework to address the complexities inherent in parking enforcement planning. This framework aims to provide a novel decision support tool for parking managers, enabling them to allocate resources effectively and design patrol routes that enhance compliance while accounting for practical constraints such as operational time limitations.

The remainder of this article proceeds as follows. Section 2 provides a comprehensive review of existing literature. Section 4.3 delves into the theoretical underpinnings of the problem. Section 4.4 presents the proposed model in detail. A novel Kernel Search (KS) matheuristic approach for solving the problem is introduced in Section 4.5. Computational results of the current study are presented and analyzed in Section 4.6. Finally, Section 4.7 offers concluding remarks and outlines potential avenues for future research.

4.2 Literature Review

The effectiveness of enforcement mechanisms and fines in shaping driver behavior has been extensively explored in transportation and urban planning research. These measures are pivotal as primary deterrents against illegal parking, influencing both compliance and broader urban traffic management. As Inci emphasizes, the level of enforcement is a critical determinant of a parking policy's success [3]. Similarly, Kim and Wang highlight the role of police enforcement in significantly reducing parking violations [95].

Sinning et al. explored the effects of deterrence (enforcement) and non-deterrence (social norms) letters on enhancing the collection of traffic and parking fines [37]. Their findings revealed that enforcement-focused messages were more effective in improving compliance rates than those based on social norms . This is particularly effective among habitual violators and younger drivers, underscoring the importance of tailoring communication strategies to specific driver profiles.

Alho et al. explored demand management strategies for freight parking, highlighting the role of infrastructural adjustments such as increasing parking availability and centralized management systems [96]. These measures complement enforcement efforts and effectively reduce violations but necessitate careful cost-benefit analyses to prevent unintended operational costs.

Researches in [75] and [97] emphasize that drivers often weigh the cost of fines against the convenience of illegal parking, particularly in high-demand scenarios. This highlights the need for context-sensitive fine structures that balance deterrence with socio-economic considerations.

Hagen et al. [98] and Albergaria and Favero, [99] presented the potential of leveraging technology, such as video surveillance and geospatial data, for real-time enforcement. These advancements enable targeted identification of violation hotspots, optimizing patrol routes while reducing operational costs.

Kladeftiras and Antoniou [48] examined the traffic and environmental effects of illegal parking through simulation-based methods, demonstrating that stricter enforcement significantly improves traffic flow, reduces delays, and lowers emissions. This evidence highlights the dual benefits of enforcement, addressing both operational efficiency and sustainability.

Nourinejad and Roorda [72] introduced an equilibrium model to evaluate the effects of enforcement policies on commercial vehicle parking. In their model, the equilibrium is characterized by a search and match function or bilateral meeting [100, 101] between inspection units and illegally parked vehicles. While increased enforcement helps reduce illegal parking, it also raises delivery costs, which can indirectly impact social welfare. These findings highlight the importance of carefully calibrating fines and enforcement intensity to achieve balanced outcomes, particularly in areas with high commercial activity. Building on this work, Nourinejad et al. [73] later developed an analytical model that explicitly accounts for both commercial and private vehicles, offering deeper insights into how different enforcement policies influence illegal parking behavior.

Lei et al. [86] developed a game-theoretic framework to model interactions between parking agencies and drivers. Their study demonstrates that optimizing patrol routes both spatially and temporally increases the perceived likelihood of detection, thereby enhancing compliance. Furthermore, implementing multiple-ticket policies effectively reduces prolonged violations with minimal enforcement effort.

Parking enforcement involves officers patrolling various areas to identify and ticket illegally parked vehicles. This task can be modeled as a variant of routing problem, where the goal is to design the patrol routes to maximize efficiency and revenue.

The field of police patrolling, which shares similarities with parking enforcement, has seen significant advancements in recent years. Samanta et al. [59] provide a comprehensive review of police patrolling problems, emphasizing the application of operations research techniques to optimize route design, resource allocation, and district planning. In the context of private security, Vidigal et al. [87] propose a multi-period orienteering model to optimize guard patrols, maximizing the total score collected while adhering to time constraints.

Applying routing problem concepts to parking enforcement offers a promising approach to enhance efficiency and effectiveness. Several studies have explored this domain. Shao et al. [57] focus on traveling officers problem, incorporating spatiotemporal probability models to predict violations and design efficient patrol routes. Their model aims to find an optimal patrol route that maximizes the probability of catching illegally parked cars, subject to a time constraint. Their approach leverages real-world data and optimization techniques to adapt to dynamic urban environments. Quin et al. [94] introduce the multiple traveling officers problem, extending the traveling officer problem to multiple officers. They employ meta-heuristics like cuckoo search, genetic algorithms, and particle swarm optimization to optimize officer routes and maximize fine collection.

Summerfield et al. [54] utilize the chinese postman problem to model parking enforcement as a revenue collection activity. Bruglieri [55] introduces the parking warden tour problem, which focuses on maximizing revenue while accounting for temporal decay. The profit from an inspection depends on the time since the last visit. Each warden's tour, a cycle where vertices and edges may repeat, must stay within the duty time and aim to maximize parking fine revenue. Revenue for a street drops to zero immediately after a visit, then increases linearly to a peak value by the turnover time. Ferreira et al. [58] address the parking enforcement routing problem, focusing on maximizing enforcement efficiency through criticality (benefit) metrics. The latter two studies employ mixed-integer linear programming (MILP) and heuristic methods to solve real-world instances.

These studies collectively demonstrate the potential of routing optimization techniques to improve parking enforcement strategies. However, the impact of enforcement planning, particularly through patrol routing, on driver behavior is critical but remains insufficiently studied. Parking enforcement agencies design patrol schedules and routing plans with the understanding that these strategies influence drivers' decisions at parking lots. Effective enforcement planning must account for the interplay between routing decisions and driver behavior to maximize compliance and minimize violations.

Lie et al. [86] study the interaction between patrol frequency and drivers' payment decisions using a game-theoretic framework. They cast the agency problem as a periodic vehicle routing problem with service choice over a spatial distribution of parking lots, while modeling drivers' decisions via a newsvendor formulation that minimizes total parking cost (fees plus expected violation cost as a function of patrol frequency). A closed-form driver response to each feasible patrol frequency allows the agency to embed driver behavior directly into its routing model. Despite these advances, the approach abstracts from important realism: it downplays geographic and temporal heterogeneity and assumes isolated drivers whose risk depends only on aggregate enforcement. In practice, the driver–enforcement system is many-to-many: higher inbound flow and crowding at a lot can dilute per-driver detection (e.g., occlusion, queueing for inspection), lowering citation probabilities at a fixed enforcement level and inducing endogenous behavioral feedbacks.

Building on these insights, we develop an operational model that integrates routing with an equilibrium-based driver–enforcement component [73]. This captures congestion dependent detection efficiency where rising illegal demand can reduce effective inspection per violator and feeds that back into route design. To tackle the resulting computational complexity, we propose two Kernel-Search-based matheuristics. The framework closes key gaps in prior work by jointly optimizing routes and enforcement intensity under behaviorally consistent driver response, offering a practical tool for designing more effective enforcement strategies.

4.3 Theoretical background

Figure 4.1 presents a coupled system of drivers and enforcement. The patrol-planning model (which selects routes and visit counts) is part of this system and can shift its equilibrium. When enforcement is intensified, the meeting rate rises directly. At the same time, that higher presence also changes the chance of receiving a citation and the likelihood that drivers choose to park illegally. Drivers then adjust their behavior, which feeds back into the system and indirectly alters the meeting rate. In the diagram, these outputs are broadcast to multiple steps in parallel rather than sending the flow down a single conditional path.

We assume vehicles arrive at a parking location at a total rate of T vehicles per hour, distributed between legal (T^n) and illegal (T^v) parking. The probability of choosing illegal parking, β, which is referred to as the non-compliance ratio, is dependent on the violator's expected utility. Employing a logit function, this probability is defined using the utilities of legal (U^n) and illegal (U^v) parking, with a Gumbel-distributed error term (θ) introducing behavioral randomness, a standard assumption in discrete choice models that yields the closed-form logit expression and is widely adopted in the transportation and enforcement literature to capture unobserved preference heterogeneity:

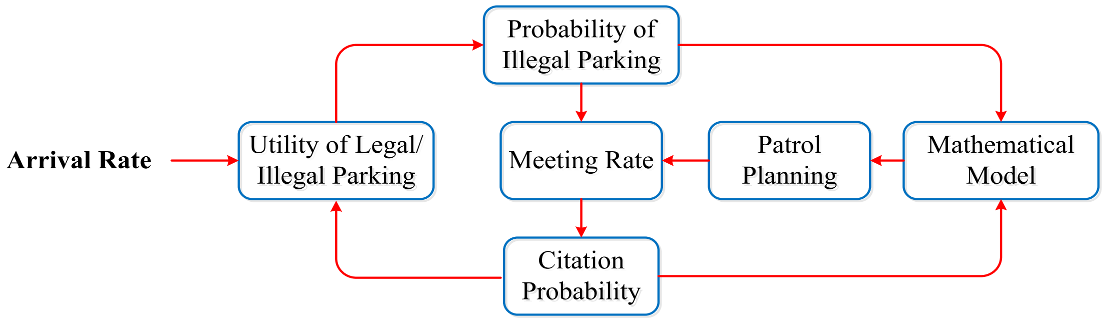

Figure 4.1 Overall relationship between different elements of enforcement planning in this paper. The figure represents information flowing to several steps in parallel, not a decision tree.

$$T^v = T\beta \tag{4.1}$$

$$T^n = T(1-\beta) \tag{4.2}$$

$$\beta = \frac{\exp(\theta U^v)}{\exp(\theta U^v) + \exp(\theta U^n)} \tag{4.3}$$

A sensitivity analysis on the dispersion parameter (θ) allows us to capture a spectrum of driver behaviors, ranging from purely deterministic ($\theta \to \infty$) to entirely random ($\theta = 0$).

The utility of parking is computed as the difference between the benefit obtained from parking and its associated costs. We define $S(l)$ as the non-negative marginal benefits a driver receives at the l^{th} hour, assuming these benefits diminish over time. Legal vehicles park for l^n hours and incur a cost of π dollars per hour, totaling $\pi \cdot l^n$ dollars. They also face a parking search cost, which we model as increasing linearly with the volume of legal arrivals at a rate of η. Thus, the search cost is $\eta \cdot T^n$, where η is a fixed parameter. Consequently, the utility for legal parking is:

$$U^n(l^n) = \int_0^{l^n} S(l)\, dl - \pi l^n - \eta T^n \tag{4.4}$$

Conversely, violators only incur a fine (φ) if they are cited, with the citation probability (α) determined by the inspection dynamics. Therefore, the utility for illegal parking is:

$$U^v(l^v) = \int_0^{l^v} S(l)\, dl - \alpha\varphi \tag{4.5}$$

The expected number of violators under steady state ($\mathcal{N}$) is linked to the arrival rate (T^v) and their average parking duration (l^v) via Little's law, a fundamental steady-state relationship widely used in queueing and service systems:

$$\mathcal{N} = T^v l^v \tag{4.6}$$

The successful identification of a violator by an officer is termed a meeting, which results in issuing a citation to the illegal vehicle. The meeting rate m (citations per hour) is governed by a meeting function $M(\cdot)$ that models the interaction between the number of violators N^v and the intensity of officer presence κ. We model the meeting rate using a Cobb–Douglas function, a standard functional form for capturing such interactions:

$$m = M(\mathcal{N}, \kappa) = A_0 \mathcal{N}^{\gamma_1} \kappa^{\gamma_2} \tag{4.7}$$

Here, $A_0 > 0$ is a meeting-efficiency parameter that accounts for spatial and operational factors. The exponents γ_1 and γ_2 represent the elasticities of the meeting rate with respect to the number of violators and the officer presence intensity, respectively. Given the meeting rate and a pool of T^v violators over the same time unit, the citation probability for a single violator is:

$$\alpha \;=\; \frac{m}{T^v} \tag{4.8}$$

A central decision for the agency is determining the optimal frequency with which officers should patrol a given parking area to inspect vehicles and enforce rules over the planning horizon. This operational decision is subject to natural constraints and directly influences both the agency's objectives and driver behavior, as drivers adjust their parking choices in response to perceived enforcement levels. In our setting, enforcement is structured as a series of periodic patrols where officers follow defined routes to inspect multiple parking lots. We now define the agency's objective function.

Given the probability of illegal parking, and the detection probability for an illegally parked vehicle, the agency's total revenue can be formulated as an objective function that maximizes total income by jointly accounting for revenues from legal parking fees and citation fines. Let $\beta(\kappa)$ be the share of drivers who park illegally under enforcement level κ, and $\alpha(\kappa)$ be the citation probability for an illegal parker. Increased enforcement typically raises $\alpha(\kappa)$ while decreasing $\beta(\kappa)$. Consequently, fee revenue scales with the legal share $1-\beta(\kappa)$, and citation revenue scales with the penalized share $\beta(\kappa)\alpha(\kappa)$. For an arrival rate T, fine amount φ, and parking fee π, the total income per period is expressed as:

$$\mathcal{R}(\kappa) = T\Big[\varphi\,\alpha(\kappa)\,\beta(\kappa) + \pi\Big(1-\beta(\kappa)\Big)\Big]. \tag{4.9}$$

Because both $\alpha(\kappa)$ and $\beta(\kappa)$ are non-linear in κ, their interaction induces multiplicative effects that may create multiple local optima, resulting in a computationally challenging optimization problem. This structure endogenously models the feedback loop between enforcement intensity and driver response within operational constraints, forming a comprehensive basis for optimizing parking enforcement planning.

Fig 4.2 (a) plots total revenue against patrol intensity κ for various arrival rates T. Revenue often follows a hump-shaped curve: as enforcement intensifies, citation income (proportional to $\beta(\kappa)\alpha(\kappa)$) initially rises because $\alpha(\kappa)$ increases faster than $\beta(\kappa)$ declines. Past a certain threshold, the reduction in $\beta(\kappa)$ dominates, causing citation income and total revenue to decrease. Fig 4.2 (b) illustrates the underlying mechanism: higher κ increases citation probability $\alpha(\kappa)$ while reducing the illegal-parking share $\beta(\kappa)$, and the balance between these two depends on the demand environment.

The revenue pattern is sensitive to T due to the multiplicative interaction between $\alpha(\kappa)$

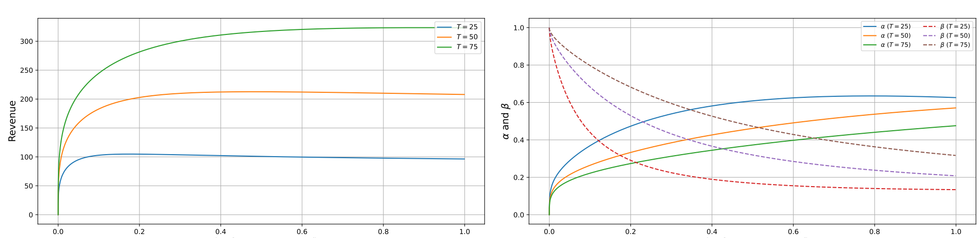

(a) (b)

Figure 4.2 An example illustrating the change in total revenue and detection probability and illegal parking rate based on varying arrival rates at a parking lot and different levels of enforcement fines.

and $\beta(\kappa)$. At lower arrival rates (e.g., $T = 25$), the marginal increase in fee revenue from improved compliance is small, while the rapid reduction in violators erodes citation income; thus, higher enforcement might reduce total revenue. Conversely, at higher arrivals (e.g., $T = 75$), the larger pool of potential violators allows revenue to increase with enforcement over a broader range before the deterrence effect becomes dominant. This demand sensitivity also explains the distinct "crisscross" patterns in Fig 4.2 (b): the point where $\alpha(\kappa)$ surpasses $\beta(\kappa)$, and the slopes of each curve, shift with T, reflecting the equilibrium feedback among congestion, detection risk, and behavioral adaptation.

From a management perspective, these results highlight a Laffer-type trade-off: under low demand, excessive patrol intensity can be fiscally counterproductive, whereas under high demand, increased enforcement can boost both revenue and compliance up to an optimal point. The optimal κ is therefore demand-dependent and requires joint calibration with fines and parking fees.

4.4 Model description

The parking enforcement agency is responsible for managing an urban service area that comprises multiple spatially dispersed parking lots. Fig 4.3 provides an illustrative example of an off-street parking network. The agency's goal is to devise optimal patrol plans for enforcement officers over a given planning horizon (P), such such as a day or a week, which is divided into a set of shorter time periods (p). Each officer operates a deterministic patrol route that must start and end at a designated depot, subject to a maximum route duration constraint $T_{\max}$ for each period. Along the route, each parking lot can be visited by at most one officer, who conducts an inspection for a fixed duration t^{in} before proceeding to the next assigned location.

Drivers are assumed to be unaware of the precise inspection schedule but form rational expectations regarding the average officer presence at a given lot over the planning horizon. This perceived enforcement level dictates their parking decisions and overall compliance behavior. Since the inspection time t^{in} is assumed constant, the total number of inspections performed at a parking lot over the planning horizon is the key factor determining the level of enforcement. As enforcement intensity κ changes, the violation probability $\beta(\kappa)$ and the detection probability $\alpha(\kappa)$ adjust accordingly. Together, these probabilities determine the agency's total income, which combines revenues from legal parking fees and citation fines.

Because these relationships are non-linear and interdependent, determining the optimal inspection level necessitates balancing operational efficiency with the resulting behavioral re-

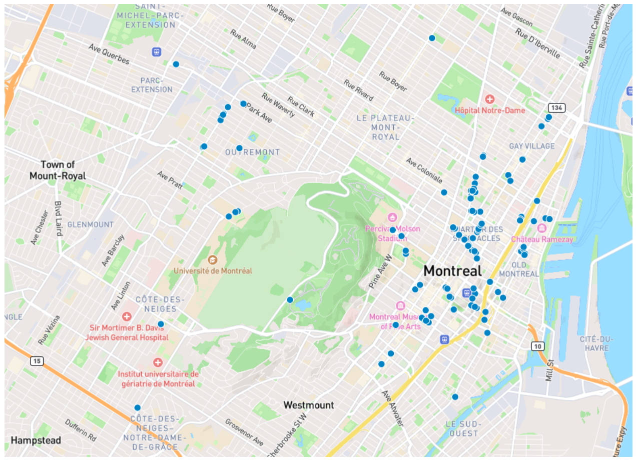

Figure 4.3 Spatial distribution of off-street parking lots in Montreal

sponse. Specifically, the decision on how frequently to inspect each parking lot directly affects the equilibrium between driver compliance and enforcement, creating a coupled problem involving both behavioral equilibrium and operational routing decisions throughout the planning horizon.

In this framework, the citation probability α and non-compliance ratio β are treated as equilibrium-derived parameters that depend on the enforcement strategy. We define the average enforcement intensity κ_i at lot i based on the total number of inspections, c_i. The optimal inspection count c_i for each parking lot emerges endogenously from the routing solution, linking the spatial patrol design with behavioral outcomes. This integrated approach allows the agency to identify not only efficient patrol routes but also the most effective enforcement intensity levels to maximize total revenue.

The total number of inspections at lot i over the planning horizon P is given by c_i. We define the average enforcement intensity κ_i (used in Equation 7) as the total inspection time divided by the length of the planning horizon, $T_{\text{horizon}} = |P| \cdot T_{\max}$:

$$\kappa_i = \frac{c_i \cdot t_i^{\text{in}}}{T_{\text{horizon}}} \tag{4.10}$$

Since t_i^{in} and T_{horizon} are constants, the continuous enforcement intensity κ_i is directly proportional to the discrete inspection count c_i. Consequently, the equilibrium outputs are functions of c_i: $\alpha_i(c_i)$ and $\beta_i(c_i)$.

The problem is defined on a complete directed graph $G = (V, A)$, where:

- $V = \{0, 1, \ldots, N^v, N^v + 1\}$ is the set of nodes. Nodes 0 and $N^v + 1$ represent the depot, while nodes $i = 1, \ldots, N^v$ correspond to the parking lots. In addition, $V' = V \setminus \{0, N^v + 1\}$ represents the set of parking lots.
- $A = \{(i, j) \mid i \neq j,\ i, j \in V\}$ is the set of arcs connecting node pairs.
- $P = \{1, 2, ..., N^p\}$ is the set of time periods, where N^p is the total number of periods.
- $O = \{1, 2, ..., N^o\}$ is the set of officers, where N^o is the total number of officers.

Each parking lot i has a non-negative arrival rate T_i. The fine for citations is φ, and the parking fee is π_i. Patrol routes must commence at node 0 and conclude at node $N^v + 1$. The travel time between nodes i and j is t_{ij}^{tr}. Each inspection requires t_i^{in} time. Patrol paths are constrained by a maximum time limit $T_{\max}$, and each parking lot can be visited at most once per period ($p \in P$).

The set of possible inspection counts is defined as $C = \{0, 1, \ldots, |P|\}$, where c represents the total number of times parking lot i is inspected over the planning horizon. The equilibrium outputs, the citation probability $\alpha_i(c)$ and the non-compliance ratio $\beta_i(c)$, are determined for each parking lot i by fixing the arrival rate T_i and the inspection count c. This count c establishes the enforcement intensity κ_i (via Equation 4.10). The resulting system of equations (1-8) is solved using a fixed-point algorithm to find the equilibrium solution. For an extensive treatment of this iterative solution procedure, please refer to [73].

Decision Variables

y_{ipr} : binary variable, equal to 1 if parking lot $i \in V$ is inspected by officer $r \in O$ in period $p \in P$, 0 otherwise.

x_{ijpr} : binary variable, equal to 1 if arc $(i, j) \in A$ is traversed in period $p \in P$ by officer $r \in O$, 0 otherwise.

z_{ic} : binary variable, equal to 1 if parking lot $i \in V$ is inspected a total of $c \in C$ times over the planning horizon, 0 otherwise.

u_{ipr}: continuous sequencing variable for $i \in V'$ (non-depot lots), $p \in P$, $r \in O$; encodes the visit order of lot i on officer r's route in period p.

Objective Function

The agency's objective is to maximize total revenue by combining revenue from legal parking fees and citation fines across all parking lots, using the equilibrium parameters $\alpha_i(c_i)$ and

$\beta_i(c_i)$ which depend on the inspection count c_i:

$$\text{Maximize} \quad \mathcal{R} = \sum_{i=1}^{N^v} T_i \Big[\varphi \cdot \alpha_i(c_i) \cdot \beta_i(c_i) + \pi_i \cdot \Big(1 - \beta_i(c_i)\Big)\Big] \tag{4.11}$$

The optimization problem is formulated as:

$$\text{Maximize:} \quad \mathcal{R} = \sum_{i \in V'} \sum_{c \in C} T_i \Big(\varphi\, \alpha_i(c)\, \beta_i(c) + \pi_i \Big(1 - \beta_i(c)\Big)\Big) z_{ic} \tag{4.12}$$

Constraints:

$$\sum_{r \in O} y_{ipr} \leq 1, \qquad \forall i \in V',\ \forall p \in P \tag{4.13}$$

$$\sum_{j \in V'} x_{0jpr} = 1, \qquad \forall p \in P,\ \forall r \in O \tag{4.14}$$

$$\sum_{i \in V'} x_{i\,(N^v+1)\,pr} = 1, \qquad \forall p \in P,\ \forall r \in O \tag{4.15}$$

$$\sum_{j \in V} x_{jipr} = \sum_{j \in V} x_{ijpr}, \qquad \forall i \in V',\ \forall p \in P,\ \forall r \in O \tag{4.16}$$

$$\sum_{j \in V} x_{jipr} = y_{ipr}, \qquad \forall i \in V',\ \forall p \in P,\ \forall r \in O \tag{4.17}$$

$$\sum_{(i,j) \in A} t_{ij}^{tr}\, x_{ijpr} \; + \; \sum_{i \in V'} t_i^{in}\, y_{ipr} \leq T_{\max}, \qquad \forall p \in P,\ \forall r \in O \tag{4.18}$$

$$u_{ipr} - u_{jpr} + N^v x_{ijpr} \leq N^v - 1 + (N^v + 1)\,(1 - y_{ipr}) \qquad \forall i, j \in V',\ i \neq j,\ \forall p \in P,\ \forall r \in O \tag{4.19}$$

$$\sum_{c \in C} z_{ic} = 1, \qquad \forall i \in V' \tag{4.20}$$

$$\sum_{c \in C} c\, z_{ic} = \sum_{p \in P} \sum_{r \in O} y_{ipr}, \qquad \forall i \in V' \tag{4.21}$$

$$\sum_{r=i+1}^{N^o} y_{ipr} = 0, \qquad \forall i \in V', \ \forall p \in P \tag{SB 1}$$

$$y_{ipr} \leq \sum_{1 \leq j \leq i-1} y_{j\,p\,r-1}, \qquad \forall i \in V', \ \forall r \in O \setminus \{1\}, \ \forall p \in P \tag{SB 2}$$

$$y_{ipr}, \ x_{ijpr}, \ z_{ic} \in \{0, 1\}, \qquad \forall i, j \in V, \ \forall p \in P, \ \forall r \in O, \ \forall c \in C \tag{4.22}$$

The optimization problem aims to maximize the total agency revenue ($\mathcal{R}$), as defined in Equation (4.12). This revenue is calculated by summing the income from all parking lots ($i \in V'$). For each lot, the revenue combines two components: the fine revenue from illegal parking, $\varphi \cdot \alpha_i(c) \cdot \beta_i(c) \cdot T_i$, and the fee revenue from legal parking, $\pi_i \cdot (1 - \beta_i(c)) \cdot T_i$. Crucially, the non-compliance ratio (β) and citation probability (α) are treated as functions of the total inspection count c, which is linked to the binary decision variable z_{ic}.

The core operational constraints manage the routing of officers. Constraint (4.13) enforces the rule that a single parking lot cannot be visited by more than one officer in any given period p. Constraints (4.14) and (4.15) define the route structure, ensuring every active officer's path must originate at the starting depot (node 0) and conclude at the ending depot (node $N^v + 1$) in each period p.

The flow constraints ensure route coherence. Constraint (4.16) maintains flow conservation requiring that if an officer enters a parking lot i, they must also leave it. Constraint (4.17) links the arc variables (x_{ijpr}) to the inspection variables (y_{ipr}), ensuring that flow enters lot i if and only if that lot is scheduled for inspection. The operational limitation on time is enforced by Constraint (4.18), which sums the total travel time along the route and the fixed inspection time for all visited lots, ensuring this sum does not exceed the maximum route duration, $T_{\max}$, for any officer in any period.

The subtour elimination constraint, Constraint (4.19), uses a modified MTZ formulation involving auxiliary variables u_{ipr} to prevent the formation of isolated cycles (subtours) that do not include the depot, thereby ensuring connected routes. This specific form is structured to apply the ordering only among nodes that are actively visited ($y_{ipr} = 1$). The subsequent constraints link the routing decisions to the equilibrium behavior. Constraint (4.20) ensures that each parking lot i is assigned exactly one total inspection count c over the entire planning horizon. Constraint (4.21) is the crucial link, equating the total number of physical inspections performed ($\sum_{p \in P} \sum_{r \in O} y_{ipr}$) to the assigned inspection level c via the binary variable z_{ic}.

Finally, the symmetry-breaking constraints are introduced to enhance computational efficiency. Constraint (SB 1) attempts to enforce an ordering on officer usage by linking their activity to the index of the lot they visit. Constraint (SB 2) enforces a hierarchical ordering by ensuring that if officer r inspects a lot i, officer $r - 1$ must inspect at least one lot with a smaller index $j < i$, thereby reducing the number of redundant, symmetric solutions. The remaining constraints define the necessary binary and non-negative conditions for all decision variables [102].

4.5 Solution approach

Our patrol planning and enforcement problem is a generalized form of the Selective Traveling Salesman Problem, which is known to be NP-hard [103]. Given the combinatorial complexity arising from the integrated routing and behavioral equilibrium decisions, solving the full mixed-integer linear programming (MILP) model to optimality for realistic instances is computationally prohibitive. Consequently, we employ the Kernel Search (KS) algorithm, a powerful matheuristic approach, to efficiently generate high-quality, near-optimal solutions.

Kernel Search is a matheuristic framework designed to find heuristic, and potentially optimal, solutions for large-scale optimization problems formulated as MILPs by exploiting the power of commercial solvers [104]. First introduced by [105] to address multi-dimensional knapsack problems, the KS algorithm operates by strategically simplifying the complex variable space. This is achieved by initially identifying a promising subset of variables called the kernel, often utilizing information derived from the linear programming (LP) relaxation of the full problem. Variables not included in the kernel are then grouped into structured buckets based on criteria such as their reduced costs.

The method proceeds in two main phases. The initialization phase uses the LP relaxation to determine the composition of the starting kernel and organizes the remaining variables into distinct buckets. This is followed by the improvement phase, which iteratively refines the solution by solving a series of restricted MILP subproblems. In each iteration of the improvement phase, the current kernel is combined with one bucket at a time, allowing the MILP solver to explore a manageable portion of the solution space. Promising variables identified during these subproblem solutions are then incorporated back into the kernel for subsequent iterations, driving the algorithm toward an enhanced final result [105].

KS has been successfully applied to a range of combinatorial optimization problems, including the machine consolidation problem [106], the minimum spanning tree problem [107], scheduling problems [108], and facility location problems [109]. Additionally, variations of KS have been effectively utilized in routing problems, such as the team orienteering problem [110], the inventory routing problem [111, 112], the capacitated vehicle routing problem [113], and the periodic vehicle routing problem [114]. A common issue with KS implementation for routing problems in aforementioned studies is that the LP relaxation cannot provide useful information to identify the promising variables.

We propose a novel approach to leverage linear programming information for bucket construction. Our observations reveal that the optimal LP solution in one-shot procedure does not provide sufficient guidance for forming a sequence of buckets. Considering that our problem, classified as a multi-period team orienteering problem, is inherently graph-based, we utilize the underlying graph structure from the relaxed LP solution to define the buckets. Specifically, for each x_{ijkr} with a positive value in the current LP solution, we remove the corresponding edge (i, j) from the network and assign it to the first bucket. The LP is then resolved on the modified network to generate subsequent buckets. This iterative process continues until an adequate number of buckets is obtained.

As noted in [112], we need to ensure connectivity within the graph. To achieve this, we follow two key steps. First, we introduce edges connecting each node directly to the depot, guaranteeing a feasible connection between all nodes and the depot. Next, for each node j, we evaluate its potential connectivity to other nodes i in the network based on a scoring mechanism. The $Score_{ij}$ is calculated using the formula:

$$Score_{ij} = \frac{T_i\Big[\varphi \cdot \alpha_i(1) \cdot \beta_i(1) + \pi_i \cdot \Big(1 - \beta_i(1)\Big)\Big]}{t_{ij}^{tr}}$$

The $Score_{ij}$ effectively represents the expected income or resource value associated with connecting nodes j and i, normalized by the travel time. For each node j, we identify and select a edges with the highest scores, where a is a predefined parameter (we consider $a = 2$ in our implementation). This ensures that the most valuable connections, in terms of potential income or resources, are prioritized for inclusion in the network.

4.5.1 Algorithm Enhancement

During the algorithmic process, specifically before reconstructing the next kernel (denoted as bucket i), we observed that in some cases, the officer route is not inefficient in terms of time (Fig. 4.4, route 2) since the underlying graph is not complete. Although time is not explicitly an objective in our model, minimizing the time taken by an officer's route is advantageous. A shorter route allows the officer to be available for additional tasks and, from an optimization perspective, might yield better results as the algorithm progresses and considers subsequent buckets.

To address this, after solving the restricted MIP, we propose enhancing the obtained route using the 2-OPT algorithm as a post-optimization step. Specifically, for each route, we apply the 2-OPT algorithm to explore possible improvements and select the route that minimizes travel time (Fig. 4.5, Step 1). Notably, this refinement does not affect revenue, as it solely optimizes the time aspect of the route.

This improvement introduces an additional opportunity: if the revised enforcement route results in extra time within the current period, it may be possible to add unvisited parking lots to the current route, provided that the route adheres to predefined time constraints (Fig. 4.5, Step 2). For instance, if there are four periods and parking lot i remains uninspected in the current period, it can potentially be included in the route, provided that the tour time constraint allows it. However, this does not necessarily guarantee an increase in the agency's income, as illustrated in Fig. 4.2. Therefore, when adding a new unvisited parking lot to the current enforcement route, we also retain the edges from the step 1. This allows the next iteration of the optimization model to evaluate all options and identify the most effective choice, ensuring improved results.

4.5.2 Proposed Algorithm Framework

Algorithm 1 outlines the general framework of our proposed algorithm. The algorithm is divided into two main phases: the *Initialization Phase* and the *Improvement Phase*. Below, we describe each phase in detail.

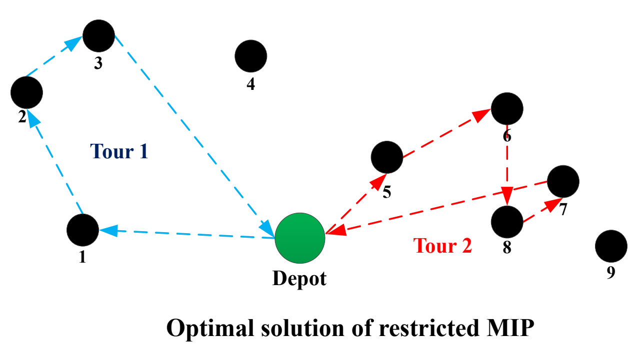

Figure 4.4 An example of an optimal routing solution for a specific period obtained after solving the restricted MIP (kernel)

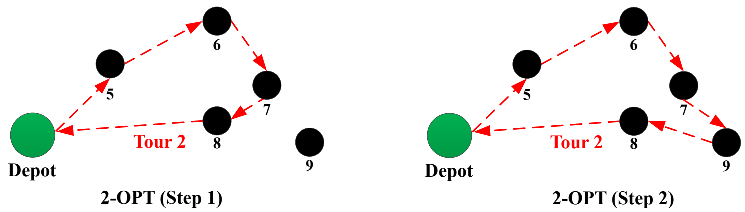

Figure 4.5 A 2-OPT improvement heuristic is applied to the optimal Restricted MILP solution

Initialization Phase

In the initialization phase, we first solve the relaxed mixed-integer problem (RMIP) of the proposed model to create buckets. During each iteration of this phase, we identify routing variables x_{ijpr} that take positive values in the RMIP. These variables correspond to the selected edges (i, j) in the network. After identifying and saving this set of promising edges, we enforce their removal from the feasible region of the RMIP for the subsequent iteration's solution. This mechanism ensures that the search explores a new, distinct set of feasible solutions (edges) when evaluating the next bucket, preventing repetition.

Next, we re-optimize the model for a predefined number of iterations (I_{max}). To maintain connectivity, specific edges are introduced back into the network E_{depot}. The edges removed during the iterations form the initial bucket structure, where the first bucket is treated as the *initial kernel* B_1, and the subsequent buckets B_{iter} are preserved for later use. This approach allows us to gradually refine the solution, ensuring that important routes are not arbitrarily eliminated while maintaining connectivity across the network.

Improvement Phase

In the improvement phase, we start with the initial kernel generated during the initialization phase. A mixed-integer model is then constructed using the edges in the kernel, along with edges selected based on a scoring mechanism, E_R, as previously described, and E_{depot}. The inclusion of a scoring mechanism ensures that only the most promising edges are reintroduced into the model, possibly improving computational efficiency. This model is solved to obtain an optimized solution. Starting from the second iteration, we also employ warm-start techniques to accelerate the solve by seeding the optimizer with the incumbent from the previous iteration. Next, the Rerouting and Enhancement algorithm (RRE), detailed in Algorithm 2, is applied to further refine the solution. This additional refinement step allows for local adjustments that *may* enhance the overall routing efficiency by addressing potential inefficiencies in the initial solution. The caveat inherent in the word *may* is critical: although

RRE aims to optimize individual patrol routes by integrating new stops, this operational improvement does not guarantee a global increase in the objective function. This is due to the non-linear nature of the revenue function: a local increase in the number of visits raises the enforcement intensity (κ), which, past the optimal threshold, can cause a sharp drop in the illegal parking share (β). Thus, locally efficient rerouting might push the system into a less profitable region of the revenue curve, reducing the global total income. Accordingly, in the final step of Algorithm 2 we take the union of both outcomes, the refined routes from Step 1 (2-OPT) and the optimized routes from Step 2, and pass this combined set to the global routes optimizer in the next KS stage. The resulting edge set defines E_{iter}, which serves as the input kernel for the subsequent iteration, where the next bucket $B_{\text{iter}+1}$ is incorporated. By iterating this procedure, the method systematically broadens the search around high-quality incumbents while preserving locally improved structures, thereby incrementally enhancing the routing plan over a diverse family of candidate solutions.

Algorithm 1 KS Algorithm for Enforcement Routing Problem

Input: Graph $G = (V, E)$, depot d, maximum iterations $I_{\max}$
Output: Optimized routing solution

Phase 1: Initialization
iter $\leftarrow 0$
$E_{\text{depot}} \leftarrow \{(d, i), (i, d) \mid i \in V\}$
while iter $< I_{\max}$ **do**
  $x^*_{ijpr} \leftarrow \arg\max_{x_{ijpr}} f(x_{ijpr})$ ▷ Solve RMIP model
  $B_{\text{iter}} \leftarrow \{(i, j) \mid \exists\, p, r \text{ s.t. } x^*_{ijpr} > 0\}$ ▷ Bucket from used edges
  $E \leftarrow E \setminus B_{\text{iter}}$ ▷ Remove edges from RMIP solution
  $E \leftarrow E \cup E_{\text{depot}}$ ▷ Add depot connections
  iter $\leftarrow$ iter $+ 1$
end while
$E_{\text{R}} \leftarrow \bigcup_{i \in V} \underset{J \subseteq V, |J|=2}{\text{argmax}} \sum_{j \in J} \text{Score}_{ij}$
Phase 2: Improvement
$E_0 \leftarrow \emptyset$, iter $\leftarrow 0$
while iter $< I_{\max}$ **do**
  $G_{\text{iter}} \leftarrow (V,\ E_{\text{iter-1}} \cup B_{\text{iter}} \cup E_{\text{R}} \cup E_{\text{depot}})$
  Solve MIP on G_{iter} with a time limit
  Pass incumbent (x^*, y^*) to Algorithm 2 (RRE) for route refinement
  $G_{\text{union}} \leftarrow \text{RRE}(G_{\text{iter}}, x^*, y^*)$
  $E_{\text{iter}} \leftarrow E(G_{\text{union}})$ ▷ Edges returned by RRE
  iter $\leftarrow$ iter $+ 1$
end while
Combined Buckets
Solve final MIP on stored edges $S \leftarrow \bigcup_{\text{iter}} E_{\text{iter}}$ with a time limit
return Optimal routing solution

Algorithm 2 Rerouting and Enhancement (RRE algorithm)

Input: Graph $G = (V, E)$; incumbent solution (x^*, y^*)
Output: Optimized union graph G_p^{union} for each period p

for all $p \in P$ **do**
  $UV_p \leftarrow \{\, i \in V \mid \sum_{r \in O} y^*_{ipr} = 0 \,\}$ ▷ Unvisited lots in period p
  $C_p \leftarrow \{\, C \subseteq E \mid C$ is a simple route induced by x^* in period $p \,\}$
  $C_p^{2\text{opt}} \leftarrow \emptyset$ ▷ Routes after 2-OPT
  $C_p^* \leftarrow \emptyset$ ▷ Optimized route set
  $G_p^{\text{union}} \leftarrow \emptyset$
  for all $c \in C_p$ **do**
    $c^{2\text{opt}} \leftarrow \text{TwoOPT}(c)$ ▷ Apply Algorithm 3 on c
    $C_p^{2\text{opt}} \leftarrow C_p^{2\text{opt}} \cup \{c^{2\text{opt}}\}$ ▷ Step 1
    $c^* \leftarrow c^{2\text{opt}}$ ▷ Start insertion from 2-OPT route
    for all $k \in UV_p$ **do**
      found $\leftarrow$ **false**
      for all $(i, j) \in c^*$ **do**
        $c' \leftarrow \Big(c^* \setminus \{(i, j)\}\Big) \cup \{(i, k), (k, j)\}$ ▷ Insert k into c^*
        if c' is feasible **then**
          $c^* \leftarrow c'$
          $UV_p \leftarrow UV_p \setminus \{k\}$
          found $\leftarrow$ **true**
          break ▷ Exit edge loop; check next lot
        end if
      end for
      if found then
        continue
      end if
    end for
    $C_p^* \leftarrow C_p^* \cup \{c^*\}$ ▷ Step 2
  end for
  $G_p^{\text{union}} \leftarrow \bigcup_{c \in C_p^{2\text{opt}} \cup C_p^*} c$ ▷ Union of both route families
end for
return G_p^{union}

Algorithm 3 2-Opt Algorithm for Route Re-optimization

Input: C : Initial route
Output: Optimized route C^*

$N \leftarrow |C|$ ▷ Number of nodes in the route
$f_s \leftarrow f(C)$ ▷ Objective value of current best
$T \leftarrow C$ ▷ Temporary solution
ImprovementFound $\leftarrow$ **true**
while *ImprovementFound* **do**
    ImprovementFound $\leftarrow$ **false**
    for $i \leftarrow 1$ **to** $N-1$ **do**
        for $j \leftarrow i+1$ **to** N **do**
            $T \leftarrow \text{Swap}(T, i, j)$ ▷ Swap node i with node j in the route
            $f_T \leftarrow f(T)$
            if $f_T < f_s$ **then**
                $C \leftarrow T; \quad f_s \leftarrow f_T$
                ImprovementFound $\leftarrow$ **true**
                break
            end if
        end for
    end for
end while
return C

4.6 Computational Study

This section presents a detailed computational study to evaluate the performance of the proposed routing-based parking enforcement model. The study focuses on analyzing the model's effectiveness under varying parameters and testing its efficiency using the proposed algorithm. The objectives are: (i) to assess the impact of routing and patrol decisions under different parameter settings, (ii) to analyze the computational approach and, (iii) to apply proposed model in real application. All computations were performed using Gurobi Optimizer 12.0.0 on a system equipped with an Intel® Core™ i7-7800X CPU @ 3.50GHz (6 physical cores, 12 logical processors), running AlmaLinux 9.5, and the implementation was carried out in Python 3.9.

4.6.1 Model analysis[1]

[1]Parameter settings used in the numerical experiments: $B_0 = 40.0$, $B_1 = 0.3$, $A_0 = 2.0$, $\gamma_1 = 0.6$, $\gamma_2 = 0.3$, $\eta = 0.02$. The marginal benefit function is $S(l) = \int_0^l B_0 \, B_1^x \, dx$.

We test the integrated model on a network of ten nodes, where nodes 0 and 9 are depots and the remaining eight nodes ($L_1, \ldots, L_8$) are parking lots. Depots are fixed at $(50, 50)$ and lots are drawn uniformly in $[0, 100] \times [0, 100]$. Lot-specific arrival rates T_i are sampled uniformly from $[40, 100]$. Each period has a maximum operational time $T_{\max} = 200$, and the inspection time is $t_i^{\text{in}} = 60$ for all lots. The planning horizon includes $|P| = 5$ periods, so the per-lot inspection count is bounded above by 5. The base parking fee is $\pi = 2$, and the fine φ takes values in $\{8, 10, 12, 14\}$.

To examine resource and behavioral effects, we run experiments for $N^v \in \{2, 3\}$ officers and $\theta \in \{0.5, 0.7\}$, where larger θ makes choices more sensitive to utility differences. Table 4.1 reports, for each lot i, the constrained optimal inspection count c_i^{OPT} found by the integrated (routing+equilibrium) model; the value in parentheses is the unconstrained per-lot maximizer ignoring routing and resource limits. The last row shows total revenue for each (φ, θ, N^v) configuration.

Table 4.1 Optimal enforcement levels c_i and revenue comparison ($N^v = 2$ vs. $N^v = 3$).

		$N^v = 2$								$N^v = 3$							
		$\theta = 0.5$				$\theta = 0.7$				$\theta = 0.5$				$\theta = 0.7$			
Lot	T_i	φ=8	φ=10	φ=12	φ=14	φ=8	φ=10	φ=12	φ=14	φ=8	φ=10	φ=12	φ=14	φ=8	φ=10	φ=12	φ=14
L1	82	2 (5)	2 (5)	2 (4)	2 (3)	3 (5)	2 (5)	2 (3)	2 (2)	4 (5)	4 (5)	3 (4)	3 (3)	4 (5)	4 (5)	3 (3)	2 (2)
L2	68	2 (5)	2 (5)	2 (3)	2 (2)	2 (5)	2 (4)	2 (2)	1 (1)	3 (5)	3 (5)	3 (3)	2 (2)	3 (5)	3 (4)	2 (2)	1 (1)
L3	89	2 (5)	2 (5)	2 (5)	3 (3)	1 (5)	2 (5)	2 (4)	2 (2)	3 (5)	3 (5)	4 (5)	3 (3)	3 (5)	3 (5)	4 (4)	2 (2)
L4	61	2 (5)	2 (5)	2 (3)	2 (2)	3 (5)	2 (4)	2 (2)	1 (1)	3 (5)	3 (5)	3 (3)	2 (2)	3 (5)	3 (4)	2 (2)	1 (1)
L5	46	1 (5)	1 (3)	1 (2)	1 (1)	1 (5)	1 (2)	1 (1)	1 (1)	1 (5)	1 (3)	1 (2)	1 (1)	1 (5)	1 (2)	1 (1)	1 (1)
L6	87	3 (5)	3 (5)	3 (5)	2 (3)	3 (5)	3 (5)	3 (4)	2 (2)	4 (5)	4 (5)	4 (5)	3 (3)	4 (5)	4 (5)	4 (4)	2 (2)
L7	56	2 (5)	2 (4)	2 (2)	2 (2)	2 (5)	2 (3)	2 (2)	1 (1)	3 (5)	3 (4)	2 (2)	2 (2)	3 (5)	3 (3)	2 (2)	1 (1)
L8	77	3 (5)	3 (5)	3 (4)	2 (2)	3 (5)	3 (5)	3 (3)	2 (2)	5 (5)	5 (5)	4 (4)	2 (2)	5 (5)	5 (5)	3 (3)	2 (2)
Revenue:		1857.0	2140.0	2308.2	2361.0	1870.7	2177.8	2338.5	2362.1	2018.1	2262.3	2358.4	2365.2	2045.4	2300.3	2366.5	2362.1

Note. Each cell reports c_i^{OPT} and, in parentheses, the unconstrained per-lot maximizer (i.e., the inspection level that maximizes the lot's revenue when routing and resource constraints are ignored).

The results reveal three patterns. First, operational scarcity pushes the solution toward coverage rather than concentration: with $N^v = 2$, constrained counts c_i^{OPT} lie mostly in $\{1, 2, 3\}$ even where the unconstrained maximizers are 3–5. Second, adding an officer ($N^v = 3$) relaxes routing limits and allows high-demand lots to approach their unconstrained targets (e.g., L_8 reaches 5(5) at $\varphi \in \{8, 10\}$), increasing total revenue albeit with diminishing marginal gains at higher φ. Third, increasing θ produces small but consistent improvements because choices react more strongly to utility differences; at very high fines, however, the system exhibits mild saturation where additional deterrence yields little or no extra revenue.

Sensitivity to $T_{\max}$ and fine level (φ)

We now examine how revenue responds jointly to the fine φ and the maximum operational time per period $T_{\max}$ with $\theta = 0.5$. Table 4.2 summarizes total revenue for $\varphi \in \{8, 9, 10, 11, 12, 13, 14\}$ across several values of $T_{\max}$.

Table 4.2 Total revenue sensitivity to $T_{\max}$ and φ at $\theta = 0.5$.

	Fine level φ						
$T_{\max}$	8	9	10	11	12	13	14
225	1859	2015	2146	2247	2317	2356	2368
250	1891	2044	2168	2261	2323	2359	2367
275	1923	2072	2190	2279	2336	2363	2367
300	1929	2082	2204	2293	2347	2368	2368
310	1916	2070	2194	2285	2342	2367	2368
320	1904	2058	2183	2277	2337	2366	2367
325	1897	2052	2178	2273	2335	2365	2369
350	1919	2069	2190	2280	2337	2365	2369

Two mechanisms explain the patterns. On the one hand, increasing $T_{\max}$ enlarges feasible tour time, which tends to raise revenue by enabling additional coverage when resources are tight or fines are modest. On the other hand, the same increase dilutes enforcement intensity because, for fixed c_i, the denominator in (4.10) grows with $T_{\max}$ (via $T_{\text{horizon}} = |P|\, T_{\max}$), thereby reducing κ_i. The first mechanism acts like adding resources; the second reduces the effectiveness of each inspection. Revenue therefore need not be monotone in $T_{\max}$.

At low fines ($\varphi \in \{8, 9\}$), the resource effect dominates. Revenue rises when $T_{\max}$ increases from 225 to about 300 because longer routes permit more lots to be visited. At high fines ($\varphi \in \{13, 14\}$), the dilution effect and behavioral saturation dominate. Revenues become nearly flat across $T_{\max}$, reflecting that optimal κ_i is already small and pushing c_i higher risks over-deterrence with little incremental gain. In short, $T_{\max}$ has a twofold role, expanding operational capacity while simultaneously attenuating enforcement intensity, and the net effect depends on φ and demand.

4.6.2 Algorithm Performance Evaluation

In this subsection, we evaluate the performance of the two proposed algorithms in identifying optimal or near-optimal solutions. To achieve this, we generated test instances with 20, 25, and 35 parking lots, varying the combinations of periods and officers. Each instance is labeled as $N/O/P$, where N represents the number of parking lots, O indicates the number of officers, and P specifies the number of periods. For example, instance 25−3−4 comprises 25 parking lots, 3 officers, and 4 periods. Additionally, we considered varying $T_{\max}$ 75, 100, and 150 to analyze the algorithms under diverse conditions. It is worth noting that a runtime of 60 seconds was allocated for each KS iteration during the improvement phase, while 120 seconds were assigned for the combined buckets stage.

Table 4.3 Obtained results of proposed algorithm on medium size instances

$T_{\max}$	N/O/P	MIP	Time (S)	KS1	Time (S)	Iter	KS2	Time (S)	Iter	Imp1 %	Imp2 %	Gap
75	20-3-3	1997.19	14.75	1997.19	37.21	1	1997.19	39.42	1	0.00	0.00	0.0%
	20-3-4	2076.30	33.45	2076.30	13.21	1	2076.30	14.83	1	0.00	0.00	0.0%
	20-4-4	2144.00	64.56	2144.00	16.65	1	2144.00	20.84	1	0.00	0.00	0.0%
	25-3-3	2544.14	29.79	2544.14	8.07	1	2544.14	9.71	1	0.00	0.00	0.0%
	25-3-4	2647.26	75.12	2647.26	14.42	1	2647.26	13.66	1	0.00	0.00	0.0%
	25-4-4	2750.18	83.72	2750.18	17.70	1	2750.18	24.70	1	0.00	0.00	0.0%
	35-3-3	2894.09	390.16	2894.09	618.42	5	2894.09	626.81	5	0.00	0.00	0.0%
	35-3-4	3047.98	3000	3047.98	666.83	5	3047.98	669.65	7	0.00	0.00	0.5%
	35-4-4	3150.57	3000	3150.57	622.58	6	3150.57	658.77	6	0.00	0.00	0.9%
100	20-3-3	2426.19	810.91	2426.19	458.62	3	2426.19	518.43	2	0.00	0.00	0.0%
	20-3-4	2546.78	3000	2546.78	608.04	2	2546.78	607.25	2	0.00	0.00	0.9%
	20-4-4	2652.25	2873.04	2645.35	617.11	7	2652.25	626.17	6	-0.26	0.00	0.0%
	25-3-3	3081.82	3000	3034.97	522.06	6	3034.97	559.91	5	-1.52	-1.52	20.8%
	25-3-4	3297.13	3000	3297.13	611.90	3	3297.13	617.66	2	0.00	0.00	25.4%
	25-4-4	3469.07	3000	3469.07	611.34	3	3469.07	614.22	2	0.00	0.00	29.7%
	35-3-3	3523.54	3000	3523.54	712.66	6	3523.54	622.34	7	0.00	0.00	46.4%
	35-3-4	3790.71	3000	3790.71	612.62	5	3790.71	614.22	5	0.00	0.00	43.5%
	35-4-4	3947.66	3000	3954.35	607.15	3	3961.75	613.62	3	0.17	0.36	47.4%
150	20-3-3	3150.03	3000	3156.18	604.67	1	3156.18	607.23	1	0.20	0.20	6.6%
	20-3-4	3315.08	3000	3323.85	721.77	3	3325.78	721.76	7	0.26	0.32	6.8%
	20-4-4	3478.22	3000	3475.48	722.25	4	3475.48	722.35	4	-0.08	-0.08	2.7%
	25-3-3	4159.60	3000	4152.52	722.09	5	4160.66	722.17	6	-0.17	0.03	7.4%
	25-3-4	4395.96	3000	4410.00	722.55	4	4410.00	722.63	7	0.32	0.32	7.7%
	25-4-4	4605.71	3000	4648.02	723.35	4	4648.02	723.08	7	0.92	0.92	8.3%
	35-3-3	4758.32	3000	4864.71	719.29	7	4900.43	729.06	6	2.24	2.99	22.0%
	35-3-4	5365.27	3000	5362.99	712.66	6	5385.57	724.30	7	-0.04	0.38	14.3%
	35-4-4	5736.18	3000	5782.86	727.26	3	5782.86	727.34	3	0.81	0.81	12.5%

The parameter settings for the tests were configured with $A_0 = 1.25$, $\theta = 1$ and $t_i^{\text{in}} = 0.2 \cdot T_i$. Table 4.3 presents a summary of the experimental results. The column *MIP* represents the solution value obtained by Gurobi, either as the optimal solution or the best solution identified within 3000 seconds. The column *Time* indicates the runtime of the algorithm in seconds. The column *KS1* shows the solution achieved by the KS algorithm without the use of the 2-OPT algorithm, while *KS2* reflects the solution obtained using the 2-OPT search. Columns *Imp1%* and *Imp2%* indicate the percentage differences between the KS algorithms and Gurobi, with positive values signifying that the KS algorithm outperformed Gurobi, and negative values indicating the opposite. The column *Iter* displays the iteration at which the algorithm first discovers the (near-) optimal solution. Finally, the column *Gap* provides the optimality gap for the solution obtained by Gurobi.

The results show that increasing the period length, number of periods, or the number of officers generally improves the revenue across the tested instances. This aligns with the intuition that more patrol resources and more inspection possibility allow for more extensive enforcement coverage.

Out of all the tested instances, Gurobi successfully solved only 9 optimally. The performance of the KS algorithms (KS1 and KS2) demonstrated their ability to identify high-quality solutions within a reasonable runtime. Notably, KS2 consistently outperformed KS1 due to its enhanced heuristic procedures, and in several cases, it closely matched the solutions obtained by Gurobi and in 9 cases it was able to report better solutions.

The computational experiments aim to evaluate the performance of the proposed KS algorithms (KS1 and KS2) compared to the MIP approach across multiple problem instances. These experiments were conducted to assess the algorithms' efficiency, scalability, and ability to provide high-quality solutions within a reasonable computational time.

We also compare the performance of the algorithms on larger-sized problems. To achieve this, we evaluated instances with 50, 75, and 100 nodes under different $T_{\max}$ (150 and 200). The results, presented in Table 4.4, indicate that both KS algorithms outperform the MIP solver while requiring significantly less computational time. Gurobi was allotted a 3,000-second runtime budget, whereas KS terminated in at most $\sim$720 seconds.

The results demonstrate that the KS algorithms, particularly KS2, consistently outperformed Gurobi in terms of scalability and solution quality, especially for larger and more complex problem instances. While the MIP solver struggled with increasing problem size and exhibited high optimality gaps (up to 57.80% for 100-4-4 with 150), KS2 reliably provided superior solutions with much smaller gaps and significant computational efficiency. For example, in the 100-4-4 instance ($T_{\max} = 200$), KS2 achieved a revenue of 15,637.59, significantly surpassing MIP's 12,722.34 while requiring a fraction of the computational resources. This trend highlights the limitations of MIP solvers like Gurobi in handling large-scale combinatorial optimization problems compared to the efficiency and practicality of KS algorithms. Furthermore, KS2 consistently delivered higher revenue than KS1, proving its robustness and reliability for solving complex scenarios. The results show that KS2 not only matched but often surpassed MIP's performance by leveraging heuristic-driven efficiency without compromising solution quality. These findings emphasize that KS algorithms, particularly KS2, are superior alternatives to classical MIP solvers for real-world applications requiring high-quality

Table 4.4 Obtained results of proposed algorithm on large instances

$T_{\max}$	**N/O/P**	**MIP**	**Gap**	**KS1**	**KS2**	**Imp1%**	**Imp2%**
150	50-3-3	6252.65	28.07%	6293.02	6362.91	0.65	1.76
	50-3-4	6751.38	25.47%	7013.14	7065.21	3.88	4.65
	50-4-4	7453.39	20.00%	7726.49	7759.31	3.66	4.10
	75-3-3	8272.90	36.45%	8762.76	8762.76	5.92	5.92
	75-3-4	9390.79	37.59%	9795.83	9824.49	4.31	4.62
	75-4-4	10167.42	33.90%	10897.31	10861.24	7.18	6.82
	100-3-3	9573.86	34.80%	10020.12	10103.43	4.66	5.53
	100-3-4	9660.11	56.90%	11152.00	11277.88	15.44	16.75
	100-4-4	10605.57	57.80%	12572.85	12705.26	18.55	19.80
200	50-3-3	7523.26	12.60%	7708.74	7744.40	2.47	2.94
	50-3-4	8069.84	10.60%	8299.09	8308.32	2.84	2.96
	50-4-4	8516.75	10.40%	8776.78	8796.84	3.05	3.29
	75-3-3	10067.44	28.34%	10423.33	10599.89	3.54	5.29
	75-3-4	11149.64	21.90%	11974.29	11992.92	7.40	7.56
	75-4-4	11092.21	29.80%	13186.92	13138.31	18.88	18.45
	100-3-3	11252.18	34.00%	12073.31	12088.35	7.30	7.43
	100-3-4	12184.05	37.30%	13635.26	13705.50	11.91	12.49
	100-4-4	12722.34	38.60%	15662.08	15637.60	23.11	22.91

solutions within practical timeframes.

It is important to highlight that while some of the solutions exhibit large optimality gaps, they can still be considered high-quality solutions for the problem. Our experimental observations indicate that the upper bound in the Branch-and-Cut algorithm tends to progress slowly toward the lower bound (MIP solution), limiting its ability to close the optimality gap effectively.

This phenomenon is not uncommon in routing problems. For instance, in [115], it was reported that for medium-sized problem instances, despite an average optimality gap of approximately 68%, the obtained solutions were still considered acceptable. A similar result can be observed in [111], where cases with 10 retailers had optimality gaps ranging from 2.8% to 50.1%.

4.6.3 Application

In this section, we present an application of the proposed model for optimizing parking lot inspections in Montreal. Data related to parking fees, the number of spaces available, and the locations of parking lots were obtained from the website of the Agence de Mobilité Durable de Montréal. For simplicity, we grouped close parking lots as a single lot by summing their spaces. Fig. 5.2 illustrates the locations of these 32 parking lots along with one central depot, highlighted in red.

To calculate travel times between lots, we utilized the Google Maps API in driving mode to estimate travel durations. For arrival rates, we defined them as a proportion of the available parking spaces. Specifically:

- For parking lots with fewer than 40 spaces, the arrival rate was equal to the number of spaces.
- For lots with 40 to 80 spaces, the arrival rate was calculated as $40 + 0.5 \cdot (\text{spaces} - 40)$.
- For lots with more than 80 spaces, the arrival rate was determined as $60 + 0.25 \cdot (\text{spaces} - 80)$.

Additionally, we considered the total ticket value to be \$91 and planned inspections across five weekdays, with two inspection periods per day and five officers. Each inspection period lasted 240 minutes, resulting in a total of 10 periods over the planning horizon. The calibration follows the model analysis; only two parameters were updated for the case study: the

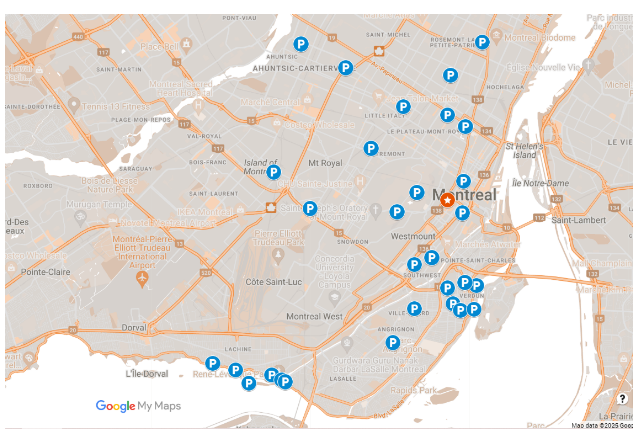

Figure 4.6 Location of Parking Lots in the case study

dispersion parameter was set to $\theta = 0.2$ and the meeting-rate scale to $A_0 = 0.7$. Inspection time per lot was scaled by relative demand,

$$t_i^{\text{in}} = 30 + 30\,\frac{T_i - \min(T)}{\max(T) - \min(T)}\ \text{(minutes)},$$

so larger and busier lots receive longer inspections. All other elements of the specification, including the marginal benefit function $S(l)$ and its coefficients, as well as the remaining behavioral and cost parameters, remain as in the model analysis.

Table 4.5 Revenue achieved by Gurobi and proposed heuristics in the case study scenario

Method	**Revenue**	**Optimality Gap**	**Time (s)**	**Improvement (%)**
Gurobi	20,519.52	4.81%	3,000.00	–
KS 1	20,655.01	–	738.08	0.66
KS 2	20,810.41	–	741.81	1.42

With a 3,000 s time limit, Gurobi obtained a revenue of 20,519.52 with a 4.81% optimality gap. KS1 reached 20,655.01 in 738.08 s (about 4× faster) for a 0.66% improvement over Gurobi, while KS2 achieved the best revenue, 20,810.41, in 741.81 s (also about 4× faster), a 1.42% improvement. These results show that both heuristics deliver higher-quality solutions in substantially less time than the exact solver, with KS2 performing best overall.

4.7 Conclusion

This research presents a significant advancement in parking enforcement planning by incorporating driver behavior and enforcement dynamics into a comprehensive decision framework. The integration of these factors allows for a more realistic and effective approach to optimizing patrol routes. The proposed model, based on equilibrium principles and incorporating a novel Kernel Search algorithm, offers several key contributions:

- **Comprehensive Modeling**: The model captures the complex interplay between driver behavior, enforcement strategies, and human resource constraints, providing a more realistic representation of the parking enforcement process.
- **Enhanced Decision Support**: The framework provides parking enforcement agencies with a valuable tool to optimize patrol routes, allocate resources effectively, and enhance compliance with parking regulations.

- **Innovative Solution Approach**: The customized Kernel Search algorithm, incorporating graph-based bucket construction and a 2-OPT post-optimization step, effectively addresses the computational challenges associated with the complex model.
- **Policy Implications**: The findings of this research have important policy implications, informing the development of more effective and efficient parking enforcement strategies that balance the needs of drivers, businesses, and the overall urban environment.

This research provides a foundation for future work in this area. Potential avenues for further research include:

- **Incorporating real-time data**: Integrating real-time data on traffic flow, parking occupancy, and violation occurrences could further enhance the model's accuracy and responsiveness.
- **Developing dynamic models**: Extending the model to incorporate dynamic factors, such as time-varying demand and changing traffic patterns, can provide more accurate predictions and improve the adaptability of enforcement strategies.

By addressing these areas, future research can further refine and enhance parking enforcement strategies, leading to more efficient urban transportation systems and improved quality of life for city residents.

CHAPTER 5 ARTICLE 2: PARKING ENFORCEMENT PLANNING WITH BEHAVIOR-RESPONSIVE PATROL ROUTING THROUGH A MULTI-VISIT ORIENTEERING PROBLEM

Authors: Mohsen Yahyaei, Michel Gendreau and Nikolaj Van Omme, Submitted on December 20, 2025, Under review at Transportation Science

Abstract: This study addresses a variant of the Multi-Visit Team Orienteering Problem (MVTOP) arising in urban parking enforcement, where multiple officers must inspect parking lots under both operational constraints and behavioral equilibrium conditions. We introduce mathematical models that incorporate recovery times between successive visits, maximum tour duration limits for officers, and the endogenous response of drivers to enforcement intensity. To tackle the inherent complexity of this setting, we design a Variable Neighborhood Descent (VND) algorithm enhanced with a temporal conflict-resolution repair procedure to guarantee feasibility. The framework is further strengthened through the integration of a knapsack-based relaxation, which provides adaptive guidance for neighborhood exploration.

To benchmark solution quality, we compare the proposed algorithm against the commercial solver Gurobi as well as classical VND approaches. Computational experiments show that our adaptive method consistently delivers superior efficiency and solution quality. Finally, the framework is validated through a real-world case study using off-street parking data from Montreal, demonstrating its ability to generate operationally feasible and practically relevant patrol schedules.

Keywords: Parking Enforcement Planning, Routing Problem, Variable Neighborhood Descent Heuristic, Multi-Visit Team Orienteering

5.1 Introduction

Urban environments are increasingly burdened by illegal parking, which exacerbates congestion, blocks traffic flow, and heightens driver frustration on already crowded streets [116]. Parking enforcement plays a vital role in regulating driver behavior, maintaining order, and optimizing public space usage. Yet, the execution of enforcement strategies faces multiple operational constraints—limited officer availability, budgetary limits, and the dynamic behavioral adaptations of drivers in response to enforcement activities. Designing effective patrol schedules requires more than simple routing optimization; it calls for a nuanced understanding of how enforcement policies influence driver compliance.

Drivers often compare parking fees with the expected cost of a fine [117]. When enforcement appears sparse or inconsistent, they might choose to park illegally despite the convenience offered by digital payment systems. The willingness to pay for adequate parking time is strongly influenced by both the perceived likelihood of being penalized and the severity of the fine for non-compliance [118]. For example, Petiot [75] showed that under high congestion, increased fine levels may paradoxically encourage non-compliance by creating conditions where risk-tolerant drivers benefit from reduced competition for parking spots. Similarly, Nourinejad and Roorda [72], and later Nourinejad et al. [73], modeled illegal parking behaviors using search-friction and equilibrium frameworks, emphasizing how citation likelihood and enforcement effort influence dwell time and compliance, particularly among commercial vehicle operators.

While these behavioral models shed light on driver decision-making, translating these insights into actionable patrol strategies across large networks remains an open challenge. The allocation of officers across large urban networks, and the integration of enforcement behavior into routing policies, remain less thoroughly examined. Although some works acknowledge that areas with higher violation frequencies merit more enforcement attention, a rigorous operational framework for determining how officers should be deployed across multiple parking zones is still lacking. In real-world contexts, limited staffing and complex urban layouts demand intelligent, scalable patrol planning strategies that balance coverage, revenue collection, and behavioral effectiveness.

To this end, several studies have attempted to formalize enforcement routing problems. Summerfield et al. [54] modeled enforcement as a variant of the Chinese postman problem, optimizing officer routes over street networks with embedded revenue and waiting options. Ferreira et al. [58] formulated the parking enforcement arc routing problem using a mixed-integer linear programming (MILP) framework that maximizes inspection benefits. Bruglieri et al. [55] introduced a time-dependent model where enforcement value diminishes immediately after a visit and recovers over time—closely mirroring the diminishing marginal return of frequent revisits. Shao et al. [57] and Qin Et al. [94] leveraged real-world sensor data to construct (multi) traveling officer problems, incorporating spatial and temporal considerations for single and multi-officer deployment.

Despite significant progress in the literature, most enforcement planning models continue to frame the problem as a static operational task, largely ignoring the adaptive responses of drivers to observed or expected enforcement activities. An important exception is the study in [86], which proposes a bi-level game-theoretic framework between drivers and inspection agency. In their model, drivers decide on their parking duration based on the perceived risk of receiving a fine, while the enforcement agency optimizes its patrol routes using a variant of the periodic vehicle routing problem (PVRP). The PVRP seeks to establish visit schedules and route plans across a multi-period planning horizon to optimize a specific objective function, such as travel cost minimization or profit maximization [119, 120]. However, in the context of enforcement planning, the model in [86] adopts simplified assumptions about driver behavior through closed-form probability functions and relies on PVRP formulations that may overlook the complex nonlinear nature of illegal parking patterns in urban environments. More recently, Yahyaei et al. [121] incorporated driver behavior and enforcement dynamics into a PVRP framework by leveraging the model of [73]. Their objective was to optimize daily enforcement operations to maximize revenue from legal parking and citations issued to violators, and they introduced a kernel-based heuristic [105] to address the problem. Nonetheless, both PVRP-based models are constrained by the assumption that enforcement actions can occur at most once per period, limiting their flexibility in capturing more dynamic enforcement strategies required in real-world urban settings.

To address these limitations, we argue that routing models based on the multi-visit routing problem (MVRP) offer a more practical and operationally aligned alternative to PVRP for enforcement planning. MVRP frameworks allow multiple visits to the same node, reflecting the real-world necessity of revisiting parking zones to deter and penalize ongoing violations. For example, Hanafi et al. [110] and Jung et al. [122] developed the multi-visit team orienteering problem (MVTOP) as subclass of the MVRP with precedence constraints, highlighting the model's capacity to capture repeated planned visits. However, existing MVTOP approaches often neglect critical behavioral dynamics, such as the ineffectiveness of back-to-back inspections. In practical parking operations, revisiting a location shortly after a prior inspection is unlikely to detect new violations, as any previously parked vehicles have already been cited. Consequently, a temporal spacing mechanism, referred to as *recovery time* between successive visits, is necessary to ensure the operational and behavioral validity of multi-visit strategies. Incorporating such recovery dynamics enables a more faithful representation of enforcement realities and enhances the effectiveness of route planning under strategic driver responses.

From a computational standpoint, the NP-hard nature of orienteering problems [123, 124] necessitates efficient solution strategies. A wide range of heuristics and metaheuristics have been developed to address (team) orienteering problems variants, including Adaptive Large Neighborhood Search [125], Simulated Annealing with Reinforcement Learning [126], Multi-level Memetic Search [127], GRASP-ILS [128], and Tabu Search [129].

Among these, Variable Neighborhood Search (VNS) [130] and its deterministic variant, Variable Neighborhood Descent (VND), are particularly appealing due to their systematic exploration of neighborhood structures and proven success in route-based optimization problems with complex constraints [131].

Research Gap and Contributions

Previous studies in parking enforcement planning with behavior responsive have primarily relied on PVRP formulations, which assume that each location can be visited at most once per time period. While such approaches offer a degree of scheduling flexibility, they fall short in capturing the operational reality of urban enforcement, where multiple inspections of the same location, appropriately spaced over time, can be necessary to maintain deterrence. In contrast, our study introduces a multi-visit routing framework specifically tailored for enforcement applications. Although multi-visit routing has been explored in the broader vehicle routing literature, to the best of our knowledge, it has not been applied within enforcement planning nor integrated with behavioral models that capture the response of drivers to inspection timing and frequency.

The distinguishing feature of our approach lies in its synthesis of multi-visit orienteering with a behaviorally grounded equilibrium framework. A key component is the incorporation of recovery time constraints, which govern the effectiveness of repeated inspections and capture the temporal sensitivity of deterrence. This formulation allows for a more realistic and dynamic representation of the interaction between inspection schedules and illegal parking behavior. To solve the problem, we propose a tailored solution methodology that couples a fixed-point algorithm for computing the equilibrium in driver responses with a VND metaheuristic for patrol routing.

By integrating behavioral equilibrium with routing optimization, our approach bridges the gap between strategic behavior modeling and operational enforcement planning. This makes the model particularly relevant for urban policy design, enabling municipalities to develop more adaptive, targeted, and resource-efficient enforcement strategies that align with driver behavior and compliance dynamics.

The structure of the paper is as follows. Section 5.2 introduces the parking enforcement equilibrium model. Section 5.3 outlines the problem setting and presents its mathematical formulation. Section 5.4 details the proposed solution methodology. Section 5.5 reports the results of numerical experiments used to evaluate the algorithm's effectiveness, along with a comparative analysis. Finally, Section 5.6 concludes the paper and discusses potential directions for future research.

5.2 Equilibrium Framework and Structural Interactions

In this section, we adopt the equilibrium model of illegal parking behavior originally introduced in [73], which is presented for localized and spatially disaggregated enforcement settings. We consider a stylized urban environment in which arriving drivers choose between legal and illegal parking options based on expected utility. The total vehicle inflow is denoted by Λ, measured per hour (this can be equivalently scaled to any other time interval). It is partitioned into compliant parkers Λ^c and violators Λ^v, where:

$$\Lambda = \Lambda^c + \Lambda^v. \tag{5.1}$$

Letting β represent the fraction of violators, we have $\Lambda^v = \beta\Lambda$ and $\Lambda^c = (1-\beta)\Lambda$. Drivers make their parking decisions by comparing the expected utility of each option. A probabilistic logit model governs the share of violators:

$$\beta = \frac{e^{\phi V^v}}{e^{\phi V^v} + e^{\phi V^c}}, \tag{5.2}$$

where V^c and V^v denote the deterministic utilities of legal and illegal parking, respectively, and ϕ is a scale parameter reflecting sensitivity to utility differences. When ϕ approaches zero, the choice probability converges to random selection. Under such conditions, the system may settle into an intermediate level of compliance.

Legal parkers pay a fee and incur a congestion-dependent search cost. Their utility is:

$$V^c(d^c) = \int_0^{d^c} r(t)\,\mathrm{d}t - \rho d^c - \zeta\Lambda^c, \tag{5.3}$$

where d^c is the legal dwell time, $r(t)$ is the marginal benefit of parking (a decreasing function), ρ is the hourly parking fee, and ζ is the search cost coefficient.

In contrast, violators do not pay for parking but risk receiving a citation. Their expected utility is:

$$V^v(d^v) = \int_0^{d^v} r(t)\,\mathrm{d}t - \pi\delta, \tag{5.4}$$

where d^v is the illegal dwell time, π is the detection probability, and δ is the fine.

Detection is modeled through a Cobb-Douglas *meeting rate* function that captures encounters between illegally parked vehicles and patrol officers. Let $\mathcal{N}^v$ be the average number of illegally parked vehicles and k represent the patrol intensity (e.g., number of inspections per hour or spatial density of patrols). The detection mechanism is given by:

$$\mu(\mathcal{N}^v, k) = \xi_0(\mathcal{N}^v)^{\lambda_1}(k)^{\lambda_2}, \tag{5.5}$$

where parameters λ_1, λ_2 denote the elasticities with respect to the number of illegal vehicles and the number of enforcement units, respectively. The detection probability becomes:

$$\pi = \frac{\mu(\mathcal{N}^v, k)}{\Lambda^v}. \tag{5.6}$$

Assuming steady-state dynamics, Little's Law, a classical result in queueing theory, gives $\mathcal{N}^v = \Lambda^v d^v$.

It should be noted that this modeling structure allows us to incorporate localized enforcement intensity and its influence on driver behavior. In urban contexts, patrol intensity varies across regions and time periods, requiring planners to allocate enforcement resources across multiple areas. This introduces spatial heterogeneity into the equilibrium system, as each region may have different demand levels, parking availability, and enforcement sensitivity.

The integration of behavioral modeling and operational planning is essential for policy relevance. A common agency's objective is to maximize total income, which combines revenue from legal parking fees and penalties collected from detected violations. This formulation links user behavior and enforcement decisions under practical constraints, such as limited inspection capacity.

The relationship between patrol intensity and total revenue may be non-monotonic as indicated in earlier studies [73, 121]. Increasing inspection effort does not always yield higher income due to behavioral feedback loops. Nonlinearity of the system can generate inflection points where additional enforcement becomes counterproductive or yields diminishing returns.

This equilibrium-informed framework supports more adaptive enforcement strategies by recognizing that revenue optimization is not solely a function of patrol volume, but of its calibration. Moreover, the spatial localization of patrol activity introduces a need for coordinated planning across regions. The next section builds on this foundation to integrate this behavioral model into a multi-region patrol routing and scheduling and optimization problem.

5.3 Problem Description and Mathematical Formulation

We investigate a patrol planning problem for a parking enforcement agency operating over a network of off-street urban parking facilities. The objective is to determine optimal patrol routes for a team of officers to maximize enforcement-driven revenue. This includes both legitimate parking payments and penalty revenue from detected violations. Notably, the revenue is endogenously influenced by the patrol intensity, as inspection frequencies affect driver compliance through a behavioral equilibrium.

The routing component of the problem is modeled as an extension of the MVTOP on an undirected graph $G = (V, E)$, where $V = 0, 1, \ldots, N$ denotes the set of nodes, and $E = (i, j) \in V \times V : i \neq j$ is the set of undirected edges. Node 0 and node N serve as the start and end depots for all officer routes. The remaining nodes represent parking lots that may be visited multiple times within a single planning period, subject to an upper bound (h_{max}). However, successive visits to the same lot by any officer are prohibited unless a minimum recovery time has elapsed between visits. This recovery time models the delay required for deterrence to take effect and for illegal parking activity to re-emerge, thereby capturing the temporal sensitivity of enforcement effectiveness. As a result, the problem requires careful coordination of both when and how often each location is inspected. This constraint, combined with time-limited tours and behavioral feedback, introduces significant complexity into the joint routing and scheduling process.

We now introduce the notation required to model our problem.

Sets

$V = \{0, 1, \ldots, N\}$	Set of all nodes (depots and parking lots)
$O = \{1, 2, \ldots, m\}$	Set of available enforcement officers
$H = \{1, \ldots, h_{\max}\}$	Set of inspection levels per lot

Parameters

$d_{i,j}$	Travel time between nodes i and j
t_i^{insp}	Inspection time required at node i
ρ_i	Parking fee collected at node i
δ_i	Fine imposed for illegal parking at node i
$\pi_{i,c}$	Citation probability at node i given c inspections (from equilibrium)
$\beta_{i,c}$	Illegal parking probability at node i given c inspections (from equilibrium)
t_{rec}	Minimum time interval required between two visits to the same node (recovery time)
$T_{\max}$	Maximum duration permitted for each officer route
Λ_i	Arrival rate at parking lot i

Decision Variables

$x_{i,h,j,h',o} \in \{0,1\}$	$1,$ if officer o travels from i at level h to j at level $h',\ 0$ otherwise
$y_{i,h,o} \in \{0,1\}$	$1,$ if officer o performs inspection at node i at level $h,\ 0$ otherwise
$z_{i,c} \in \{0,1\}$	$1,$ if node i is inspected exactly c times, 0 otherwise
$u_{i,h,o} \in \mathbb{Z}$	Auxiliary variable to eliminate subtours in officer o's route at i level h
$a_{i,h,o} \in \mathbb{R}_+$	Arrival time of officer o at node i during visit level h

Objective Function

The objective function is designed to maximize revenue for the agency.

$$\max \sum_{i=1}^{N-1} \Lambda_i \left[\sum_{c=0}^{h_{\max}} \delta_i \, \pi_{i,c} \, \beta_{i,c} \, z_{i,c} \ + \ \sum_{c=0}^{h_{\max}} (1 - \beta_{i,c}) \, \rho_i \, z_{i,c} \right] \tag{5.7}$$

Constraints

$$\sum_{j\in V\setminus\{N\}}\sum_{h\in H} x_{0,1,j,h,o} = 1 \qquad \forall o \in O \tag{5.8}$$

$$\sum_{i\in V\setminus\{0\}}\sum_{h\in H} x_{i,h,N,1,o} = 1 \qquad \forall o \in O \tag{5.9}$$

$$\sum_{o\in O} y_{i,h,o} \leq 1 \qquad \forall i \in V \setminus \{0, N\},\ \forall h \in H \tag{5.10}$$

$$\sum_{j\neq i}\sum_{h'\in H} x_{i,h,j,h',o} = y_{i,h,o} \qquad \forall i \in V \setminus \{0, N\},\ \forall h \in H,\ \forall o \in O \tag{5.11}$$

$$\sum_{i\neq j}\sum_{h\in H} x_{i,h,j,h',o} = y_{j,h',o} \qquad \forall j \in V \setminus \{0, N\},\ \forall h' \in H,\ \forall o \in O \tag{5.12}$$

$$a_{0,1,o} = 0 \qquad \forall o \in O \tag{5.13}$$

$$a_{N,1,o} \leq T_{\max} \qquad \forall o \in O \tag{5.14}$$

$$a_{i,h,o} + t_{rec} + t_i^{\text{insp}} \leq a_{i,h+1,o'} + M(1 - y_{i,h+1,o'}) \qquad \forall i \in V \setminus \{0, N\},\ \forall h \in H\ \forall o, o' \in O \tag{5.15}$$

$$\sum_{o\in O} y_{i,h+1,o} \leq \sum_{o'\in O} y_{i,h,o'} \qquad \forall i \in V \setminus \{0, N\},\ \forall h \in H \tag{5.16}$$

$$a_{i,h,o} + d_{i,j} + t_i^{\text{insp}} - a_{j,h',o} \leq M(1 - x_{i,h,j,h',o}) \qquad \forall i, j \in V,\ h, h' \in H,\ \forall o \in O \tag{5.17}$$

$$u_{i,h,o} + 1 - u_{j,h',o} \leq (|V| - 1)h_{max}(1 - x_{i,h,j,h',o}) \qquad \forall i \neq j,\ \forall h, h' \in H,\ \forall o \in O \tag{5.18}$$

$$\sum_{c=0}^{h_{max}} z_{i,c} = 1 \qquad \forall i \in V \setminus \{0, N\} \tag{5.19}$$

$$\sum_{c=0}^{h_{max}} c \cdot z_{i,c} = \sum_{h\in H}\sum_{o\in O} y_{i,h,o} \qquad \forall i \in V \setminus \{0, N\} \tag{5.20}$$

Variable Domains

$$x_{i,h,j,h',o} \in \{0, 1\} \qquad \forall i, j \in V,\ h, h' \in H,\ o \in O$$

$$y_{i,h,o} \in \{0, 1\} \qquad \forall i \in V,\ h \in H,\ o \in O$$

$$z_{i,c} \in \{0, 1\} \qquad \forall i \in V,\ c \in H$$

$$u_{i,h,o} \in \mathbb{Z} \qquad \forall i \in V,\ h \in H,\ o \in O$$

$$a_{i,h,o} \in \mathbb{R}_+ \qquad \forall i \in V,\ h \in H,\ o \in O$$

The objective function (7) aims to maximize total revenue, which is composed of two elements: expected citation income from violators (the first term) and legal parking payments (the second term).

Constraints (8) and (9) ensure that each officer's route starts and ends at the depot. Constraints (10) guarantee that each parking lot can be visited at most once at level h. Constraints (11) and (12) enforce flow conservation, guaranteeing that entries and exits to visited nodes are balanced within each officer's route. Constraints (13) enforce that all patrol routes are initiated at the beginning of the planning horizon, i.e., time 0. Constraints (14) define the route timing, where all patrols must finish within the time limit $T_{\max}$. Constraints (15) guarantees that sufficient time (i.e., recovery time t_{rec}) separates consecutive inspections of the same lot. Constraints (16) enforce hierarchical precedence at each node i by allowing a visit at level $h+1$ only if the node is also visited at level h. Constraints (17) ensure the feasibility of the time sequence across consecutive node visits. Constraints (18) are classical subtour elimination conditions used to prevent disconnected cycles in each officer's route. Constraints (19) ensure that each parking lot is assigned a unique inspection frequency level across the planning horizon, and Constraints (20) link these inspection frequencies with actual visits performed by officers.

5.3.1 Behavioral Linkage: Translating Discrete Visits to Continuous Intensity

The core challenge in integrating the operational patrol routing model (MVTOP) with the economic behavioral model [73] is the translation of scales. The behavioral equilibrium is defined by a continuous enforcement rate (k), while the operational output is a set of discrete, integer visits (c_i) to each lot i over the planning horizon $T_{\max}$. This translation is necessary as the routing problem delivers discrete events, while rational agents in a steady state require a continuous expectation of risk.

To establish a mathematically consistent profit function, we introduce average enforcement intensity ($\bar{k}_i$) as a strategic proxy that links the operational decisions to the drivers' perceived risk:

Definition of Average Enforcement Intensity

We define $\bar{k}_i$ as the fractional time utilization of enforcement resources at lot i over the total planning horizon $T_{\max}$. This factor represents the long-term average resource commitment, which is assumed to be the basis for rational driver expectation in a steady state: This is consistent with the core assumption that drivers are aware of the enforcement agency's operational window ($T_{\max}$) as a fixed policy variable.

$$\bar{k}_i = \frac{c_i \cdot t_i^{\text{insp}}}{T_{\max}}$$

where c_i is the total number of visits, t_i^{insp} is the fixed inspection time, and $T_{\max}$ is the total duration of the planning horizon. It should be noted that this average intensity abstracts from the minute-by-minute visit sequence, which is consistent with the model's scope under static demand and drivers' imperfect knowledge of the tactical schedule.

This average intensity $\bar{k}_i$ is substituted into the Cobb-Douglas meeting rate function to derive the citation probability π used in the behavioral equilibrium:

$$\mu_i(N^v, \bar{k}_i) = \xi_0 (N^v)^{\lambda_1} (\bar{k}_i)^{\lambda_2}$$

Structural Consequence and Methodological Defense

The chosen linkage strategy introduces a critical structural consequence: the planning horizon $T_{\max}$ functions not only as a resource constraint but also as a deterrence divisor.

This duality leads to a non-classical effect on the enforcement intensity for a fixed set of visits:

$$\uparrow T_{\max} \implies \downarrow \bar{k}_i \quad \text{(for fixed } c_i\text{)}$$

This relationship is not a model deficiency, but rather a direct and realistic representation of deterrence economics. It accurately models the dilution of deterrence: when the same quantum of physical enforcement ($c_i \cdot t_i^{\text{insp}}$) is spread over a longer period, the resulting lower intensity ($\bar{k}_i$) can correspond to a lower perceived risk (π), thus capturing the principle of diminishing marginal returns in enforcement effectiveness.

Crucially, while $\bar{k}$ decreases for a fixed set of visits, the optimal MVTOP solution will leverage the increase in $T_{\max}$ to select a greater total number of profitable visits ($\sum c_i$) whose added revenue offsets the marginal loss in deterrence efficiency, resulting in a net increase in the overall optimal objective value, consistent with resource expansion in classical optimization. The robustness and magnitude of this structural trade-off will be quantified in the sensitivity analysis of the parameter $T_{\max}$.

5.3.2 Knapsack Based Upper Bound (KUB) Model

The relationship between orienteering problems and knapsack formulations is well established in the literature, both for deriving upper bounds [103] and for guiding constructive heuristics [132]. Building on this connection, we enhance our benchmarking process by introducing a knapsack-based relaxation to obtain a tighter upper bound. In addition, we develop an adaptive solution method that leverages this relaxation to guide the search process and improve solution quality. In the knapsack relaxation, routing and sequencing decisions are omitted to yield a computationally efficient approximation. Parking lots ($\ell \in L$) are treated as knapsack items, and permitting multiple inspections corresponds to selecting multiple copies of each item, thereby capturing the revenue potential of the system without imposing route constraints.

Each item's weight is defined as its inspection time plus an access-time proxy, computed as the average of the shortest outgoing and incoming travel arcs to and from the lot:

$$t_\ell = \tfrac{1}{2} \min_{\ell' \in L \setminus \{\ell\}} \{d_{\ell\ell'}\} + \tfrac{1}{2} \min_{\ell' \in L \setminus \{\ell\}} \{d_{\ell'\ell}\} + t_\ell^{\text{insp}} = \min_{\ell' \in L \setminus \{\ell\}} \{d_{\ell\ell'}\} + t_\ell^{\text{insp}}.$$

Revenues remain identical to those in the original model. The officer working time defines the knapsack capacity, the number of officers corresponds to the number of knapsacks, and the maximum number of inspections allowed per lot restricts the number of copies of each item.

Sets

$L = \{1, \ldots, N-1\}$	Set of parking lots (items)
$O = \{1, \ldots, m\}$	Set of enforcement officers (knapsacks)
$H = \{1, \ldots, h_{\max}\}$	Allowable number of inspections per lot

Parameters

t_ℓ	Inspection time plus the average of the shortest outgoing and incoming travel to/from lot ℓ
$p_{\ell,h}$	Revenue from visiting lot ℓ exactly h times
$T_{\max}$	Officer working time (knapsack capacity)
$h_{\max}$	Maximum number of inspections allowed per lot

Decision Variables

$$x_{\ell,o,h} = \begin{cases} 1 & \text{if copy } h \text{ of lot } \ell \text{ is assigned to officer } o \\ 0 & \text{otherwise} \end{cases}$$

$$w_{\ell,h} = \begin{cases} 1 & \text{if lot } \ell \text{ is inspected } h \text{ times across all officers} \\ 0 & \text{otherwise} \end{cases}$$

Objective Function

$$\max \quad \sum_{\ell \in L} \sum_{h \in H} p_{\ell,h} \cdot w_{\ell,h}$$

Constraints

$$\sum_{o \in O} \sum_{h \in H} x_{\ell,o,h} = \sum_{h \in H} h \cdot w_{\ell,h} \qquad \forall \ell \in L \tag{5.21}$$

$$\sum_{h \in H} w_{\ell,h} \leq 1 \qquad \forall \ell \in L \tag{5.22}$$

$$\sum_{o \in O} x_{\ell,o,h} \leq 1 \qquad \forall \ell \in L,\ \forall h \in H \tag{5.23}$$

$$\sum_{\ell \in L} \sum_{h \in H} t_\ell \cdot x_{\ell,o,h} \leq T_{max} \qquad \forall o \in O \tag{5.24}$$

$$\sum_{h \in H} (t_\ell^{\text{insp}} \cdot h + t_{\text{rec}} \cdot (h-1)) \cdot w_{\ell,h} \leq T_{max} \qquad \forall \ell \in L \tag{5.25}$$

Variable Domains

$$x_{\ell,o,h} \in \{0,1\} \qquad \forall \ell \in L,\ o \in O,\ h \in H$$

$$w_{\ell,h} \in \{0,1\} \qquad \forall \ell \in L,\ h \in H$$

Constraints (21) ensures that the total number of inspections assigned to a lot across all officers matches the selected overall inspection level. Constraints (22) guarantees that each lot is inspected at most once, at a specific level h. Constraints (23) ensures that officers can inspect a copy h of lot ℓ at most once. Constraint (24) enforces that the total inspection time assigned to each officer does not exceed their maximum time budget T_{max}. Constraint (25) impose the model if a multiple copies of an item is selected, minimum time space between copes should be respected.

5.4 Methodology

The mathematical routing model outlined above relies on two key parameters: the citation probability and the rate of illegal parking. These values are determined through the solution of an equilibrium system. By solving the equilibrium for each node i under varying levels of enforcement intensity, the corresponding revenue can be calculated. The routing component then selects the optimal enforcement levels across the network by accounting for operational constraints and the revenues of other lots. In the following, we first describe the method for finding the fixed point of the equilibrium system [133], and then present our proposed approach for generating an operational enforcement plan.

5.4.1 Solution Methodology for the Parking Equilibrium

The parking choice equilibrium is modeled as a two-stage game (interested readers are referred to [73] for details). Drivers make two sequential decisions:

1. **Lower-level:** given the parking mode (legal or illegal), choose the dwell time to maximize expected utility;
2. **Upper-level:** choose between legal or illegal parking based on the utilities obtained from the lower-level problem.

This setup is solved using backward induction: first the best-response dwell times are derived, then these results feed into the parking mode-choice decision to determine the aggregate equilibrium.

Lower-level: Best-response Dwell Times

At the lower level, each driver optimizes dwell time conditional on their parking mode. Let d^c denote legal dwell time and d^v denote illegal dwell time. Drivers solve:

$$d^{m*} = \arg\max_{d^m} V^m(d^m, \cdot), \quad m \in \{c, v\},$$

where V^c and V^v are the expected utilities of legal and illegal parking, respectively.

Legal parking. The first-order condition gives:

$$\frac{\partial V^c}{\partial d^c} = 0 \quad \Rightarrow \quad r(d^{c*}) - \rho - \zeta \frac{\partial \Lambda^c}{\partial d^c} = 0 \quad \Rightarrow \quad d^{c*} = r^{-1}(\rho). \tag{5.26}$$

Hence, legal dwell time depends only on the parking fee ρ (and not on congestion from other drivers).

Illegal parking. For violators, the first-order condition yields:

$$\frac{\partial V^v}{\partial d^v} = 0 \quad \Rightarrow \quad r(d^{v*}) - \delta \frac{\partial \pi}{\partial d^v} = 0. \tag{5.27}$$

After rearranging and simplifying (see the proof provided in [73]), this becomes:

$$r(d^{v*})\, d^{v*} = \delta\, \mu(\Lambda^v, k), \tag{5.28}$$

where $\mu(\cdot)$ captures the expected fine as a function of the illegal arrival rate Λ^v and enforcement intensity k.

Thus, the best-response illegal dwell time is obtained as:

$$d^{v*} = L(\Lambda^v), \tag{5.29}$$

where $L(\cdot)$ is an implicit function solved numerically (e.g., using root-finding methods).

Upper-level: Parking Mode Choice

At the upper level, given the lower-level optimal dwell times d^{c*} and d^{v*}, drivers choose between legal and illegal parking based on expected utility. The aggregate illegal arrival rate Λ^v must be consistent with the fraction of drivers who prefer illegal parking.

This condition can be expressed using a logit-type mapping:

$$\Gamma\left(\Lambda^v\right) = \frac{\Lambda}{1 + \exp\left(\phi\left[V^c(\Lambda^v) - V^v(\Lambda^v, d^{v*})\right]\right)}, \quad \forall\, \Lambda^v \in [0, \Lambda], \tag{5.30}$$

where $V^c(\Lambda^v)$ and $V^v(\Lambda^v, d^{v*})$ are computed using the lower-level dwell times.

The equilibrium illegal arrival rate Λ^{v*} is then defined as the fixed point of this mapping:

$$\Lambda^{v*} = \Gamma(\Lambda^{v*}). \tag{5.31}$$

Intuitively, this ensures that the number of drivers choosing illegal parking, as implied by the logit choice model, is consistent with the total inflow of illegal drivers. This fixed-point formulation bridges the lower-level dwell time optimization with the upper-level mode choice decision, completing the two-stage parking choice equilibrium.

Algorithm for Computing the Equilibrium

Based on the lower- and upper-level conditions, the equilibrium is computed via an iterative fixed-point algorithm. This procedure updates illegal dwell times and arrival rates until convergence (Algorithm 4).

Algorithm 4 Compute Parking Choice Equilibrium (d^{v*}, Λ^{v*})

Initialization: Set outer iteration counter $g := 0$ and initialize number of illegal vehicles $\mathcal{N}_g^v$ randomly.
repeat
  Solve Illegal Dwell Time: Compute $d_g^v = L(\mathcal{N}_g^v)$ by solving the implicit equation

$$r(d^v)\, d^v = \delta\, \mu(\mathcal{N}_g^v, k)$$

  using a numerical root-finding method (e.g., Newton-Raphson).
  Compute Realized Rate: $\Lambda_g^v = \mathcal{N}_g^v / d_g^v$
  Inner Loop: Update Illegal Arrival Rate via Logit Model
  Set inner counter $h := 0$ and initialize trial illegal arrival rate $\hat{\Lambda}_h^v$
  repeat
    Compute utilities $V^v(\hat{\Lambda}_h^v, d_g^v)$ and $V^c(\hat{\Lambda}_h^c, d^{c*})$, with $\hat{\Lambda}_h^c = \Lambda - \hat{\Lambda}_h^v$
    Update illegal arrival rate:

$$\hat{\Lambda}_{h+1}^v = \frac{\Lambda}{1 + \exp\left(\phi\left[V^c - V^v\right]\right)}$$

    $h := h + 1$
  until $|\hat{\Lambda}_{h+1}^v - \hat{\Lambda}_h^v| \leq \epsilon$
  Update Number of Illegal Vehicles: $\mathcal{N}_{g+1}^v = \hat{\Lambda}_{h^*}^v \cdot d_g^v$
  $g := g + 1$
until $|\mathcal{N}_g^v - \mathcal{N}_{g-1}^v| \leq \epsilon$

5.4.2 Solution Methodology for optimization

Initial Solution Construction: Greedy Heuristic with Hierarchical Enforcement

In addition to the multi-visit possibility, the inclusion of a recovery time introduces bidirectional dependencies between routes, as a minimum interval t_{rec} must be respected between successive inspections of the same parking lot. As we mentioned previously, successive visits to the same lot must be sufficiently spaced in time, and successive visits that violate this spacing are disallowed. This constraint introduces a strong temporal coupling between routing and scheduling, as the feasibility of each inspection depends on the timing of previous visits, even across different officer tours. Consequently, the inspection time of a lot in one tour may be directly affected by earlier visits to the same location in another tour. This interdependence significantly increases the complexity of feasibility checking and requires coordinated scheduling across all officers to ensure compliance with recovery constraints and to avoid overlaps. It is also worth noting that when multiple visits to the same lot are assigned to a single officer within a single tour, the recovery constraint still applies; however, in this case it does not create inter-route dependencies. Such situations can be handled more straightforwardly, for example, by applying local time adjustments or shifts along the route, and are therefore not discussed in detail here.

To construct an initial solution with behaviorally motivated constraints, we implement a greedy insertion heuristic that incrementally builds officer routes based on marginal revenue gain. The insertion logic accounts for feasibility under the recovery time constraint and ensures the maximum number of allowed visits $h_{\max}$ is respected. At each step, the best candidate insertion is selected to maximize revenue while preserving schedule feasibility. The hierarchical nature of enforcement (i.e., deterrence strength based on inspection timing and frequency) is respected by prioritizing spatial-temporal diversity in visit distribution.

Motivating Example. Consider two officer tours:

- Tour 1 (π_{o_1}): $0 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 0$
- Tour 2 (π_{o_2}): $0 \rightarrow 5 \rightarrow 3 \rightarrow 2 \rightarrow 0$

Here, lots 2 and 3 appear in both tours. The inspection time of lot 2 in Tour 1 depends not only on its local sequence within the tour, but also on the inspection schedule of lot 2 in Tour 2. This is due to the recovery time t_{rec} between consecutive inspections and the constraint that each lot can be inspected by only one officer at a time. Furthermore, lot 2 precedes lot 3 in Tour 1, while the reverse occurs in Tour 2, creating a bidirectional timing dependency.

Hierarchical Relaxation To eliminate such cyclic dependencies and ensure temporal feasibility, we impose a hierarchical visitation rule: for any parking lot i, a visit by officer o_1 (i.e., the officer with the lower index) must precede that by officer o_2 whenever $o_1 < o_2$. Formally, if $i \in \pi_{o_1} \cap \pi_{o_2}$, then the inspection time by o_2 (i.e., the officer with the higher index) must satisfy

$$a_{o_2}(i) \geq \max_{v \in \pi_{o_1}, v=i} \{a_{o_1}(v)\} + t_{\text{rec}} + t_i^{\text{insp}}$$

where the maximum operator selects the latest inspection time of lot i in officer o_1's tour. This accounts for the possibility of multiple visits to the same lot within a single officer's route. The transformation thus converts the bidirectional recovery constraint into a well-ordered precedence structure, simplifying the enforcement of temporal feasibility during solution construction.

Greedy Heuristic Our procedure differs in its specific steps, but follows the general framework of [134] and is based on a greedy insertion heuristic. We initialize all officer tours as empty and iteratively insert lots $i \in V$ into routes, selecting at each iteration the insertion with the best revenue–time ratio [135, 136]:

$$\text{Ratio} = \frac{\Delta R(i)}{\Delta T(i)}$$

where $\Delta R(i)$ is the increase in revenue and $\Delta T(i)$ the increase in total route time caused by inserting i at a given position. Insertions are only accepted if:

- Total duration $T_o \leq T_{\max}$ for all officers;
- Lot i is not yet visited $h_{\max}$ times;
- Hierarchical and recovery constraints are preserved.

Each iteration evaluates possible insertions for all lots not yet scheduled $h_{\max}$ times, and selects the best one. This process continues until no feasible improvement is found.

Overview of the VND

The method begins with an initial solution generated via the greedy insertion algorithm designed to maximize expected revenue while respecting time and visiting constraints. The core local search component is based on VND, which systematically applies a sequence of neighborhood structures. Whenever a local optimum is reached in the current neighborhood, the search moves to the next one. If no improvement is found across all neighborhoods, a shaking procedure is invoked to diversify the search and escape local optima.

Neighborhood Structures

To enhance a feasible solution $\mathcal{S}$, composed of patrol routes $\{\pi_o\}_{o \in O}$, we define a set of neighborhood structures $\mathcal{N}_k(\mathcal{S})$ for $k \in \{1, 2, 3, 4\}$. Each neighborhood introduces a specific

Algorithm 5 Compute total revenue with hierarchical structure

$R \leftarrow 0$
Initialize $T[o] \leftarrow 0 \ \ \forall o \in O$
Initialize $V_t[i] \leftarrow 0, \ v[i] \leftarrow 0 \ \ \forall i \in V$
for each officer $o \in O$ **do**
  route $\leftarrow$ tour of officer o
  $t_{\text{curr}} \leftarrow 0$
  prev $\leftarrow$ route[0]
  for position p in route **do**
    $i \leftarrow$ route$[p]$
    $t_{\text{cand}} \leftarrow \begin{cases} \max(t_{\text{curr}} + d_{prev,i}, \ V_t[i] + t_{rec}), & \text{if } V_t[i] \neq 0, \\ t_{\text{curr}} + d_{prev,i}, & \textit{otherwise.} \end{cases}$
    $V_t[i] \leftarrow t_{\text{cand}}$ ▷ Updating start time of inspection
    $t_{\text{curr}} \leftarrow t_{\text{cand}} + t_i^{\text{insp}}$
    $v[i] \leftarrow v[i] + 1$ ▷ Updating number of visits
    prev $\leftarrow i$
  end for
  $t_{\text{curr}} \leftarrow t_{\text{curr}} + d_{prev,0}$ ▷ Return to depot
  $T[o] \leftarrow t_{\text{curr}}$
end for
for each lot $i \in V$ **do**
  Let $c = v[i]$ ▷ Number of visits to lot i
  $R \leftarrow R + \Lambda_i(\delta_i \cdot \pi_{i,c} \cdot \beta_{i,c} \cdot + (1 - \beta_{i,c}) \cdot \rho_i)$
end for
return R, T

Algorithm 6 Greedy insertion algorithm to construct initial solution

while True **do**
    Initialize best_ratio $\leftarrow -\infty$, best_insertion $\leftarrow \emptyset$
    Compute current revenue R_{old}, durations T^o_{old}
    for lot $i \in L$ not yet visited $h_{\max}$ times **do**
        for officer $o \in O$, route route$_o$ **do**
            for position p in route **do**
                Temporarily insert i at position p
                Compute $R_{\text{new}}, T^o_{\text{new}}$ via **Algorithm 5**
                if All $T^o_{\text{new}} \leq T_{\max}$ and $R_{\text{new}} > R_{\text{old}}$ **then**
                    $Ratio = \frac{(R_{\text{new}} - R_{\text{old}})}{\sum_o (T^o_{\text{new}} - T^o_{\text{old}})}$
                    if $Ratio >$ best_ratio **then**
                        Update best insertion (i, p, o)
                        Update best_ratio
                    end if
                end if
                Undo insertion
            end for
        end for
    end for
    if best_insertion $= \emptyset$ **then break**
    else Insert i^* at p^* in route of o^*
    end if
end while

type of move designed to improve the total collected revenue $R(\mathcal{S})$ or to reduce the aggregate tour duration $\sum_{o \in O} T_o(\mathcal{S})$, while strictly maintaining feasibility. A *first-improvement* strategy is employed within each neighborhood. As soon as a feasible move that improves the solution is found, the current solution is updated, and the search restarts from the first neighborhood to promote intensification. After each move, we invoke Algorithm 7 to ensure feasibility and to evaluate the resulting solution.

Neighborhood 1: Insertion ($\mathcal{N}_1$) This neighborhood inserts a parking lot i into a different position within an officer's tour π_o, provided the lot has not reached its maximum allowed visits across all officers. The lot is placed between two consecutive nodes in the tour. This operator can increase coverage of valuable lots or improve travel efficiency.

Neighborhood 2: Intra-route Swap ($\mathcal{N}_2$) This neighborhood swaps the positions of two non-depot nodes within the same route. The set of visited lots is unchanged, but reordering can reduce travel distance or improve the timing of visits with recovery constraints in mind.

Neighborhood 3: Replace ($\mathcal{N}_3$) This operator introduces a lot i that is not currently in any tour by removing an existing lot j at some position in a tour π_o and replacing it with i, keeping the rest of the route intact.

Neighborhood 4: Cross-Exchange ($\mathcal{N}_4$) This neighborhood exchanges short contiguous segments (two or three consecutive lots, excluding depots) between two different officer routes. Redistributing sub-sequences can improve workload balance, reduce travel, or create more efficient spacing between repeated visits.

Acceptance Condition for All Neighborhoods. Any candidate solution $\mathcal{S}' \in \mathcal{N}_k(\mathcal{S})$ is accepted if:

$$R(\mathcal{S}') > R(\mathcal{S}) \quad \text{or} \quad \left(R(\mathcal{S}') = R(\mathcal{S}) \ \wedge \ \sum_{o \in O} T_o(\mathcal{S}') < \sum_{o \in O} T_o(\mathcal{S}) \right) \tag{5.32}$$

5.4.3 Feasibility Repair Procedure

A solution that is initially feasible may become infeasible after the application of local search operators. Such modifications can exceed the maximum allowable tour duration, or fail to respect the minimum recovery time t_{rec} between consecutive inspections of the same lot. To ensure the feasibility of the operational planning, any modification to the solution must be followed by a feasibility repair procedure. Our proposed method consists of two interlinked phases: the first resolves recovery time conflicts across multiple routes by iteratively delaying conflicting visits; the second enforces the maximum tour duration $T_{\max}$ by removing jobs with the highest time burden from overlong routes. This iterative process guarantees that all constraints are satisfied. The complete approach is described in Algorithm 7.

Initially, each officer's route is processed independently to compute the inspection schedule using the travel times and node service durations. For each tour, the arrival times are calculated from the depot onward. After this, we check for recovery time violations across all routes. For every pair of inspections on the same lot by officers, if difference of inspection start times are less than the minimum required t_{rec}, the earlier visit is retained, and the later conflicting one is shifted forward in time to restore feasibility. This shif will be store in *ShiftMap*. Once a shift is made, the downstream schedule of that officer must be recomputed recursively, and all inter-route conflicts re-evaluated.

After resolving all recovery time conflicts, we verify whether any officer's tour exceeds the maximum allowed tour duration $T_{\max}$. If so, the inspection of a lot contributing the highest from/to travel time and inspection time in the tour is removed. This process, inspection scheduling, conflict resolution, and duration enforcement, is repeated iteratively until a fully feasible solution is obtained.

5.4.4 Shaking Mechanism: Random Drop and Reinsert Strategy

While VND is effective at intensifying the search within promising regions of the solution space, several studies [137–140] have shown that hybridizing VND with shaking procedures—such as those used VNS — can significantly improve solution quality. To this end, we integrate a dedicated shaking mechanism into our VND-based solution approach for our problem.

In metaheuristic paradigms, shaking is used to introduce controlled diversification and to escape local optima by perturbing the incumbent solution. Here, we implement a *Random Drop and Reinsert* strategy specifically designed for the multi-visit nature of our problem. The key idea is to remove a subset of parking lot visits from the current solution and reinsert them in a randomized yet feasible manner, thus exploring alternative configurations while maintaining solution validity.

Formally, given the current set of routes $\{\pi_o\}_{o \in \mathcal{O}}$, the shaking procedure performs the following steps:

- **Drop Phase:** A set of k visits are selected uniformly at random from the current solution

Algorithm 7 Conflict Resolution and Feasibility Repair

Input: A set of officer routes $\mathcal{R} = \{\pi_1, \pi_2, \ldots, \pi_m\}$; recovery time t_{rec}; maximum tour duration $T_{\max}$
Output: A feasible set of routes with no recovery or time violations
ShiftMap $\leftarrow \emptyset$
while True **do**
    Phase 1: Schedule and Conflict Resolution
    for each route $\pi_o \in \mathcal{R}$ **do** ▷ disregard other routes
        Compute arrival and completion times of all visits in π_o considering *ShiftMap*
    end for
    Identify all recovery time conflicts between tours
    if conflicts exist **then**
        Select the earliest conflict; let v_1 and v_2 be visits such that v_1 precedes v_2
        Shift v_2 and all subsequent visits in its tour forward to satisfy t_{rec}
        Save the shift in *ShiftMap*
        restart Phase 1
    end if
    Phase 2: Tour Duration Enforcement
    for each route $\pi_o \in \mathcal{R}$ **do**
        if total tour time of π_o exceeds $T_{\max}$ **then**
            Identify the visit with the maximum (inbound travel + outbound travel + service time)
            Remove that visit from π_o
            ShiftMap $\leftarrow \emptyset$
            return to Phase 1 ▷ restart conflict checking
        end if
    end for
    break ▷ All constraints satisfied, exit loop
end while
Update $\mathcal{R}$
Phase 3: Final Evaluation
Calculate tour times (T) and total revenue (R)
return $\mathcal{R}$, T and R

and removed. Each visit is defined by a triplet (o, i, ℓ) indicating that officer o visits lot ℓ at position i in their route.

- **Reinsertion Candidates:** After removal, we construct the set of unvisited lots $\mathcal{U}$, defined as the complement of lots still present in the routes. If $|\mathcal{U}| \geq k$, we randomly choose k lots from $\mathcal{U}$. Otherwise, we take all unvisited lots and supplement with randomly selected already-visited lots to reach k.

- **Add Phase:** Each selected lot is assigned to a randomly chosen officer $o \in \mathcal{O}$ and inserted at a randomly chosen valid position within π_o, preserving the feasibility of the route (e.g., maintaining depot as the terminal node).

This shaking mechanism is particularly well-suited for the problem because it allows for dynamic redistribution of visits across officers while also creating opportunities to explore regions of the solution space that deterministic local search might overlook. Moreover, by prioritizing unvisited lots for reinsertion, the method encourages the discovery of high-reward configurations that may have been missed in previous iterations.

Adaptive Search Strategy Based on Upper Bound Analysis

In this section, we present an adaptive variant of our solution approach specifically tailored to address the complexity introduced by the multi-visit feature. Allowing multiple visits per node significantly expands the solution space, as the number of visits becomes an additional decision variable that must be optimized alongside routing.

To manage this complexity, we propose a dynamic adaptation mechanism that gradually expands the allowable visit levels during the search process. This mechanism aims to balance diversification and intensification by allocating computational effort more strategically across visit levels.

As a guiding principle, we leverage insights from the KUB model introduced earlier. This relaxed model provides an estimate of the potential number of visits per node by abstracting away routing constraints and focusing solely on visit feasibility under a global time budget. The average number of visits suggested by the KUB serves as a proxy for expected inspection density across the network.

We define the ratio M as Total number of visits/Number of nodes, representing the expected visit density. This value is used as the mode of a triangular probability distribution over the interval $[0, h_{\max}]$. The triangular distribution is designed to allocate more search effort to

Algorithm 8 Adaptive-VND based on Knapsack Upper Bound

Solve the Knapsack Upper Bound model and compute $M \leftarrow \text{TotalVisits}/|V|$
Compute schedule $\{(H'_1, T_1), (H'_2, T_2), \dots\}$ using triangular distribution on $[0, h_{\max}]$ with mode M
Initialize solution $S \leftarrow \text{GreedyInsertion}()$
for each (H', T) in the adaptive schedule **do**
  for $t = 1$ to T **do**
    Set $k \leftarrow 1$, $k_{\max} \leftarrow 4$
    repeat
      $S' \leftarrow \text{ApplyNeighborhood}(S, k, H')$
      if S' improves S **then**
        $S \leftarrow S'$, $k \leftarrow 1$
      else
        $k \leftarrow k + 1$
      end if
      if $k > k_{\max}$ **then**
        $S \leftarrow \text{Shake}(S)$, $k \leftarrow 1$
      end if
    until local termination criterion for T iterations is met
  end for
end for
return S

visit levels near the mode M, while still allowing for exploration of the entire range of visit levels.

The triangular distribution is defined as:

$$f(h) = \begin{cases} \frac{2h}{h_{\max} M}, & \text{if } 0 \leq h \leq M, \\ \frac{2(h_{\max} - h)}{h_{\max}(h_{\max} - M)}, & \text{if } M < h \leq h_{\max}. \end{cases}$$

This function concentrates probability mass around lower visit levels when M is small and shifts the emphasis to higher visit levels as M increases.

To distribute the total number of iterations T among discrete visit levels $H' = 1, 2, \dots, h_{\max}$, we compute the probability mass for each interval $[h-1, h)$ using the cumulative distribution function $F(\cdot)$ associated with the triangular distribution:

$$P_h = F(h) - F(h-1), \quad h = 1, \dots, h_{\max}.$$

The number of iterations assigned to each level $H' = h$ is then:

$$T_h = \text{round}\left(T \cdot P_h\right).$$

This adaptive iteration allocation mechanism prioritizes search around the most promising visit levels as informed by the upper bound model. When M is small, the search intensifies around lower visit levels; as M increases, higher visit levels receive proportionally more attention. This approach embeds a data-informed, gradual expansion strategy into the VND framework, balancing early exploitation with later-stage exploration. To preserve the value of the initial solution, we first execute a single probe iteration at small visit levels and skip the remaining search budget for that level whenever the resulting solution is not sufficiently close to the initial one in terms of revenue. The complete integration of this adaptive control mechanism into our metaheuristic is described in Algorithm 8.

5.5 Computational Study

We evaluate the proposed routing model for parking enforcement through a systematic computational study designed (i) to assess performance across key parameter settings and (ii) to examine the efficiency of the solution algorithm. All experiments were run with Gurobi Optimizer 12.0.0 on a machine with an Intel® Core™ i7–7800X @ 3.50 GHz (6 physical cores, 12 logical processors) under AlmaLinux 9.5, using Python 3.9.

With the computational environment established, we begin with a focused experiment on $T_{\max}$ to illustrate the interaction between operational limits and behavioral feedback. As established in Section 5.3.1, $T_{\max}$ acts simultaneously as (i) a hard operational limit on route duration and (ii) a deterrence divisor that scales the average enforcement intensity perceived by drivers. To expose these mechanisms, Deterrence Dilution (economic feedback) and Resource Expansion (classical routing gain), we solve a controlled small-network instance while gradually increasing $T_{\max}$. In this setting, small increases in $T_{\max}$ that leave the visit pattern unchanged dilute enforcement intensity and reduce revenue, whereas larger increments that enable an additional feasible visit can more than compensate for dilution through added operational benefit.

The test case consists of five serviceable nodes, two officers, and fixed behavioral parameters ($\lambda_1 = 0.6$, $\lambda_2 = 0.3$, $\xi_0 = 2.0$, $\delta = 10$, $\phi = 2$, $t^{\text{insp}} = 45$ minutes). Table 5.1 reports the optimal revenue and the total number of visits (C_{total}).

The results reveal a non-monotonic, step-wise response of the optimal revenue to increases

Table 5.1 Sensitivity of Optimal Revenue to Maximum Tour Duration ($T_{\max}$)

$T_{\max}$ (min)	Revenue	C_{total}	$(c_1, c_2, c_3, c_4, c_5)$	Trend	Interpretation
240	1469.27	6	(1,2,1,1,1)	—	Baseline solution
245	1464.47	6	(1,2,1,1,1)	↓	Dilution dominates
250	1459.71	6	(1,2,1,1,1)	↓	Continuing dilution
255	1483.08	7	(1,2,1,2,1)	↑	Resource expansion threshold reached
260	1478.82	7	(1,2,1,2,1)	↓	Dilution resumes
265	1474.59	7	(1,2,1,2,1)	↓	Continuing dilution
270	1493.77	8	(2,2,1,2,1)	↑	Second expansion threshold reached

C_{total} is the sum of all node visits (excluding depots).

in $T_{\max}$. When $T_{\max}$ increases from 240 to 250 minutes, the optimal visit pattern remains unchanged ($(c_1, c_2, c_3, c_4, c_5) = (1, 2, 1, 1, 1)$). Because the same enforcement effort is spread over a longer horizon, the average enforcement intensity $\bar{k}_i$ decreases, reducing citation probability and revenue (1469.27 → 1459.71). This illustrates the deterrence dilution effect for small step sizes.

When $T_{\max}$ crosses a feasibility threshold (250 → 255 minutes), an additional profitable visit becomes possible (at Node 4), producing a discrete jump in revenue and in the total number of visits C_{total}. This is the classical resource expansion effect: a sufficiently large increase in $T_{\max}$ that alters the visit vector converts marginal capacity into additional enforcement benefit. The same pattern repeats at $T_{\max} = 270$, which enables a second visit to Node 1.

Now, we conduct computational experiments to evaluate the performance of the proposed solution approaches and benchmark them against Gurobi under controlled instance families. The problem instances are generated under the following parameter settings. The recovery time between consecutive inspections of the same parking lot is set at 50 minutes. The maximum operational time per enforcement officer during a shift is limited to 250 minutes. Each violation incurs a fixed fine of 10 currency units while the parking fee is 3 currency units.

The problem instances are based on complete graphs consisting of N nodes, where each node represents either a parking lot or the depot. Node locations are randomly generated over a 100x100 grid using a uniform distribution. The depot is fixed at the center of the grid, positioned at coordinates (50, 50). Travel times between nodes are determined using the Euclidean distance. A random vector is generated to represent the parking arrival rates at each lot, with values ranging between 50 and 100. The inspection time at lot i is set to $t_i^{\text{insp}} = 0.20\,\Lambda_i$ minutes. This scaling is a pragmatic heuristic to increase inspection time with demand and is not intended to be dimensionally exact.

We benchmark VND and VND+S against Gurobi on small- and medium-sized instances. Instances are characterized by the number of lots (N), the maximum inspections per lot (H), and the number of officers (O), denoted $N–H–O$. For each configuration, we generate 10 random instances and report averages of total revenue and runtime, along with Gurobi's optimality gap at termination ("Gap") and the gap to a knapsack-based upper bound ("KUB Gap"). Gurobi uses a 3000 s time limit; VND+S runs for 100 iterations.

We quantify relative objective differences as

$$\text{Dif}(\%) = \ 100 \times \frac{\text{Algorithm} - \text{MIP}}{\text{MIP}}$$

so positive values indicate the heuristic achieved a higher objective than the MIP incumbent within the time limit. The performance of the VND-based algorithms was recorded without any time restrictions, and their average computational times are reported.

Table 5.2 Comparison of Gurobi, VND, and VND+Shaking

Ins.	Gap	KUB Gap	Ave. Time	VND Dif	VND Time	VND+S Dif	VND+S Time
8-2-2	8.37	6.17	1546.76	-4.18	0.04	1.36	2.94
8-3-3	7.82	6.11	1783.37	-11.34	0.11	0.08	8.55
8-3-4	3.98	3.80	1277.70	-14.50	0.08	-0.04	17.24
10-2-2	14.52	8.92	1480.51	-3.86	0.07	0.26	3.21
10-3-3	11.75	7.16	2173.86	-8.24	0.25	0.15	11.98
10-3-4	6.75	5.94	2043.50	-12.34	0.14	0.09	26.99
15-2-2	28.96	14.85	1800.16	-3.79	0.13	0.89	6.55
15-3-3	21.31	10.18	2199.87	-3.05	0.38	0.72	19.44
15-3-4	13.78	9.06	1666.72	-6.51	3.22	1.18	50.56
20-2-2	50.08	21.47	2292.65	-0.07	0.33	4.10	11.78
20-3-3	33.91	14.68	2101.26	0.39	0.55	3.86	26.86
20-3-4	22.84	11.90	1975.55	-0.42	6.55	1.80	64.15
30-2-2	78.22	22.49	2483.89	2.68	0.78	5.80	23.32
30-3-3	65.64	25.89	2290.53	7.16	1.04	10.89	45.33
30-3-4	46.10	20.99	2129.49	6.86	1.70	11.57	77.89

The computational results indicate that the proposed VND and VND+S algorithms perform effectively in solving the parking enforcement problem within negligible computational time compared to Gurobi. While Gurobi requires a substantial amount of time (typically over 1500 seconds per instance and up to the limit of 3000 seconds for larger instances), both VND-based algorithms deliver solutions within seconds.

VND alone is fast but exhibits variable quality: on smaller instances it is often slightly below the MIP incumbent, yet it occasionally surpasses it in select settings). VND+S consistently improves solution quality over VND across all settings while preserving low runtimes.

As the problem size increases (e.g., $N = 20$ and $N = 30$), Gurobi's optimality gap also grows, reflecting the difficulty in closing the gap to optimality within the imposed time limit. In contrast, VND+S maintains reasonable solution quality, with gaps up from -0.04% to 11.57% relative to Gurobi's solutions. Moreover, the computational time for VND+S scales well, remaining under 100 seconds even for the largest instances.

In these regimes, the basic VND algorithm also shows acceptable performance (from -0.42% below to $+6.86\%$ above the MIP incumbent across settings) with low computation times (max 6.55 s). Overall, the results validate the effectiveness and scalability of the VND-based heuristics, particularly when fast solutions are required. The shaking mechanism adds robustness, especially in avoiding poor local solutions for larger and more constrained instances. From this point onward, VND+S is adopted as the reference algorithm for our analysis.

5.5.1 Local Search Heuristic Component Analysis

To evaluate the individual contributions of each local search operator, we perform an ablation study in which each neighborhood, *Insertion*, *Swap*, *Replace*, *Cross-Exchange*, and *Swap+Replace* together, is removed from the full VND configuration and the results are compared to the baseline. The experiments are conducted on three problem instances (15-3-3, 20-3-3, and 25-3-3) under two route time limits ($T_{\max} = 200$ and $T_{\max} = 300$) with a fixed recovery time of 25 units. Table 5.3 reports the average runtime, objective value, time savings, and revenue loss for each configuration.

The *Insertion* operator is the most critical for achieving high-quality solutions: its removal yields the largest revenue losses (on average 7.56% for $T_{\max} = 200$ and 7.43% for $T_{\max} = 300$), despite substantial time savings (88.99% and 80.97%, respectively). *Swap* has minimal impact on revenue (below 1% loss) but yields moderate time savings, indicating its role in fine-tuning sequences rather than restructuring routes. *Replace* provides a balance between quality and speed, with moderate losses (around $1-2\%$) and high time savings under tighter limits. *Cross-Exchange* delivers intermediate effects, supporting diversification through structured multi-edge moves. Removing *Swap+Replace* lowers runtime significantly but at the cost of higher revenue loss, indicating complementary value when used together.

Across problem sizes, the trends are consistent: *Insertion* dominates in smaller instances, *Swap* and *Replace* gain importance in medium instances, and *Replace* and *Cross-Exchange* become increasingly valuable in larger, more complex settings. These results confirm that while *Insertion* is indispensable, a balanced mix of neighborhoods enhances both robustness and efficiency of the search.

Table 5.3 Component analysis results for different local search operators

Problem	Configuration	$T_{\max} = 200$				$T_{\max} = 300$			
		Time (s)	Obj.	Time Sav. (%)	Loss (%)	Time (s)	Obj.	Time Sav. (%)	Loss (%)
15-3-3	Baseline	11.93	2294.37	–	–	41.45	2537.67	–	–
	Insertion	2.06	2122.50	82.76	7.49	8.21	2245.49	80.20	11.51
	Swap	7.94	2291.22	33.45	0.14	24.64	2526.67	40.55	0.43
	Replace	6.30	2245.49	47.16	2.13	38.74	2536.72	6.56	0.04
	Cross-Exchange	6.75	2251.12	43.42	1.89	25.50	2487.95	38.48	1.96
	Swap+Replace	3.87	2245.49	67.60	2.13	25.16	2522.31	39.30	0.61
20-3-3	Baseline	33.02	2774.06	–	–	59.07	3208.06	–	–
	Insertion	2.95	2469.08	91.06	10.99	14.17	2955.30	76.02	7.88
	Swap	29.37	2747.45	11.07	0.96	33.11	3169.10	43.95	1.21
	Replace	7.90	2719.11	76.08	1.98	47.35	3149.83	19.85	1.81
	Cross-Exchange	17.54	2683.65	46.87	3.26	35.57	3129.94	39.78	2.43
	Swap+Replace	4.29	2623.48	87.01	5.43	29.15	3098.07	50.65	3.43
25-3-3	Baseline	62.08	3098.47	–	–	95.09	3953.73	–	–
	Insertion	4.26	2968.74	93.14	4.19	12.64	3839.07	86.70	2.90
	Swap	58.61	3050.24	5.59	1.56	77.17	3906.80	18.84	1.19
	Replace	4.80	3050.86	92.28	1.54	32.50	3898.04	65.82	1.41
	Cross-Exchange	31.06	3040.28	49.97	1.88	50.32	3881.62	47.08	1.82
	Swap+Replace	2.68	2974.96	95.68	3.99	18.54	3766.20	80.50	4.74
Average	Insertion			88.99	7.56			80.97	7.43
	Swap			16.71	0.88			34.45	0.94
	Replace			71.84	1.88			30.74	1.09
	Cross-Exchange			46.75	2.34			41.78	2.07
	Swap+Replace			83.43	3.85			56.82	2.93

"Time Sav. %" and "Loss %" are relative to the baseline (full set of operators).

5.5.2 Effect of Arrival Rate and Parking Price on Enforcement Planning

This section investigates how arrival rates and parking prices influence enforcement decisions in a parking network. A network with 6 parking lot and a depot is considered. Three distinct price scenarios are examined: (i) strong correlation between price and arrival rate (price range 1–4), (ii) moderate correlation (price range 2–3), and (iii) no correlation (fixed price at 2.5). In the first case, prices are generated through a linear scaling with arrival rates. For all scenarios, the maximum number of inspections per node is set to $H = 2$, the maximum route duration is $T_{\max} = 250$, and two recovery time settings are considered: $t_{rec} = 10$ and $t_{rec} = 100$.

To isolate the effect of arrival rates, we vary only the arrival rate at Node 1, $\Lambda_1 \in \{25, 50, 75, 100\}$, while others are held constant. The results, including the tours, objective values, and the frequency of inspections at Node 1, are summarized in Table 5.4.

Table 5.4 demonstrates that as the arrival rate at Node 1 increases, the number of inspections at this node also increases. Specifically, when $\Lambda_1 = 25$, Node 1 is either omitted or visited once, whereas for $\Lambda_1 = 75$ and $\Lambda_1 = 100$, it is visited twice regardless of the pricing policy or recovery time.

This finding aligns with practical enforcement policies where areas with higher arrival rates, such as downtown districts, are typically prioritized for inspections. These areas often feature both higher parking demand and higher prices, contributing to increased enforcement activity. Although pricing differences impact the objective function values, arrival rates more strongly dictate inspection frequency.

Interestingly, even under a fixed pricing scheme, the enforcement pattern adapts to the higher demand at Node 1, indicating the robustness of the enforcement strategy to prioritize high-demand zones regardless of price variations. This suggests that, in practice, while pricing can influence revenue outcomes, enforcement frequency is more tightly linked to demand intensity.

Table 5.4 Effect of Arrival Rate and Parking Price on Tours, Objective Values, and Node 1 Inspections

Price Interval	**Arrival rate at Node 1(Λ_1)**	$t_{rec} = 10$		$t_{rec} = 100$		**Visits to**	**Visits to**
		Tour	**Obj**	**Tour**	**Obj**	**Node 1** ($t_{rec} = 10$)	**Node 1** ($t_{rec} = 100$)
1–4	25	0-6-4-2-3-2-4-0	641.27	0-6-4-2-3-4-0	629.80	0	0
	50	0-6-1-4-2-3-0	697.54	0-6-1-4-2-3-0	697.54	1	1
	75	0-1-5-1-4-2-3-0	772.07	0-6-1-4-2-3-0	768.76	2	1
	100	0-1-5-1-4-2-3-0	849.57	0-6-1-4-2-3-0	834.22	2	1
2–3	25	0-6-4-2-3-2-4-0	641.65	0-6-3-2-4-1-0	639.72	0	1
	50	0-6-1-4-2-3-0	702.32	0-6-1-4-2-3-0	702.32	1	1
	75	0-1-5-1-4-2-3-0	759.84	0-6-1-4-2-3-0	753.35	2	1
	100	0-1-5-1-4-2-3-0	815.12	0-6-1-4-2-3-0	800.47	2	1
2.5–2.5	25	0-6-3-2-4-1-0	647.77	0-6-3-2-4-1-0	647.77	1	1
	50	0-6-1-4-2-3-0	703.29	0-6-1-4-2-3-0	703.29	1	1
	75	0-1-5-1-4-2-3-0	751.96	0-6-1-4-2-3-0	743.96	2	1
	100	0-1-5-1-4-2-3-0	792.82	0-6-1-4-2-3-0	778.89	2	1

5.5.3 Comparison Between Adaptive and Classic Algorithms

To evaluate the performance differences between the adaptive and classic routing algorithms for parking enforcement, extensive experiments were conducted under varying conditions. These experiments consider network sizes ($N \in \{10, 15, 20, 25\}$), number of officers ($O \in \{2, 3\}$), inspection limits ($H \in \{1, 2, 3, 4\}$), and route time limits ($T_{\max} \in \{250, 500\}$).

Both algorithms were run independently with 10 replications per scenario. Average objective values and computational times were recorded. The results are presented in two tables. Table 5.5 summarizes the performance of the adaptive method, while Table 5.6 shows the results for the classic approach.

To highlight the relative performance differences between the two methods, Figure 5.1 visu-

Table 5.5 Performance Results of Adaptive Algorithm

			N=10		N=15		N=20		N=25	
	H	O	Time	Obj	Time	Obj	Time	Obj	Time	Obj
$T_{max} = 250$	1	2	1.06	1467.08	4.12	2245.49	8.76	2625.86	16.44	3221.36
		3	0.98	1467.08	4.50	2245.49	12.95	2955.30	22.38	3913.13
	2	2	2.58	1571.12	5.65	2245.49	9.77	2628.18	16.58	3221.85
		3	6.97	1668.53	14.92	2413.96	23.55	2964.67	29.07	3907.13
	3	2	2.73	1569.77	5.64	2245.49	10.04	2625.86	16.86	3221.74
		3	11.85	1695.92	20.21	2410.74	26.13	2958.99	31.69	3909.01
	4	2	2.75	1570.29	5.09	2245.49	10.62	2627.88	17.52	3221.85
		3	12.85	1701.95	19.20	2399.63	26.23	2965.13	33.02	3911.21
$T_{max} = 500$	1	2	1.27	1467.08	6.14	2245.49	17.44	2955.30	40.40	4131.44
		3	1.19	1467.08	6.50	2245.49	20.79	2955.30	42.35	4131.44
	2	2	5.86	1668.53	37.95	2549.85	66.17	3275.57	82.51	4275.13
		3	6.99	1668.53	42.35	2549.85	135.78	3351.95	284.57	4662.09
	3	2	24.42	1794.70	62.13	2609.43	70.76	3270.58	89.59	4297.49
		3	31.93	1794.70	242.74	2731.51	401.63	3511.98	398.91	4650.25
	4	2	36.74	1816.15	63.81	2609.25	73.25	3262.63	154.77	4288.70
		3	99.92	1877.62	341.74	2794.99	453.50	3507.29	430.27	4627.41

Table 5.6 Performance Results of Classic Algorithm

			N=10		N=15		N=20		N=25	
	H	O	Time	Obj	Time	Obj	Time	Obj	Time	Obj
$T_{max} = 250$	1	2	1.06	1467.08	4.15	2245.49	8.91	2625.86	16.63	3223.46
		3	0.98	1467.08	4.44	2245.49	12.41	2955.30	22.11	3913.13
	2	2	3.42	1571.76	7.38	2199.09	11.80	2563.82	19.39	3138.27
		3	8.06	1668.53	20.08	2422.44	30.74	2973.21	41.15	3719.46
	3	2	3.91	1568.78	8.06	2208.10	12.62	2542.48	20.00	3145.10
		3	16.06	1700.69	27.10	2408.23	35.41	2962.62	47.23	3681.25
	4	2	4.04	1568.27	7.83	2218.15	12.75	2587.11	20.21	3189.01
		3	18.30	1702.11	27.75	2405.67	37.86	2948.04	45.45	3739.60
$T_{max} = 500$	1	2	1.28	1467.08	5.98	2245.49	17.14	2955.30	37.14	4131.44
		3	1.31	1467.08	6.57	2245.49	21.44	2955.30	44.61	4131.44
	2	2	8.75	1668.53	56.56	2549.85	91.35	3289.22	106.31	4288.67
		3	10.00	1668.53	61.02	2549.85	194.00	3351.95	375.42	4655.72
	3	2	40.64	1794.70	86.60	2609.80	106.78	3249.38	119.17	4288.88
		3	55.68	1794.70	360.26	2731.51	463.69	3514.91	501.86	4584.09
	4	2	57.22	1819.43	90.63	2601.04	108.99	3244.94	126.81	4288.67
		3	219.93	1877.62	470.84	2792.47	772.35	3470.76	559.38	4586.65

alizes the percentage improvement (indicated by circles) or decline (indicated by crosses) of the adaptive approach over the classic approach. Differences smaller than 1% are omitted for clarity. In the figure, the Y-axis labels are formatted such that the first value denotes the visit level H, while the second value indicates the number of officers. For example, the label "4–3" represents a scenario with $H = 4$ possible visit levels and 3 officers. The X-axis corresponds to the problem size. The percentage differences are computed as follows:

$$\text{Difference (\%)} = 100 \times \frac{\text{Adaptive} - \text{Classic}}{\text{Classic}} \tag{5.33}$$

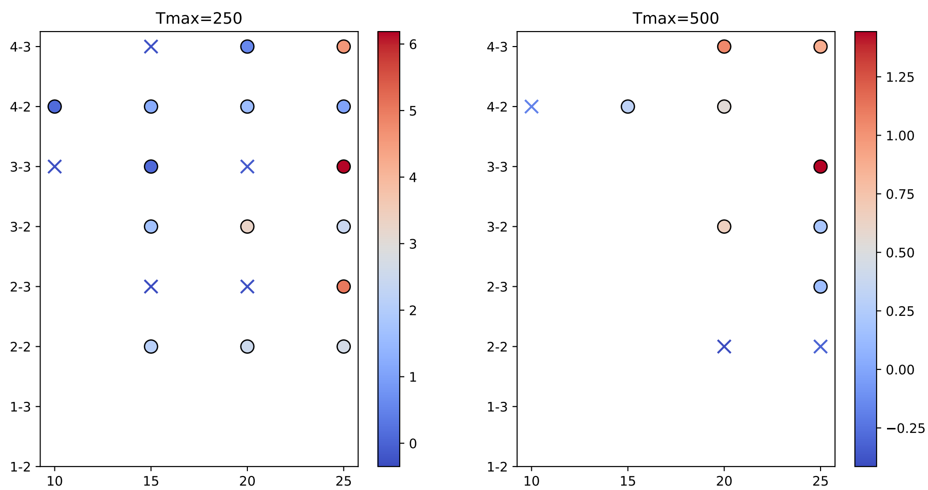

Figure 5.1 Relative Differences (%) Between Adaptive and Classic Methods

The adaptive algorithm consistently outperforms the classic method in most scenarios with larger networks or tighter constraints (e.g., higher H or higher N) for both $T_{\max} = 250$ and $T_{\max} = 500$. The improvements, although modest in some cases, highlight the adaptive algorithm's ability to better exploit routing flexibility under operational constraints. When the visit level is restricted to $H = 1$, both algorithms perform similarly since no adaptability is possible. Moreover, for $T_{\max} = 500$ in smaller networks ($N = 10$ and $N = 15$), the differences are not significant.

5.5.4 Case Study: Optimizing Parking Enforcement in Montreal

To evaluate the practical effectiveness of our proposed model, we apply it to a real-world case study in the city of Montreal, Canada. The relevant dataset was assembled using information obtained from the Agence de Mobilité Durable de Montréal, which provides public data on parking fees, lot capacities, and locations.

For spatial simplification, adjacent lots were aggregated by summing their capacities, thereby reducing redundant parking lots. The resulting network comprises 30 lots. The coordinates of these lots, together with a single centralized depot (Agence de mobilité durable de Montréal, shown by the red pin), are illustrated in Fig. 5.2.

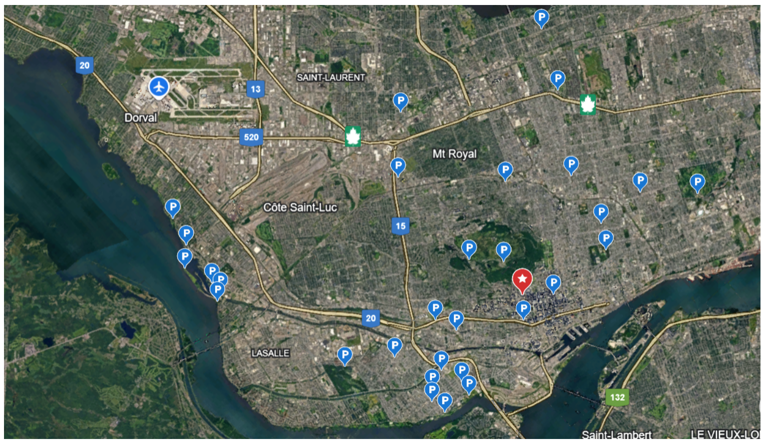

Figure 5.2 Location of Parking Lots in Montreal

Driving times between all pairs of parking lots and the central depot were computed using the Google Maps API in driving mode to reflect typical congestion. These travel times account for city traffic patterns and road network constraints.

Arrival rates at each lot were derived based on the number of parking spaces using the following tiered logic:

- If the lot has fewer than 40 spaces, its arrival rate equals the number of spaces.
- For lots with 40–80 spaces, the arrival rate is computed as $40 + 0.5 \times (\text{spaces} - 40)$.
- For lots exceeding 80 spaces, the arrival rate becomes $60 + 0.25 \times (\text{spaces} - 80)$.

Each violation detected results in a fine of \$91 CAD. The case study considers:

- 10 enforcement officers available for deployment.
- A maximum of 3 inspection visits allowed per parking lot.
- Each tour is limited to 420 minutes per officer.

- A recovery time of 30 minutes is enforced between consecutive visits to the same lot.

Our proposed solution method, based on proposed VND, was applied to this setup. The algorithm achieved an objective function value of 25,614.59 within 1128.56 seconds. Figure 5.3 presents the resulting Gantt chart of inspection schedules, indicating the allocation of visits across periods and officers.

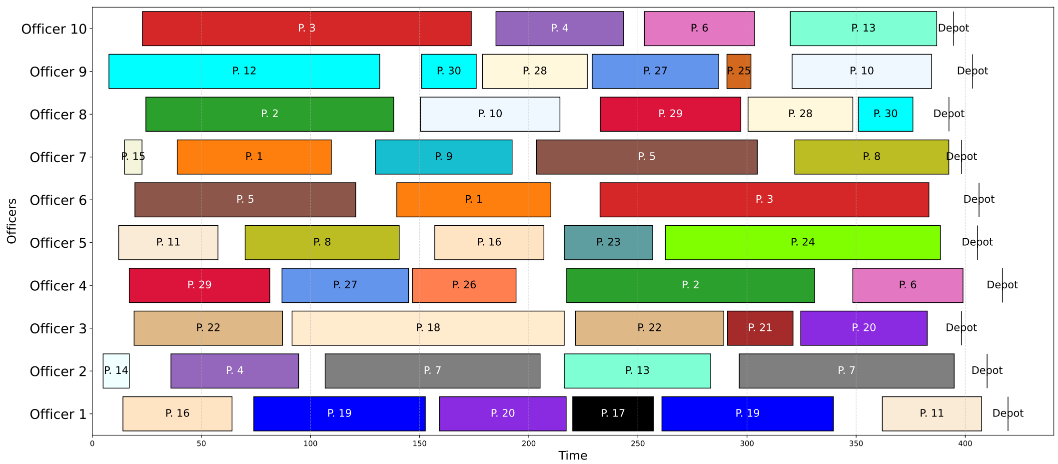

Figure 5.3 Inspection Schedules across Officers and Periods (Gantt Chart)

The inspection schedule tends to prioritize lots with higher arrival rates and parking fees, which aligns with the objective of maximizing revenue. For instance, Lot 3, which has the highest arrival rate (150.7), receives two inspections. In contrast, lots with fewer visitors or lower fees, such as Lot 15 (arrival rate = 8, fee = $15), are inspected only once. However, this pattern is not necessarily generalizable due to the complexity of the problem and the influence of other factors, such as the spatial distribution.

The case study demonstrates the applicability of our model in a realistic urban setting. By combining spatial, behavioral, and operational dimensions, our approach successfully distributes enforcement resources to maximize deterrence and revenue. The VND algorithm proved effective in handling real-scale data with multiple constraints while delivering interpretable and high-quality solutions.

5.6 Conclusions

This paper presents a novel approach to solving the Multi-Visit Team Orienteering Problem, tailored to the operational and behavioral complexities of parking enforcement. By modeling driver responses to inspection frequency and incorporating recovery time and tour duration constraints, we extend traditional orienteering frameworks to better reflect real-world enforcement dynamics. Our solution methodology integrates a Variable Neighborhood Descent metaheuristic with a dedicated repair procedure that enforces temporal feasibility through conflict resolution and route pruning.

The proposed conflict resolution algorithm plays a central role in maintaining feasibility, especially after local search modifications that may cause inter-route dependencies to violate recovery time rules. By hierarchically coordinating inspection schedules and eliminating excessive tours, the approach yields robust and behaviorally-informed patrol plans.

Future work can extend the framework by modeling stochastic arrival and dwell patterns, endogenizing compliance via equilibrium-based demand, and deploying adaptive (learning-based) inspection policies. Integrating real-time traffic feeds and predictive analytics would further support dynamic, on-the-fly adjustment of patrol schedules as urban conditions evolve.

CHAPTER 6 ARTICLE 3: TACTICAL PARKING ENFORCEMENT PLANNING: OPTIMIZING RESOURCE ALLOCATION UNDER DRIVER RESPONSE

Authors: Mohsen Yahyaei, Michel Gendreau and Nikolaj Van Omme, Submitted on November 2, 2025, Under review at Transportation Research Part A: Policy and Practice

Abstract: Municipal parking enforcement shapes driver compliance, equity, and revenue, yet most analytical models emphasize local equilibria or operational routing rather than the *tactical* allocation of resources across heterogeneous regions. We develop an integrated behavioral–optimization framework linking enforcement intensity, driver choice, and revenue. Using the Lambert–W function, we derive a closed-form critical staffing threshold $n^\star$ that delineates the transition from weak to effective deterrence. This behavioral mechanism is embedded in a dynamic program that allocates limited patrol units across regions while enforcing a uniform minimum detection-probability constraint as an equity requirement. Computational experiments on synthetic city archetypes and a six-region case study show that (i) *moderate fines* and *balanced staffing* maximize long-run revenue, (ii) *overstaffing* and *excessively high fines* yield diminishing returns, (iii) *technology that raises detection efficiency* can maintain deterrence with fewer officers, and (iv) *equity constraints* reallocate resources toward underserved areas at modest fiscal cost. The framework provides a tractable, data-informed basis for setting staffing, fines, and technology investments that unifies behavioral deterrence with tactical enforcement planning.

Keywords: Parking enforcement, Driver compliance, Behavioral modeling, Resource allocation, Urban mobility policy

6.1 Introduction

Municipal parking enforcement serves a dual function: a regulatory mechanism for managing congestion [141] and an essential source of public revenue [142]. The economic scale of these operations is significant, exemplified by New York City's collection of approximately $1.09 billion in fines from over 16 million violations in fiscal year 2024 [143]. This significant revenue stream highlights how enforcement policy directly influences municipal finances, yet its effectiveness hinges on the tactical deployment of resources. The central problem is how a limited, discrete budget of enforcement units should be allocated across a heterogeneous urban space to maximize long-run revenue while maintaining adequate deterrence. The allocation of enforcement units determines the driver's expected cost of violation (ECV), a perceived likelihood of detection that fundamentally shapes parking choices. A higher enforcement presence increases short-term collections by elevating detection rates, but excessive intensity can induce behavioral adaptation, reducing future violations and thus long-run revenue. Conversely, limited enforcement undermines deterrence. These interlinked mechanisms, resource allocation, driver response, and revenue outcomes, emphasize the need for adaptive, data-driven strategies that sustain compliance while ensuring predictable public income streams.

Understanding driver behavior is central to optimizing enforcement. The decision to park illegally is fundamentally an economic gamble, a choice formally rooted in deterrence theory, where motorists seek to minimize their total cost, balancing the search time and parking price to a legal space against the expected cost of violation. Opportunistic behavior arises from these underlying attitudes and is shaped by the relative scale of parking fees, fines, and the perceived likelihood of detection [117].

The foundational literature models parking choice as an equilibrium that balances parking price, search time, and the deterrent effect of enforcement. Arnott and Rowse [144] shows how providing information about parking availability affects driver cruising behavior and final choice, a critical factor in urban areas.

More recent empirical studies have concentrated on specialized driver behaviors, particularly those related to urban logistics. For example, Amaya et al. [145] found that the pressure for on-time delivery often compels delivery drivers to engage in illegal parking or double-parking when legal curb space is limited. Building on this understanding, analytical work has developed refined models for this specialized cohort. This includes discrete choice models for commercial vehicles that analyze how congestion, queuing, and enforcement intensity influence parking decisions [146]. Similarly, continuum models have been employed to investigate the system dynamics of on-street parking under the unique demands of urban deliverers, exploring optimal management strategies such as dedicated delivery bays and dynamic pricing [147].

The dimension of driver behavior extends beyond individual economic choice to encompass broader societal and spatial concerns. Enforcement policies must be viewed not only as revenue instruments but also through the lens of fairness and public trust. Empirical research shows that parking fines can function as a regressive tax when enforcement is uneven across neighborhoods, disproportionately affecting lower-income and renter-dense areas [47, 148]. Moreover, illegal parking behavior itself varies spatially, influenced by factors such as population density, curb availability, and local traffic dynamics [149]. These findings underscore the importance of incorporating equity considerations into enforcement planning, particularly when allocating limited patrol resources across heterogeneous regions.

Building on the behavioral foundation, several analytical studies have formalized the interaction between enforcement intensity and driver decision-making. Early economic frameworks, such as [144], conceptualized parking choice as an equilibrium between congestion, pricing, and the ECV. Petiot [75] modeled enforcement as a rational decision problem where drivers explicitly weigh the expected benefit of illegal parking against the ECV. This demonstrated that simply raising penalties does not always enhance compliance, as behavioral adaptation may occur. Complementing these equilibrium approaches, Martens and Benenson [97] used an agent-based simulation to capture how inspection frequency and parking fees jointly affect spatial parking patterns and municipal revenue.

Recent advances in sensing and artificial intelligence have transformed parking enforcement from manual inspection into a data-driven and adaptive activity. Shao et al. [57] formulated the traveling officer problem, showing how information from in-ground sensors can be combined with probabilistic routing algorithms to help officers prioritize areas with a higher chance of detecting violations within limited patrol time. Peng et al. [150] further expanded this idea by applying deep learning to videos captured by in-vehicle cameras, introducing a *minimal illegal units* labeling approach that allows the system to identify several types of parking offences in real time while lowering manpower and installation expenses. Chen et al. [151] later proposed a roadside LiDAR monitoring framework that tracks curbside use through three-dimensional point clouds, operating reliably in any lighting condition and preserving privacy by avoiding image collection. Collectively, these studies illustrate a clear shift from static surveillance toward intelligent, sensor-based enforcement systems that combine detection, prediction, and routing to support more efficient, automated, and privacy-aware management of urban parking.

Despite extensive research on parking behavior, analytical models focusing on the optimal strategic spatial allocation of enforcement resources over a long-run planning horizon remain scarce. Existing studies primarily address localized or operational aspects. Nourinejad and Roorda [72] and Nourinejad and et al. [73] develop equilibrium models linking enforcement intensity to driver compliance, using inspector–vehicle meeting functions and search frictions to capture the stochastic nature of detection. Both studies compute behavioral equilibria via an iterative procedure grounded in Brouwer's fixed-point theorem [152]. A limitation in [72] is the assumption of a fixed cost for any legal parking duration, whereas in practice the cost typically increases with time. By contrast, [73] explicitly specifies a utility formulation for legal parking. Notably, Nourinejad [153] demonstrated the superiority of the meeting-function approach over classical game-theoretic formulations such as inspection games [154, 155] emphasizing the latter's inability to capture sensitivity to the number of enforcement units deployed, and therefore its inadequacy for determining the optimal level of enforcement.

The approach proposed by Hernandez and Neeman [156] uses a Bayesian persuasion framework [157] to determine the optimal way an authority (principal) should communicate about the allocation of uncertain enforcement resources. The goal is to maximize the deterrence of unwanted behavior, like illegal parking, across different locations. This is achieved by committing to a randomized communication policy that selects the optimal distribution of Bayes-plausible posterior beliefs for agents (drivers). The persuasion technique is shown to improve deterrence by leveraging the agents' uncertainty. Furthermore, several studies adopt a purely operational perspective, aiming to maximize detection or influence driver behavior for a fixed patrol staff. This includes the work of Lei et al. [86] on optimizing officer routing to influence illegal dwell time, along with related inspection-planning problems [54, 55, 58]. While these routing models are crucial for day-to-day efficiency, they operate on a short-term, hour-by-hour basis and treat the total number of enforcement units as a constant input.

The literature, while rich on localized behavior and enforcement equilibria, lacks a prescriptive framework for *city-wide* allocation of enforcement resources across heterogeneous regions. Most prior studies emphasize steady-state equilibria or inspection/routing games at an operational scale, offering limited guidance on how a municipality should distribute a fixed patrol budget across space over a long-run horizon. In addition, legal-parking cost is often modeled coarsely (e.g., as a fixed charge), whereas in practice it *rises with dwell time* under posted tariffs. As a result, existing models do not fully capture how a constrained patrol budget interacts with spatial variation in demand, dwell times, and detection efficiency.

This paper addresses that gap by formulating an explicit *cost-based* planning model that links driver choice and city allocation under an *increasing legal-parking cost* schedule. Unlike iterative equilibrium approaches (e.g., single-region fixed-point methods), we achieve analytical tractability and strategic integration through two advances:

1. **Closed-form behavioral mechanism.** An exponential-hazard citation model, coupled with cost-minimizing driver choice under time-varying legal cost, yields a closed-form critical staffing threshold $n^\star$ via the Lambert–W function, delineating weak versus effective deterrence regimes.

2. **Strategic integration.** The behavioral mechanism embeds in a nonlinear, multi-region dynamic program that prescribes how to allocate a fixed patrol budget across regions. The structure directly accommodates managerial requirements, such as equity

or minimum-coverage constraints, and it can prioritize specific area (e.g., school zones, freight streets) via region weights when desired.

By unifying the threshold $n^\star$, enforcement intensity, and spatial allocation, the framework provides a tractable tool for optimizing deterrence and revenue at the tactical scale. It also quantifies the value of technology (through changes in detection efficiency) and establishes a foundation for extensions that incorporate broader policy objectives (e.g., social costs, fairness, and compliance equity).

The rest of the paper proceeds as follows. Section 6.2 introduces the enforcement and detection mechanisms governing driver behavior under probabilistic patrols. Section 6.3 formulates the city's long-run revenue optimization problem and describes the dynamic programming framework for strategic resource allocation. Section 6.4 presents the experimental setup and computational results using synthetic and real urban regions that reflect heterogeneous parking environments. Finally, Section 6.5 concludes with key managerial insights and potential avenues for future research.

6.2 Enforcement and Detection of Illegal Drivers

A municipality divides its curbside network into distinct regions, each patrolled by one or more officers. Officers circulate continuously and can issue citations to illegally parked vehicles encountered during inspection.

6.2.1 Patrol and Detection Mechanics

Let s denote patrol speed (km/min), m the curb length of a region (km), and $d \in (0,1]$ the detection probability per pass, capturing both technology effectiveness and officer attentiveness. Following [64], the pass-by detection model treats patrol passes as opportunities to "intercept" an observable event with probability d. With n officers assigned to the region, the aggregate detection (citation) hazard per minute is

$$\lambda = \kappa n \quad \text{where} \quad \kappa \triangleq \frac{s\,d}{m}.$$

The coefficient κ is a composite efficiency term: faster patrols (s), higher per-pass acuity (d), or shorter curb length (m) all raise the effective inspection intensity per officer.

Although Chelst's original formulation targeted crime interception, the same pass-by mechanism applies to on-street parking enforcement with minimal reinterpretation: an "interception" is a *citation* event when an illegally parked vehicle is observed by a passing patrol.

Under standard assumptions, independent passes, stationarity over the short horizon of a single stay, and homogeneous coverage within the region, the citation arrival process is Poisson with rate $\lambda = \kappa n$.

A more general formulation could include a demand or congestion adjustment factor $\mu \geq 1$ to reflect that higher curb occupancy reduces per-pass visibility (e.g., denser regions may dilute patrol effectiveness). However, in our framework, regional differences in occupancy and parking activity are already represented explicitly through the demand parameter D_i and the dwell-time distribution a_i at the allocation stage. Accordingly, we set $\mu = 1$ to avoid redundant scaling and to keep κ interpretable in terms of directly measurable physical characteristics (patrol speed, curb length, and detection efficiency). This isolates behavioral and allocation effects without introducing an additional calibrated parameter. If an illegally parked driver remains for t minutes, the probability of receiving a citation (single-citation regime) is

$$P^{\text{cite}}(t \mid n) \;=\; 1 - e^{-\lambda t} \;=\; 1 - \exp(-\kappa n\, t). \tag{6.1}$$

We assume parking durations are exponentially distributed and citation opportunities follow a Poisson process, yielding the memoryless hazard used in (6.1) and enabling closed-form thresholds and revenue expressions. These assumptions are standard in analytical parking and curbspace models and queueing formulations [116, 158], while empirical work shows that real parking durations can vary by context and often warrant parametric survival modeling [159]. The exponential formulation is adopted for analytical tractability at the tactical planning scale.

This exponential-hazard structure implies a time-invariant enforcement risk, meaning each additional minute parked illegally carries the same instantaneous risk of detection. Drivers observe patrol conditions and adjust their parking decisions accordingly. Practices for reticketing vary by jurisdiction, but in many settings a first citation triggers prompt behavioral adjustment: drivers typically depart, legitimize the stay (e.g., pay at the meter), or otherwise cease the illegal episode shortly after being cited. We therefore adopt a first-citation departure assumption (single-citation regime), which captures the principal deterrence and revenue effects observed in practice while avoiding locality-specific calibration (e.g., reticketing cooldowns or maximum daily fines).

We work at a tactical (long-run average) level, where region-level inputs summarize typical conditions over a planning horizon (e.g., a week or month). The allocation n_i should thus be interpreted as an *average deployment* per region. Operational, within-day scheduling (e.g., peak vs. off-peak patrols, weekday vs. weekend adjustments) is handled separately once these baseline staffing levels are determined.

The next subsection introduces the competing parking options and the behavioral thresholds that govern switching between them.

6.2.2 Driver Options and Behavioral Thresholds

Each driver faces three mutually exclusive options:

- **(i) Illegal parking:** risk a fine F under patrol intensity $\lambda = \kappa n$.
- **(ii) Metered parking:** pay $r = p/60$ per minute with a fixed overhead K (search and payment time cost).
- **(iii) Daily pass:** pay a lump sum Dp, valid for unlimited stay.

Several cities (including Montréal[1]) offer day-long visitor stickers that permit parking in residents-only zones; availability is borough-specific. We include the daily-pass option to keep the model policy-complete, but it is typically attractive only for long stays and thus contributes little revenue.

$$\text{Illegal cost: } C^{\text{ill}}(t;n) = F\left(1 - e^{-\kappa n t}\right), \qquad \text{Parking Metered cost: } C^{\text{leg}}(t) = K + rt,$$
$$\text{Daily pass cost: } C^{\text{pass}}(t) = Dp.$$

The illegal cost $C^{\text{ill}}(t;n)$ is *concave* and increasing in t, bounded above by the maximum fine F:

$$\frac{\partial C^{\text{ill}}}{\partial t} = F\kappa n e^{-\kappa n t} > 0, \qquad \frac{\partial^2 C^{\text{ill}}}{\partial t^2} = -F(\kappa n)^2 e^{-\kappa n t} < 0,$$

which implies strict concavity for $n > 0$. Equivalently, since $e^{-\kappa n t}$ is convex, $-Fe^{-\kappa n t}$ and thus $F(1 - e^{-\kappa n t})$ are concave. The metered cost $C^{\text{leg}}(t) = K + rt$ is linear (both convex and concave), and the daily-pass cost $C^{\text{pass}}(t) = Dp$ is constant.

Intuitively, the concavity of $C^{\text{ill}}(t;n)$ reflects diminishing marginal expected penalties as stay duration increases: the longer a driver remains parked, the smaller the additional detection risk per extra minute, since the cumulative probability of being caught saturates toward one.

We assume risk-neutral, cost-minimizing drivers with a reasonable perception of patrol intensity, as is typical on a tactical (long-run average) horizon. Given their realized stay t and perceived deployment $\bar{n}$, they choose the least-cost option, which induces threshold behavior in t.

[1] https://montreal.ca/en/how-to/get-daily-or-monthly-parking-sticker

6.2.3 Critical Staffing Level and the Lambert–W Expression

The indifference between illegal and legal parking options occurs when

$$C^{\text{ill}}(t;n) = C^{\text{leg}}(t) \quad \Longleftrightarrow \quad D(t) = (K - F) + rt + Fe^{-\kappa n t} = 0.$$

Because $D(t)$ is convex in t, tangency (a single crossing) arises when

$$D(t^*) = D'(t^*) = 0, \qquad \text{so that} \quad e^{-\kappa n t^*} = \frac{r}{F\kappa n}.$$

Letting $x = e^{-\kappa n t^*}$ and $\alpha = 1 - \frac{K}{F}$ gives

$$x - x \ln x = \alpha,$$

whose relevant solution lies on the W_{-1} branch of the Lambert function. Substituting back yields the *critical continuous staffing level*

$$n_{\text{crt}} = -\frac{r}{F\kappa\alpha}\, W_{-1}\!\left(-\tfrac{\alpha}{e}\right), \qquad \alpha = 1 - \tfrac{K}{F}.$$

Since patrol units are discrete, the operational threshold is

$$n^\star = \lfloor n_{\text{crt}} \rfloor + 1,$$

the smallest integer number of officers sufficient to induce at least one intersection between C^{ill} and C^{leg}.

For $n < n_{\text{crt}}$, the illegal option always dominates (*weak deterrence*); for $n > n_{\text{crt}}$, two intersections appear, corresponding to distinct behavioral regimes in which short-stay drivers may still park illegally while long-stay drivers comply.

To maintain monotone driver responses, the daily pass price D_p is assumed to be chosen within the interval defined by the two intersection costs at $n^\star$that is,

$$C^{\text{leg}}(t_1) = C^{\text{ill}}(t_1; n^\star) < D_p < C^{\text{leg}}(t_2) = C^{\text{ill}}(t_2; n^\star),$$

ensuring that the overall cost ordering (illegal → meter → pass) remains consistent as dwell time increases.

6.2.4 Driver Switching Logic

Let $t_1(n)$ and $t_2(n)$ denote the two solutions of $C^{\text{ill}}(t;n) = C^{\text{leg}}(t)$ when they exist, and define $t_{\text{pass}} = (Dp - K)/r$. Drivers follow a monotone rule:

$$\arg\min\{C^{\text{ill}}, C^{\text{leg}}, C^{\text{pass}}\} = \begin{cases} \text{Illegal}, & 0 \le t < t_1(n), \\ \text{Meter}, & t_1(n) \le t < t_{\text{pass}}, \\ \text{Pass}, & t \ge t_{\text{pass}}. \end{cases}$$

If $n < n^\star$, there is no crossing, and the meter region vanishes. The choice collapses to illegal vs. pass, with the threshold

$$t_{Dp}(n) = \frac{1}{\kappa n} \ln\!\left(\tfrac{F}{F - Dp}\right),$$

indicating the minimum stay at which purchasing a daily pass becomes cheaper than risking enforcement.

6.2.5 Behavior Under Different Enforcement Level

When staffing is insufficient, $C^{\text{ill}}(t;n) < C^{\text{leg}}(t)$ for all t, so metered parking is never optimal. The only switch occurs at $t_{Dp}(n)$, beyond which drivers prefer a pass. This threshold increases with Dp and decreases with n:

$$\frac{\partial t_{Dp}}{\partial n} < 0, \quad \frac{\partial t_{Dp}}{\partial Dp} > 0.$$

Hence, weaker enforcement or more expensive passes expand the illegal regime.

Figure 6.1 illustrates how a driver's cost evolves under varying enforcement intensities. This example highlights the dynamic interplay between enforcement strength and driver parking choices, showing how changes in patrol allocation directly influence behavioral responses. The critical staffing level is $n^\star = 10$, representing the minimum number of patrol units required for the illegal parking cost curve (orange line) to intersect with the metered parking cost (black dashed line). When enforcement is below this threshold (e.g., $n = 9$), no intersection occurs, the illegal option consistently yields a lower expected cost, as shown by the blue curve. This corresponds to an under-enforced regime where deterrence is ineffective and drivers find it optimal to park illegally regardless of dwell time.

As enforcement increases beyond the critical level (e.g., $n = 13$), the illegal parking curve becomes steeper and crosses the metered parking cost earlier. This means that the expected cost of remaining illegally parked rises faster, inducing drivers with moderate or long stays to

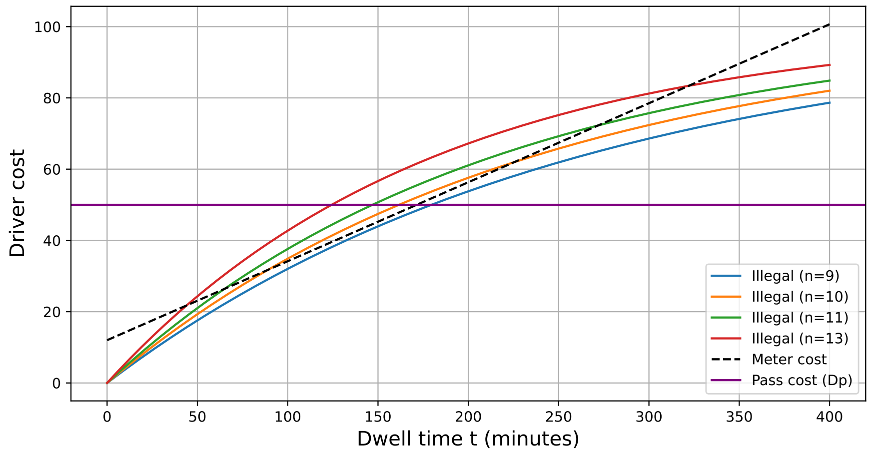

Figure 6.1 Driver cost curves under varying enforcement levels and dwell times ($n^{\star}$=10)

switch to legal alternatives. The purple horizontal line represents the daily pass cost, which remains constant regardless of stay duration. In this example, the daily pass becomes the optimal choice for drivers with long expected stays, as it caps their total cost below that of illegal or metered parking.

6.3 City Revenue Optimization Problem in the Long Run

Let dwell time $T \sim \mathrm{Exp}(1/a)$ with mean a. The city's expected revenue per driver combines citations, meter payments, and passes:

$$\mathrm{Rev}(n) = \int_0^{\infty} \min\{C^{\mathrm{ill}}(t;n),\, C^{\mathrm{leg}}(t),\, C^{\mathrm{pass}}(t)\} \frac{1}{a} e^{-t/a}\, dt.$$

Depending on the enforcement regime, this integral reduces to:

(A) Low enforcement (two-option rule). The expected revenue thus becomes the integral of illegal cost up to t_{Dp} and the pass payment afterward:

$$\mathrm{Rev}^{(A)}(n) = \int_0^{t_{Dp}} F(1 - e^{-\kappa n t}) \frac{1}{a} e^{-t/a}\, dt + Dp\, e^{-t_{Dp}/a}.$$

Carrying out the integration yields the closed-form expression:

$$\mathrm{Rev}^{(A)}(n) = F\Big[(1 - e^{-t_{Dp}/a}) - \tfrac{1 - e^{-(\kappa n + 1/a) t_{Dp}}}{1 + \kappa n a}\Big] + Dp\, e^{-t_{Dp}/a}.$$

(B) High enforcement (three-option rule). The expected revenue decomposes into three integrals:

$$\mathrm{Rev}^{(B)}(n) = \int_0^{t_1} F(1 - e^{-\kappa n t}) \frac{1}{a} e^{-t/a} \, dt + \int_{t_1}^{t_{\text{pass}}} r t \frac{1}{a} e^{-t/a} \, dt + Dp \, e^{-t_{\text{pass}}/a}.$$

Evaluating these integrals analytically gives

$$\mathrm{Rev}^{(B)}(n) = F\Big[(1 - e^{-t_1/a}) - \tfrac{1 - e^{-(\kappa n + 1/a)t_1}}{1 + \kappa n a}\Big] + r\Big[-(t + a)e^{-t/a}\Big]_{t=t_1}^{t=t_{\text{pass}}} + Dp \, e^{-t_{\text{pass}}/a}.$$

Both (A) and (B) therefore admit closed-form solutions, enabling direct sensitivity analysis with respect to fine levels, detection intensity, or dwell-time mean a without numerical integration. The model excludes driver search cost from the city's revenue function, as it reflects user inconvenience rather than municipal income.

Higher staffing levels (n) or technological improvements that increase the effective detection rate (κ) raise expected citation revenue up to a saturation point, beyond which additional enforcement yields diminishing returns. Once enforcement becomes sufficiently strong, most drivers anticipate detection and thus switch to legal or pass-based options, reducing the pool of violators. In contrast, under prolonged low enforcement ($n < n^\star$), the system collapses into a two-option regime, illegal parking versus the daily pass, resulting in widespread noncompliance and limited municipal revenue.

This transition across the critical threshold $n^\star$ constitutes a fundamental behavioral bifurcation: below $n^\star$, deterrence is weak and illegal parking dominates; above $n^\star$, compliance emerges through two observable thresholds separating illegal, metered, and pass-paying drivers. Hence, $n^\star$ represents not only an operational target for enforcement agencies but also a behavioral tipping point in the parking ecosystem.

Figure 6.2 (a) illustrates the city's expected revenue as a function of the number of officers n and the mean parking duration (the parameter a of the exponential dwell-time distribution). The vertical dashed line marks the critical staffing level $n^\star$. Before this threshold, revenue increases with n, as every additional officer raises detection frequency and fine collection. After $n^\star$, however, the marginal revenue depends strongly on driver behavior: longer average stays (a larger) shift the revenue curve upward because both legal payments and fines accumulate over longer durations, whereas shorter stays reduce overall exposure to enforcement and consequently diminish revenue. This relationship highlights that both enforcement intensity and parking duration jointly shape the equilibrium between compliance and violation.

Figure 6.2 (b) illustrates how changes in κ alter the equilibrium between enforcement and driver behavior. The figure illustrates that improvements in detection technology or patrol speed allow the city to achieve effective deterrence with fewer officers. As shown in the revenue curves, higher κ values shift the revenue peak toward smaller n and flatten its left-hand portion, indicating an earlier transition from low to high deterrence. Conversely, when κ decreases (e.g., due to slower patrols or limited detection capability), more officers are needed to maintain compliance, and total revenue declines as illegal parking persists longer (black dot in figure shows $n^\star$). Hence, κ drives a nonlinear trade-off between technology and manpower: a larger κ compresses the enforcement range needed for optimal revenue, whereas a smaller κ widens it and delays behavioral shifts.

Overall, the figure reveals two key insights: (i) drivers' dwell-time distribution critically affects the optimal staffing level, and (ii) technological or behavioral changes that increase κ or a can shift the revenue peak to the left, allowing cities to maintain high revenue with fewer officers.

6.3.1 Effect of Fine Parameter on Optimal Revenue

The optimal revenue per driver, $\mathrm{Rev}_{\mathrm{opt}}(F)$, represents the upper envelope of all feasible staffing configurations:

$$\mathrm{Rev}_{\mathrm{opt}}(F) = \max_{n \in \mathcal{N}} \mathrm{Rev}(F, n), \qquad n_{\mathrm{opt}}(F) \in \arg\max_{n} \mathrm{Rev}(F, n),$$

where $\mathcal{N}$ denotes the admissible staffing levels. The corresponding *optimizer line* $F \mapsto \mathrm{Rev}_{\mathrm{opt}}(F)$ traces the best attainable revenue given the city's fine level (see Figure 6.3).

The function $\mathrm{Rev}_{\mathrm{opt}}(F)$ is generally *non-monotonic* in F. Counterintuitively, increasing the fine level F does not guarantee higher revenue; it may even lead to a local decline in $\mathrm{Rev}_{\mathrm{opt}}(F)$. This behavior arises from the piecewise structure of $\mathrm{Rev}(n)$ and the discontinuous shifts in the optimal staffing level $n_{\mathrm{opt}}(F)$, which depend on the regime boundary $n^\star(F)$. As F increases, the critical staffing threshold $n^\star(F)$ decreases, potentially triggering regime changes and discontinuous jumps in the optimal solution. Thus, while individual terms (e.g., citation revenue) tend to rise with F, the total expected revenue may fall if deterrence effects dominate—fewer violations, earlier switching to paid options, or over-deterrence reducing enforcement yield.

Each fixed-n revenue curve $\mathrm{Rev}(F \mid n)$ is smooth but may exhibit local concavity and turning points due to competing effects: higher fines increase the payout per citation but simultaneously reduce the illegal-parking probability and thus the number of citations. The optimizer line $\mathrm{Rev}_{\mathrm{opt}}(F)$ is the pointwise maximum over these fixed-n curves, forming a continuous,

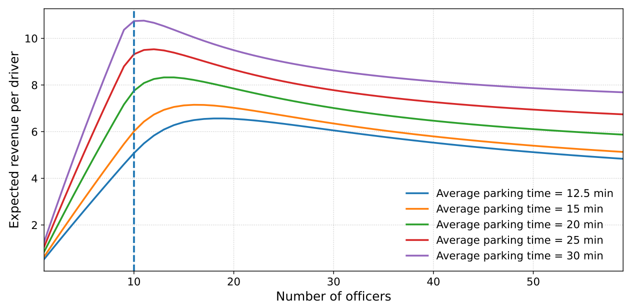

(a)

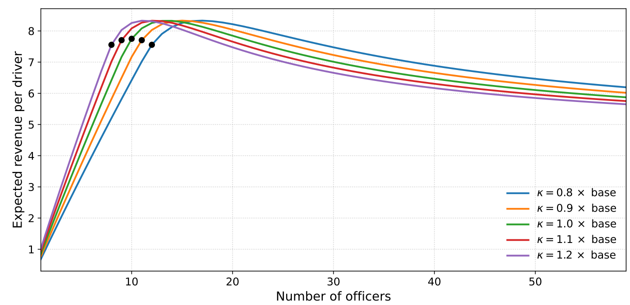

(b)

Figure 6.2 Expected municipal revenue under varying enforcement levels and (a) mean parking durations ($n^{\star}$=10) (b) detection technology factor (κ)

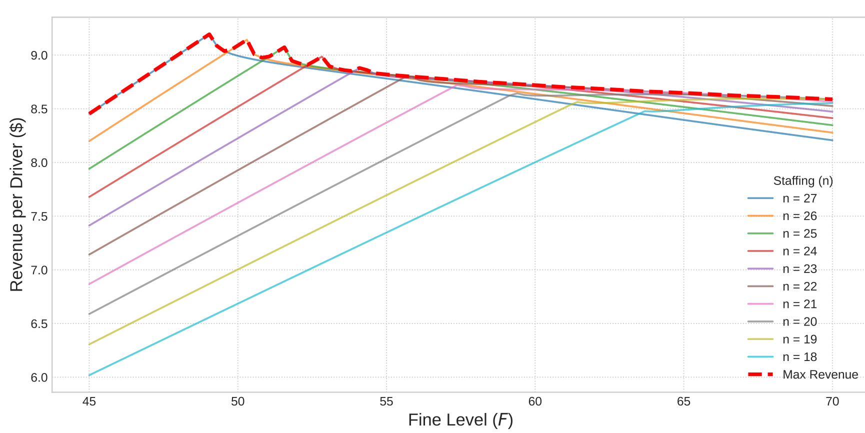

Figure 6.3 Revenue per driver under different level of fine

piecewise-smooth envelope with visible kinks at fine levels where the optimal staffing level $n_{\text{opt}}(F)$ switches.

A particularly noticeable example of this is the significant jaggedness and sudden vertical steps observed in the individual fixed-n curves around $F \approx \$50$ in Figure 6.3. This is the point where the critical staffing threshold $n^\star(F)$ drops, causing the fixed staffing levels (e.g., $n = 27, 26$) to suddenly switch from the lower-yielding Regime A to the higher-yielding Regime B. This regime-switching discontinuity is a direct result of $n^\star(F)$ being defined as a discrete integer, which creates steps that the continuous $\text{Rev}_{\text{opt}}(F)$ envelope smooths but must follow.

From a policy standpoint, the non-monotonicity of $\text{Rev}_{\text{opt}}(F)$ underscores that simply raising fines is not a reliable strategy for maximizing revenue. Beyond a certain range, the system transitions into a high-deterrence, low-citation regime, reducing marginal returns from further fine increases. Therefore, optimal fine calibration should be conducted jointly with staffing optimization, balancing deterrence effectiveness and operational yield.

Proposition 1 (Possibility of Multiple Regime Shifts)**.** *The global optimizer of the piecewise revenue function*

$$R(F,n) = \begin{cases} \text{Rev}^{(A)}(F,n), & n < n^\star(F), \\ \text{Rev}^{(B)}(F,n), & n \geq n^\star(F), \end{cases}$$

may switch non-monotonically between regimes as the fine level F varies, producing transitions such as B$\to$A$\to$B.

Proof of Possibility. Let $n^\star(F)$ denote the tangency (critical staffing) level where $C^{\text{ill}}(t;n)$ and $C^{\text{leg}}(t)$ are tangent. From the closed form $n^\star(F) = -\frac{r}{F\kappa(1-K/F)} W_{-1}\Big(-\frac{1-K/F}{e}\Big)$, one verifies that $n^\star(F)$ is *strictly decreasing* in F. Hence, as the fine rises, the feasible domain of Regime B ($n \geq n^\star(F)$) *expands*, while the domain of Regime A ($n < n^\star(F)$) contracts.

Define the within-regime optima

$$A^*(F) = \max_{n<n^\star(F)} \text{Rev}^{(A)}(F,n), \qquad B^*(F) = \max_{n\geq n^\star(F)} \text{Rev}^{(B)}(F,n),$$

so that the global maximum revenue is $\max\{A^*(F), B^*(F)\}$ and the corresponding optimizer $n_{\text{opt}}(F)$ lies in the regime attaining this maximum.

Because both $\text{Rev}^{(A)}$ and $\text{Rev}^{(B)}$ depend nonlinearly on F (e.g., via F, Dp, and κn) and the boundary $n^\star(F)$ itself shifts with F, the difference $\Delta(F) = A^*(F) - B^*(F)$ need not be monotone. Multiple sign changes of $\Delta(F)$ produce multiple regime switches. In particular, suppose there exist $F_1 < F_2$ such that

$$B^*(F_1) > A^*(F_1), \qquad A^*(F_2) > B^*(F_2),$$

then the optimal regime shifts from B at F_1 to A at F_2, yielding a B $\to$ A transition. A further crossing of $\Delta(F)$ at some $F_3 > F_2$ would restore Regime B, giving B$\to$A$\to$B.

These non-monotonic shifts arise from (i) the decreasing boundary $n^\star(F)$, which alters each regime's feasible range, and (ii) the non-concavity of the revenue functions, which allows their envelopes to intersect multiple times. Integer-valued staffing n can accentuate the discontinuities. $\square$

6.3.2 City Staffing Model Across Regions

The city is administratively divided into regions $i \in I$, each with D_i expected parking events per period and mean dwell time a_i ($T_i \sim \text{Exp}(1/a_i)$).

Each region has a per-minute parking rate r_i, a fixed overhead K_i, a daily pass price Dp_i, and a fine F_i. Enforcement efficiency κ_i captures patrol speed, detection probability, and curb length. Allocating n_i officers to region i yields patrol intensity $\lambda_i = \kappa_i n_i$.

Expected revenue per driver. With exponential dwell density $f_i(t) = (1/a_i)e^{-t/a_i}$:

$$\text{Rev}_i^{\text{per}}(n_i) = \begin{cases} F_i\left[(1 - e^{-t_{i,Dp}/a_i}) - \dfrac{1 - e^{-(\kappa_i n_i + 1/a_i)t_{i,Dp}}}{1 + \kappa_i n_i a_i}\right] \\ \qquad + Dp_i\, e^{-t_{i,Dp}/a_i}, & n_i < n_i^\star, \\ F_i\left[(1 - e^{-t_{i,1}/a_i}) - \dfrac{1 - e^{-(\kappa_i n_i + 1/a_i)t_{i,1}}}{1 + \kappa_i n_i a_i}\right] \\ \qquad + r_i\Big[-(t + a_i)e^{-t/a_i}\Big]_{t_{i,1}}^{t_{i,\text{pass}}} + Dp_i\, e^{-t_{i,\text{pass}}/a_i}, & n_i \geq n_i^\star. \end{cases}$$

Expected revenue per region is $\text{Rev}_i(n_i) = D_i\,\text{Rev}_i^{\text{per}}(n_i)$. We interpret D_i as short-run parking demand at the tactical planning scale, meaning travelers still make the trip to the region and only their parking choice (legal, illegal, or pass) responds to enforcement intensity.

City allocation problem. The municipality allocates limited officers B across all regions to maximize total expected revenue:

$$\max_{\{n_i \in \mathbb{Z}_+\}} \sum_{i \in I} \mathrm{Rev}_i(n_i) \quad \text{s.t.} \quad \sum_{i \in I} n_i \leq B.$$

We treat B as a fixed tactical staffing level determined outside the model (e.g., through annual budgeting or labor agreements). The optimization therefore allocates a given number of officers across regions to maximize long-run revenue. Incorporating the salary or labor cost of officers would shift the objective from revenue to net revenue. In that case, the model could be used to determine the optimal total number of officers B itself. However, choosing that overall staffing level is a higher-level strategic planning decision and is outside the scope of this tactical deployment model.

6.3.3 Equity on Average Detection Probability

To avoid systematically over-enforcing some neighborhoods while neglecting others, we impose an equity constraint on enforcement coverage. Specifically, we require that the average illegally parked driver in each region experiences at least a minimum detection probability, ensuring a baseline level of deterrence citywide. Basing equity on expected detection probability ties fairness directly to model inputs ($\kappa_i n_i$ and a_i) and avoids reliance on rare or extreme outcomes. This approach is motivated by findings that spatial disparities in parking citations often reflect structural differences in curb supply, commercial density, and overlapping patrol routes rather than differences in behavior alone [47]. By enforcing a consistent minimum level of enforcement presence, the model helps reduce the risk that fines disproportionately burden specific communities while preserving flexibility to allocate additional officers where they are most effective.

In our framework, equity is operationalized through the detection probability experienced by an *average driver*, defined as a representative driver with mean parking duration.. Let the mean parking duration in region i be a_i. The probability that an illegally parked vehicle is detected before leaving, given n_i officers and patrol efficiency κ_i, is

$$\bar{P}_i(n_i) = 1 - e^{-\kappa_i n_i a_i}.$$

To guarantee equitable coverage, the city enforces a minimum threshold $\rho \in (0, 1)$ such that

$$\bar{P}_i(n_i) \geq \rho, \qquad \forall i \in I.$$

This ensures that drivers, regardless of location, face a consistent minimum probability of inspection under average parking behavior.

City problem with average detection equity. The staffing allocation problem becomes

$$\max_{\{n_i \in \mathbb{Z}_+\}} \sum_{i \in I} \mathrm{Rev}_i(n_i) \quad \text{s.t.} \quad \sum_{i \in I} n_i \leq B, \qquad 1 - e^{-\kappa_i n_i a_i} \geq \rho \;\; \forall i \in I.$$

Equivalently, the equity constraint can be written as an integer lower bound:

$$n_i \geq \left\lceil \frac{-\ln(1-\rho)}{\kappa_i a_i} \right\rceil, \qquad \forall i \in I.$$

6.3.4 Dynamic Programming Formulation and Interpretation

The staffing problem can be viewed as a discrete nonlinear resource allocation problem. The municipality has a limited budget B enforcement units (officers) and must decide how many units n_i to assign to each region $i \in I$, where each region's expected revenue $\mathrm{Rev}_i(n_i)$ depends nonlinearly on staffing. Because $\mathrm{Rev}_i(\cdot)$ is typically increasing but not concave, gradient-based optimization is impractical. Dynamic programming (DP) offers an exact and interpretable approach to solving this class of problems.

Baseline model (without equity). At each stage i, representing region i, the decision variable is the number of officers n_i allocated to that region, subject to $0 \leq n_i \leq b$, where b is the number of remaining officers. The state variable (b, i) captures the remaining budget and the current region under consideration. Let $V(b, i)$ denote the maximum total revenue attainable from regions $i, i+1, \ldots, |I|$ when b officers remain. The Bellman recursion is

$$V(b, i) = \max_{0 \leq n_i \leq b} \Big\{ \mathrm{Rev}_i(n_i) + V(b - n_i,\, i+1) \Big\}, \qquad V(b, |I|+1) = 0.$$

This recursion enumerates all feasible budget splits among the remaining regions, caching intermediate results to avoid recomputation. The optimal staffing vector $(n_1^{\mathrm{opt}}, n_2^{\mathrm{opt}}, \ldots, n_{|I|}^{\mathrm{opt}})$ is obtained by backtracking from the final state $(B, 1)$.

DP with equity constraint To incorporate the equity constraints efficiently, pre-allocate the mandatory minimum units L_i and define residual variables $m_i \geq 0$ with $n_i = L_i + m_i$. Let $B' = B - \sum_i L_i \geq 0$ be the residual budget available for discretionary allocation. Define the incremental revenue function

$$\widetilde{\mathrm{Rev}}_i(m_i) = \mathrm{Rev}_i(L_i + m_i) - \mathrm{Rev}_i(L_i).$$

The residual DP recursion becomes

$$\widetilde{V}(b, i) = \max_{0 \le m_i \le b} \Big\{ \widetilde{\mathrm{Rev}}_i(m_i) + \widetilde{V}(b - m_i,\, i+1) \Big\}, \qquad \widetilde{V}(b, |I|+1) = 0.$$

The optimal staffing levels are recovered as $n_i^{\mathrm{opt}} = L_i + m_i^{\mathrm{opt}}$. This formulation ensures that the minimum detection probability constraint for the average parker is always satisfied while maintaining the computational efficiency and structure of the standard DP algorithm.

Complexity and tractability. For $|I|$ regions and total budget B, the computational complexity is $O(|I|Bn_{\max})$, where $n_{\max}$ is the maximum number of officers assignable to a region (often $n_{\max} = B$) [160]. This is tractable for moderate-scale instances (tens of regions and budgets of a few hundred officers). For large-scale systems, decomposition, Lagrangian relaxation, or metaheuristics can approximate the DP solution while maintaining the equity feasibility guarantees.

6.4 Computational Experiments

We evaluate the model in two stages: (i) a *synthetic* five-region city designed to stress-test allocation behavior and conduct sensitivity analyses, and (ii) a *real* case study for Montréal (Plateau–Mont-Royal) that uses empirically grounded parameters. For the synthetic setting, we first describe how the instance is generated (regions, shares, dwell, search costs, and patrol encounter rates), then report the resulting allocations and revenues. The real case maps the same modeling pipeline to observed demand, network length, and tariff/fine levels.

6.4.1 Synthetic Instance Generation

To evaluate the allocation model in a controlled environment, we generate a single five-region synthetic city using one random seed to ensure reproducibility. The regions, *Downtown Core*, *Business District*, *Mixed Residential–Retail*, *Entertainment Area*, and *Suburban Zone*, capture contrasting parking environments. For each region i, we specify:

- the share of street length and daily demand (s_i, w_i),
- an average dwell time a_i,

- a search cost K_i,
- a patrol encounter rate $\kappa_i = \frac{v_{\text{patrol}} p_{\text{det}}}{L_i}$ from patrol speed v_{patrol}, detection probability p_{det}, and effective length L_i.

Meter rate r, fine F, and daily pass price D_p are held constant across regions with mild noise to reflect heterogeneity. All parameters are drawn once (single seed) and kept fixed across experiments to isolate policy effects. The dynamic program then allocates a total staff budget B across regions to maximize expected revenue, optionally subject to equity constraints.

6.4.2 Baseline Experiment: Synthetic Multi-Region Allocation

We apply the model to the generated instance from Section 6.4.1. Table 6.1 reports the parameters used to compute equilibrium outcomes, and Table 6.2 summarizes the optimal allocation and revenue composition under $B = 70$.

Table 6.1 Synthetic regional characteristics and critical enforcement levels.

Region	L_i [m]	D_i	κ_i	a_i [min]	r_i [\$/min]	F_i	K_i	$n_i^\star$
Downtown	11,487	2,020	2.54e–4	21.5	0.217	324.0	20	4
Business	31,578	2,025	9.2e–5	53.3	0.198	297.8	20	11
MixedRR	35,032	1,351	8.3e–5	72.7	0.186	305.9	20	12
Entertainment	14,064	1,077	2.07e–4	90.1	0.185	310.3	20	5
Suburban	43,840	531	6.7e–5	180.3	0.114	294.5	20	9

Shorter dwell times and compact networks in *Downtown*/*Entertainment* produce higher encounter rates (κ_i) and lower $n_i^\star$, whereas dispersed areas (*Business*, *MixedRR*) require more staff to deter effectively. *Suburban* exhibits low demand but wide coverage.

Table 6.2 Optimal allocation and revenue composition under baseline parameters ($B = 70$).

Region	n_i	Total Rev [$]	Citation	Meter	Pass	Rev/Driver [$]
Downtown	9	21,861	17,748	4,113	0	10.8
Business	15	34,203	25,668	8,508	27	16.9
MixedRR	14	26,947	19,120	7,823	4	20.0
Entertainment	5	25,088	18,084	6,997	7	23.3
Suburban	9	14,739	12,698	1,614	427	27.8

The optimizer concentrates staff where marginal gains remain high (*Business*, *MixedRR*), while *Downtown*/*Entertainment* show diminishing returns beyond $n_i^\star$. Citywide revenue totals \$122,838 using 52 of 70 available units, indicating rapidly diminishing marginal returns once deterrence thresholds are reached.

6.4.3 Sensitivity to Fine Levels

To examine how the level of fines influences both driver and agency behavior, we conducted a fine-sensitivity analysis by varying the unit fine F from 50 to 1000. For each level of F, the equilibrium detection threshold, staffing allocation, and resulting revenue composition were recomputed under identical regional conditions.

Table 6.3 summarizes key performance indicators. As fines increase, fewer patrol units are required to induce compliance. At low fine levels ($F < 100$), almost all revenue originates from citations, as deterrence is weak and illegal parking remains dominant. Intermediate fines ($F \approx 200$–400) produce the highest total revenue, where both enforcement and voluntary compliance coexist. Beyond this point, higher fines over-deter illegal parking, reducing citation frequency and total income despite the higher penalty per violation.

Table 6.3 Summary of total revenue and staffing response to increasing fine F.

F	n_{total}	Total Rev [\$]	Citation Share [%]	Meter Share [%]	Main Behavior
50	70	39,861	99.9	0.0	Fully illegal
100	70	76,598	99.8	0.0	Fully illegal
150	70	104,415	99.2	0.6	Fully illegal
200	70	123,202	92.9	6.0	Weak deterrence
300	52	122,838	84.7	13.1	Weak deterrence
400	39	121,672	78.1	21.8	Partial deterrence
600	28	120,164	69.1	30.9	Partial deterrence
1000	17	119,050	58.3	41.7	Balanced mix

Overall, the results indicate that excessively high fines can paradoxically reduce municipal revenues by discouraging illegal parking to the extent that enforcement generates little marginal return. This adjustment reflects a strategic substitution between monetary and operational deterrence: higher penalties reduce the need for frequent patrols, allowing the municipality to preserve revenues by balancing legal and illegal contributions across regions.

Switching Behavior from Illegal to Legal Parking

To further explore the behavioral response of drivers to enforcement and monetary deterrence, we examine how the switching time from illegal to legal parking varies with the level of the fine F. This time, denoted as $t_{\text{switch}}(\text{ill} \to \text{meter})$, represents the threshold dwell duration after which a rational driver chooses to comply (i.e., pay for parking) rather than risk receiving a citation.

Table 6.4 reports the observed switching times across regions under different fine levels, along with the number of enforcement units allocated to each region in parentheses. This joint presentation highlights the interplay between enforcement allocation and monetary penalties.

Table 6.4 Switching time from illegal to legal parking and allocated officers per region in () as a function of fine F.

Fine F	Downtown (n)	Business (n)	MixedRR (n)	Entertainment (n)	Suburban (n)
50	— (31)	— (13)	— (0)	— (26)	— (0)
100	127.5 (18)	— (27)	— (6)	— (19)	— (0)
150	94.5 (12)	— (27)	— (21)	— (10)	— (0)
200	59.0 (11)	171.4 (20)	— (18)	— (7)	— (14)
300	40.8 (9)	107.8 (15)	139.3 (14)	185.4 (5)	555.6 (9)
400	38.1 (7)	108.5 (11)	149.5 (10)	144.3 (4)	370.8 (7)
500	44.5 (5)	101.0 (9)	144.7 (8)	165.4 (3)	580.0 (5)
600	34.0 (5)	87.5 (8)	126.4 (7)	105.9 (3)	272.9 (5)
700	37.3 (4)	114.0 (6)	124.9 (6)	192.7 (2)	327.2 (4)
800	46.6 (3)	86.2 (6)	138.4 (5)	134.2 (2)	225.1 (4)
900	39.0 (3)	97.2 (5)	106.3 (5)	103.6 (2)	355.8 (3)
1000	33.6 (3)	79.2 (5)	136.6 (4)	84.5 (2)	260.0 (3)

At low fine levels ($F \leq 50$), switching never occurs, and drivers consistently choose illegal parking since the expected cost of being caught remains lower than the cost of compliance. Enforcement allocation is consequently concentrated in high-activity zones such as Downtown and Entertainment, yet deterrence is insufficient to induce behavioral change.

When fines reach moderate levels ($F \approx 100$–300), switching behavior begins to emerge in the core areas. The threshold t_{switch} decreases substantially in Downtown and Business districts, indicating that drivers start paying for shorter stays rather than risking a citation. Simultaneously, enforcement allocation becomes more evenly distributed, suggesting an adaptive redistribution effect across regions.

It is noteworthy that, for cases with a fixed level of allocated units (e.g., $F = 500$–600 with five officers assigned to Downtown), increasing the fine level reduces the duration of illegal parking and alters driver behavior as they adjust to the city's enforcement policy (from 44.5 to 34.0 minutes). However, when fines increase further (e.g., $F = 700$), the city's response changes, fewer officers are deployed to maintain revenue efficiency. This adjustment creates a new equilibrium point: with four officers, drivers reoptimize their parking behavior, leading to a longer illegal stay duration (37.3 minutes) under the new enforcement configuration. A similar interpretation applies to other regions, where interactions between fine levels and staffing adjustments jointly determine the equilibrium duration of illegal parking and overall compliance dynamics.

Overall, these results reveal that fine intensity and enforcement capacity interact nonlinearly: modest fines require strong enforcement to ensure compliance, whereas high fines can maintain deterrence even with limited patrol resources. However, extreme fines produce diminishing behavioral effects, highlighting the need for calibrated, region-specific penalty structures rather than uniformly increasing fines across the network.

6.4.4 Impact of Detection Technology and Patrol Efficiency

In this section, we examine how technological improvements in detection systems and patrol efficiency influence enforcement effectiveness, revenue generation, and driver behavior. The parameters governing the patrol process are the patrol speed v_{patrol} and the detection probability per pass p_{det}, which jointly determine the encounter intensity $\kappa = (v_{\text{patrol}} \cdot p_{\text{det}})/L_i$. Higher detection probabilities may represent the use of advanced sensors, license plate recognition, or smart-meter integration, while faster patrol speeds can result from motorized enforcement vehicles or route optimization software.

We explore nine technology scenarios formed by the combinations of patrol speed multipliers $v_{\text{mult}} \in \{0.75, 1.00, 1.25\}$ and detection probability multipliers $p_{\text{mult}} \in \{0.75, 1.00, 1.25\}$. Each combination alters the effective κ_i values across the five synthetic regions, and the dynamic programming model reallocates the available enforcement units under a fixed budget $B = 70$ to maximize total expected revenue.

Table 6.5 summarizes the outcomes across the technology levels. The results show that an increase in p_{det} (detection capability) has a stronger impact on total revenue than a proportional increase in v_{patrol}. For instance, improving p_{det} from 0.7 to 0.875 while keeping $v_{\text{patrol}} = 0.25$ km/h increases total expected revenue by approximately 1%, even though fewer officers are needed to reach similar detection coverage (41 units vs. 52 under baseline). Conversely, doubling patrol speed with no technological improvement yields minimal additional benefit, as officers already saturate the effective encounter rate.

As technology improves, the enforcement system becomes more efficient, enabling similar or higher revenue levels with fewer staff. This is primarily because higher κ_i values increase the expected detection rate, lowering the marginal value of adding additional officers in saturated regions such as Downtown and Business districts. Consequently, the model reallocates fewer

Table 6.5 Effect of patrol speed (v_{patrol}) and detection probability (p_{det}) on enforcement outcomes.

v_{mult}	p_{mult}	v_{patrol} [km/h]	p_{det}	Total Revenue [$]	Staff Used
0.75	0.75	0.188	0.525	113,273	70
0.75	1.00	0.188	0.700	122,775	68
0.75	1.25	0.188	0.875	122,765	55
1.00	0.75	0.250	0.525	122,775	68
1.00	1.00	0.250	0.700	122,838	52
1.00	1.25	0.250	0.875	122,945	41
1.25	0.75	0.312	0.525	122,765	55
1.25	1.00	0.312	0.700	122,945	41
1.25	1.25	0.312	0.875	122,720	34

officers to peripheral areas (e.g., Suburban) while maintaining or even improving deterrence effectiveness.

Furthermore, technological improvements shift driver behavior by reducing the switching time t_{switch} between illegal and legal (metered) parking. Under weak technology ($p_{\text{det}} = 0.525$), drivers tend to remain illegally parked for longer durations before switching, whereas in high-technology scenarios ($p_{\text{det}} = 0.875$), switching occurs earlier, leading to a higher share of legal-meter payments and fewer prolonged illegal stays. This highlights the complementarity between patrol automation and behavioral deterrence, showing that investments in detection technology can achieve comparable financial outcomes while reducing staffing needs and exposure costs.

6.4.5 Equity-Constrained Revenue and Staffing Sensitivity

We examine five fine levels ($F = \{200, 300, 400, 500, 600\}$) under a full staffing budget of $B = 100$, and then analyze an under-staffed case with $B = 30$. In each configuration, the dynamic programming (DP) allocation model is solved for increasing minimum average detection probabilities $\rho \in [0.01, 0.20]$, where higher ρ imposes stricter detection equity across regions.

Full staffing ($B = 100$). At the lowest fine level ($F = 200$), deterrence is weak and patrol intensity must remain high: *Downtown* and *Business* require 14 and 23 officers, respectively, with total revenue of $125,239. As fines increase to $F = 300$ and $F = 400$, optimal allocations contract to $(9, 15, 14, 5, 9)$ and $(7, 11, 10, 4, 7)$, while revenues stabilize near $121,000–$123,000. Introducing $F = 500$ maintains this trend: only $(5, 9, 8, 3, 5)$ officers are needed to achieve a comparable outcome of \$121,171. At the highest fine level ($F = 600$), deterrence effects dominate—allocations of $(5, 8, 7, 3, 5)$ suffice to generate nearly the same financial performance (\$120,164). This progression confirms that increasing citation fines effectively substitutes monetary penalties for enforcement labor, maintaining deterrence with fewer patrol resources.

Equity constraints ($P_i(\tau; n_i) \geq \rho$) impose minimum detection thresholds that limit concentration of staff in high-revenue regions. At small thresholds ($\rho \leq 0.05$), the constraints are nonbinding and allocations coincide with the baseline. Beyond $\rho = 0.07$, equity requirements trigger rebalancing toward *Downtown*, *Business*, and *MixedRR*, leading to diminishing returns. For example, at $F = 300$ and $\rho = 0.09$, the allocation $(19, 21, 17, 6, 9)$ lowers total revenue from \$122,838 to \$116,910. Similarly, for $F = 200$, enforcing $\rho = 0.09$ yields \$124,422 with $(19, 23, 20, 7, 15)$. At higher fine levels, analogous redistributions occur, for instance, $(19, 21, 17, 6, 9)$ under $F = 500$ or $F = 600$, with total revenue dropping below \$103,000 and \$99,000, respectively. When $\rho \geq 0.12$, the model becomes infeasible under $B = 100$, as the required number of officers exceeds the available budget, highlighting the tight trade-off between equity enforcement and operational feasibility.

Table 6.6 Total Revenue (\$) and Regional Allocations under Different Fine Levels and Equity Constraints ($B = 100$)

Equity Level ρ	$F = 200$	$F = 300$	$F = 400$	$F = 500$	$F = 600$
0.00 (no equity)	\$125,239 (14,23,20,7,15)	\$122,838 (9,15,14,5,9)	\$121,672 (7,11,10,4,7)	\$121,171 (5,9,8,3,5)	\$120,164 (5,8,7,3,5)
0.05	\$125,239 (14,23,20,7,15)	\$122,764 (10,15,14,5,9)	\$120,565 (10,11,10,4,7)	\$118,091 (10,11,9,3,5)	\$114,224 (10,11,9,3,5)
0.07	\$125,239 (14,23,20,7,15)	\$121,424 (14,16,14,5,9)	\$115,382 (14,16,13,5,7)	\$109,343 (14,16,13,5,7)	\$104,662 (14,16,13,5,7)
0.09	\$124,422 (19,23,20,7,15)	\$116,910 (19,21,17,6,9)	\$108,707 (19,21,17,6,9)	\$102,996 (19,21,17,6,9)	\$98,995 (19,21,17,6,9)
0.12	\$123,197 (23,25,21,7,15)	\$112,042 (23,25,21,7,11)	\$104,121 (23,25,21,7,11)	\$99,064 (23,25,21,7,11)	\$95,612 (23,25,21,7,11)
0.14+	Infeasible	Infeasible	Infeasible	Infeasible	Infeasible

Figure 6.4 displays the Revenue–Equity frontier across fine levels, showing that total revenue declines as ρ increases. The decline is steeper at higher fines, indicating that high-penalty regimes are more sensitive to equity constraints.

Under-staffed system ($B = 30$). A restricted budget amplifies the trade-off between equity and revenue. Table 6.7 summarizes results for $\rho \in \{0.01, 0.02, 0.03\}$; the problem

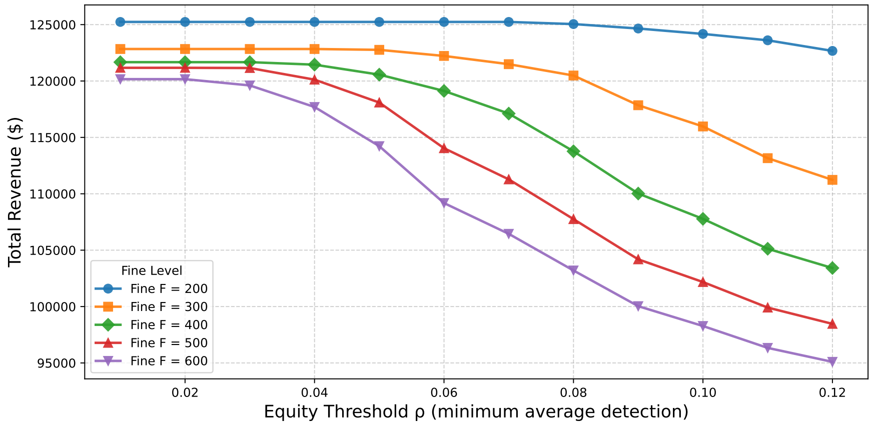

Figure 6.4 Revenue–Equity Frontier under different fine levels with staff budget $B = 100$.

becomes infeasible at $\rho \geq 0.04$. In the base case, all resources are concentrated in *Downtown*, *Business*, and *MixedRR*, while *Suburban* receives no coverage. Imposing equity progressively shifts patrols toward less profitable regions, modestly reducing revenue from $95,164 to $93,725 as ρ rises to 0.03. In the table, the notation "(min #)" indicates the minimum number of officers required in that region to satisfy the equity constraint at the stated ρ.

Table 6.7 Optimal allocations and revenues for the under-staffed budget ($B = 30$)

ρ	Downtown	Business	MixedRR	Entertainment	Suburban	Total revenue
Base	6	11	9	4	0	$95,164.30
0.01	6 (min 2)	11 (min 3)	8 (min 2)	4 (min 1)	1 (min 1)	$94,754.03
0.02	6 (min 4)	11 (min 5)	7 (min 4)	4 (min 2)	2 (min 2)	$94,274.08
0.03	6 (min 6)	11 (min 7)	6 (min 6)	4 (min 2)	3 (min 3)	$93,725.49

6.4.6 Case Study

This experiment revisits the six-region case study of Le Plateau–Mont-Royal (PMR) under an updated, empirically grounded economic setting. The PMR comprises six main arrondissements (subdivisions): Lorimier, Mile End, Milton–Parc, Parc–Lafontaine, Parc–Laurier, and Saint–Louis. The aggregate number of on-street paid parking events is approximately 80,500 per observation period, of which only 8.7% are associated with parking meters and other parking categories such as private or interior lots are excluded from the calibration [161]. According to [162], the distribution of paid parking spaces across these subdivisions are: Lorimier (366), Mile End (1757), Milton–Parc (300), Parc–Lafontaine (703), Parc–Laurier (302), and Saint–Louis (942). Accordingly, the total effective demand of $80{,}522 \times 0.087$ parking events is proportionally allocated across the six regions based on their number of paid spaces. The total street network length in PMR is about 136 km [163]; the relative length per subdivision is estimated according to their proportional area shares.

The monetary parameters reflect current enforcement conditions in Montréal: the citation fine is set to $F = \$91$, and the hourly meter cost is $r = 4.25/60$ \$/min, following the reported parking tariff structure [162]. The baseline daily pass cost is $D_p = \$25$, and the convenience offset for legal parking remains $K = \$5$. A nominal detection probability $p_{\text{det}} = 0.5$ is assumed to capture typical officer attentiveness and technology use in this borough.

The region-specific patrol encounter rate is computed as

$$\kappa_i = \frac{v_{\text{patrol}}\, p_{\text{det}}}{L_i}, \tag{6.2}$$

where $v_{\text{patrol}} = 2$ km/h represents the average patrol speed and L_i is the estimated effective street length of Arrondissement i. The driver decision structure thus uses the updated parameters:

$$F = \$91, \qquad r = \tfrac{4.25}{60}\ \$/\text{min}, \qquad K = \$5, \qquad D_p = \$25, \qquad a = 60\ \text{min}, \qquad p_{\text{det}} = 0.5.$$

Table 6.8 presents the equilibrium outcomes and optimal staffing allocation for the six-region instance of the PMR borough. Each region differs in parking demand, critical staffing threshold (n_i^*), and feasible daily-pass interval $[C_1, C_2]$. The dynamic programming algorithm was solved under a total staffing budget of $B = 20$, with the system optimally allocating 12 officers, yielding a total expected revenue of $40,414.

Table 6.8 Optimal staffing and regional outcomes under budget $B = 20$.

Region	n_i^*	$[C_1, C_2]$ [$]	Alloc. n_i	Demand	Prob. legal	Revenue [$]
Milton–Parc	1	(16.89, 46.74)	1	480.8	0.061	2,824.07
Mile End	3	(10.41, 69.71)	3	2,815.8	0.280	16,217.01
Saint–Louis	1	(15.11, 51.57)	1	1,509.7	0.093	8,895.02
Parc–Laurier	2	(8.12, 81.89)	2	484.0	0.479	2,629.49
Lorimier	4	(15.11, 51.57)	4	586.6	0.093	3,456.03
Parc–Lafontaine	1	(9.59, 73.91)	1	1,126.6	0.340	6,392.27
Total			**12**			**40,413.89**

The results reveal clear spatial heterogeneity in enforcement efficiency. Mile End, with the highest demand, receives the largest allocation ($n_i = 3$) and contributes over 40% of total revenue. Lorimier, characterized by long street coverage and low detection efficiency, requires four officers to sustain comparable outcomes. In contrast, compact areas such as Milton–Parc and Parc–Lafontaine operate efficiently with minimal staffing while maintaining moderate compliance. Legal parking probabilities remain below 0.5 in all regions, confirming that moderate detection rates preserve a profitable level of noncompliance. Overall, the model allocates staff only where marginal gains remain positive, ensuring near-optimal utilization of the enforcement budget.

Analysis of Fixed vs. Dynamic Staffing Scenarios

The analysis of fixed staffing scenarios holds significant operational value, not as a confirmation of the optimal solution (which is self-evident), but to assess the cost-effectiveness and robustness of near-optimal fixed policies across the varying fine landscape.

Figure 6.5 compares total revenues under dynamic allocation and three fixed-staff configurations (10, 12, and 15 officers) as fine levels increase. The dynamic model endogenously adjusts staffing to maximize total revenue, showing a clear substitution between fine severity and patrol effort. Officer deployment decreases monotonically from 20 to 9 as fines rise, while revenues peak near $F = 55$ ($41,942) and then gradually stabilize around $40,414 at the current fine level of $F = 91$. This indicates that the marginal enforcement value of additional officers diminishes as fines strengthen deterrence. The system thus naturally balances deterrence and coverage: higher fines reduce required staffing without compromising financial outcomes.

When the staffing level is fixed at $n = 12$ officers, which matches the dynamically optimal level ($n_{\text{opt}} = 12$) under the current fine $F = 91$, the resulting total revenue ($40,414) is identical to the dynamic optimum. We examine this case to understand how the existing staffing configuration performs as the fine level F varies. The fixed–15 allocation outperforms under low fines ($F \leq 70$) due to increased patrol presence, but exhibits diminishing returns at higher fines, where additional enforcement yields minimal benefit once deterrence becomes sufficiently strong.

The introduction of the $n = 10$ scenario reveals the potential for further cost optimization. For very high fines ($F \geq 100$), the fixed $n = 10$ revenue curve matches the dynamic optimum, confirming that $n = 10$ is the efficient staffing level in this higher fine range.

This stability suggests that for management, the current $F = 91, n = 12$ policy is highly cost-effective. The slight revenue deficit compared to the overall peak at $F = 55$ must be weighed against the significant reduction in personnel costs gained by dropping 8 officers (from $n = 20$

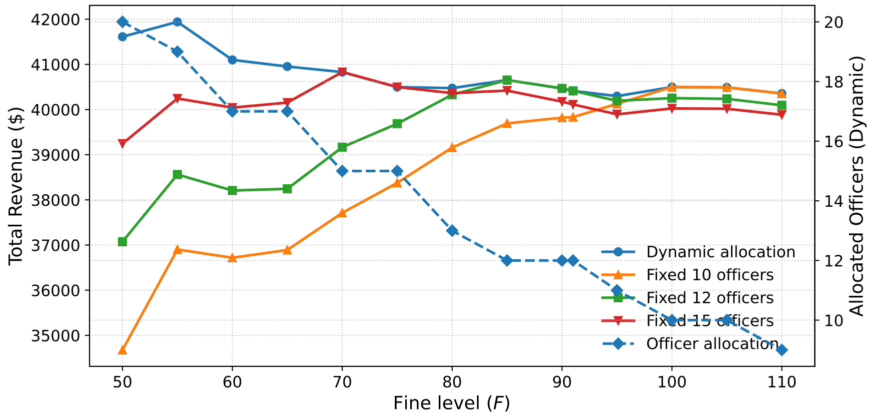

Figure 6.5 Revenue comparison across fine levels for dynamic, fixed-12, and fixed-15 officer allocations. The current operational configuration ($F = 91$, 12 officers) aligns closely with the dynamic optimum.

down to $n = 12$). This fine-induced staff reduction represents the most critical operational insight, allowing the city to sustain high financial outcomes with fewer resources.

Across all regimes, the revenue difference between $n = 12$ and $n = 10$ at $F = 91$ is minimal (\$40,414 vs. \$39,832, a gap of $\sim 1.4\%$). This indicates that while $n = 12$ is optimal, a small reduction to $n = 10$ would maintain $near\text{-}optimal\ performance$ at a lower personnel cost. However, the dynamic model would become valuable under weaker fines or degraded detection conditions, where adaptive resource allocation could recover up to 10–12% additional revenue (e.g., at $F = 50$, where $n_{\text{opt}} = 20$).

6.5 Conclusion

This paper integrates a closed-form behavioral deterrence mechanism with a tractable resource-allocation model to support tactical parking enforcement planning. The analytical threshold derived in Section 6.2 characterizes the transition between weak and effective deterrence, and the dynamic program allocates patrol effort across regions to maximize long-run revenue subject to an equity requirement.

The results yield several actionable insights for enforcement planning. First, *staffing focus*: optimal deployments operate near the effectiveness threshold rather than far above it; once deterrence is achieved, additional staffing yields sharply diminishing returns. Second, *fine calibration*: intermediate fines tend to maximize revenue and shift optimal staffing downward, indicating that fines should be calibrated jointly with deployment rather than adjusted in isolation. Third, *technology versus labor*: increases in detection efficiency κ (e.g., via ALPR or improved routing) can substitute for additional officers while maintaining deterrence. Finally, *equity considerations*: enforcing a minimum average detection probability promotes more balanced enforcement across regions with only modest revenue impact, revealing a transparent revenue–equity policy frontier.

The framework yields a practical workflow: (i) estimate region-level parameters (κ_i, a_i, r_i, F_i, $D_{p,i}, K_i, D_i$); (ii) compute behavioral thresholds and revenue curves; (iii) solve the allocation problem via dynamic programming, with or without equity constraints; and (iv) evaluate fine and technology sensitivity to guide policy calibration of (n_i, ρ, F) under agency objectives.

The analysis assumes that short-run parking demand is inelastic to enforcement, that drivers receive at most one citation per stay, that dwell times follow an exponential distribution, and that drivers minimize expected monetary costs. These assumptions support analytical tractability. Extensions could incorporate empirical dwell time distributions (e.g., lognormal), perceived rather than actual detection rates, or behavioral probability weighting, which would smooth but not eliminate threshold-based switching. The model fixes total staffing B; introducing explicit labor or technology cost functions would endogenize the optimal B itself, shifting the problem from tactical allocation to strategic resource planning. Additional extensions could incorporate curb-choice displacement, time-of-day scheduling, or region-specific policy priorities.

Overall, the results provide a decision-support structure linking behavioral deterrence, financial outcomes, and distributional fairness, allowing agencies to coordinate staffing, fine levels, and technology investments in a unified manner.

CHAPTER 7 SUMMARY OF WORK

This chapter synthesizes the main findings from the three research studies that form this dissertation and highlights their managerial and policy implications. While the studies differ in scope and decision horizon, they share a common premise: parking enforcement is not only a logistical task of assigning officers to locations, but also a behavioral and policy-driven activity that shapes how drivers choose between legal and illegal parking. The studies collectively show that enforcement strategies are more effective when they explicitly account for how drivers respond to enforcement visibility, when they respect the time dynamics of violation recurrence, and when they are designed with spatial fairness in mind.

7.1 Integrated View of the Research Contributions

Municipal parking enforcement agencies must manage limited officer capacity, heterogeneous parking demand across neighborhoods, and evolving public expectations regarding fairness and mobility. Traditionally, patrol decisions have been driven either by static coverage rules or by complaint-based deployment. Both approaches ignore the strategic behavior of drivers and the long-term consequences of enforcement patterns.

This research program addresses this gap by developing planning methods at two complementary time scales:

- **Operational horizon (daily):** how to route officers during a shift to maintain compliance.
- **Tactical horizon (seasonal or yearly):** how to allocate enforcement resources across neighborhoods.

The three studies move progressively along this spectrum, building a comprehensive decision-support framework.

7.2 Key Contributions and Practical Insights

7.2.1 Operational Enforcement Planning with Behavioral Response

The first study demonstrates that the perceived likelihood of receiving a citation strongly shapes illegal parking behavior. Areas with frequent officer visits tend to exhibit higher compliance because drivers adapt and avoid risking penalties. This means that patrol scheduling directly influences compliance levels, not simply through citations issued, but through the deterrence signal it creates.

The study proposes an operational planning framework that determines officer patrol routes while accounting for this deterrence effect. A matheuristic solution method enables the approach to scale to realistic networks. Computational experiments show that schedules which incorporate behavioral response lead to significantly improved compliance and fewer violations compared to schedules optimized for coverage alone.

Managerial message: Patrol routes should be actively optimized to leverage the deterrence effect by maximizing perceived enforcement presence, as this leads to significantly improved compliance and fewer violations compared to simple coverage-based schedules.

7.2.2 Multi-Visit Routing with Temporal Recovery Considerations

The second study extends the operational perspective by recognizing that illegal parking rates increase again after an area is inspected. However, revisiting too soon yields little enforcement value because the opportunity for new violations to occur has not yet materialized. The key is therefore not only *how many times* an area is visited, but *when*.

To address this, the study introduces a recovery-time constraint that prevents back-to-back visits to the same location. A Variable Neighborhood Descent approach is used to generate patrol plans that balance the number of visits and the timing of return. The results show that the most effective schedules include multiple well-timed inspections of high-turnover areas.

Managerial message: Enforcement resources should be deployed efficiently by focusing on strategically timing revisits, particularly in high-turnover areas, using a recovery-time constraint to maximize the window for new violations to occur, thereby maximizing compliance and revenue.

7.2.3 Tactical Resource Allocation with Equity Requirements

The third study moves to the tactical scale, where the problem is no longer how to route officers, but how many officers to allocate to different regions. At this horizon, the goals include long-run compliance stability, spatial fairness, and fair access to public space.

The study identifies a threshold level of enforcement presence required to shift behavior from persistent violation to sustained compliance. This threshold is embedded in a dynamic programming allocation framework that ensures each neighborhood receives at least a minimum level of enforcement visibility.

Managerial message: Tactical resource allocation must ensure every region meets the critical enforcement threshold required for sustained compliance. Resource allocation should be equitable, guaranteeing a uniform minimum level of enforcement visibility across all neighborhoods, even if it involves a modest trade-off with maximum revenue.

7.3 Overall Managerial Insights

Across the three studies, several high-level lessons emerge for enforcement agencies and urban mobility planners:

- **Compliance is shaped by expectations, not punishment.** Increasing perceived enforcement presence reduces violations more effectively than increasing fines.
- **Time matters.** Inspection frequency and spacing determine whether enforcement resources are used efficiently.
- **Fairness is operationally relevant.** Uneven deployment leads to predictable shifts in illegal parking patterns.

7.4 Closing Remarks

Taken together, the three studies provide a foundation for more adaptive, equitable, and behaviorally informed parking enforcement strategies. They demonstrate how rigorous operational research methods can support day-to-day scheduling decisions and long-term resource planning, while aligning enforcement practice with wider urban mobility goals. The work suggests that future improvements should focus on real-time adaptation and integration with digital curb monitoring systems, enabling even more responsive and targeted enforcement strategies.

CHAPTER 8 CONCLUSION

This dissertation addressed the planning of parking enforcement in urban environments where driver behavior responds to perceived enforcement intensity. Chapter 1 introduced the problem context and articulated the overarching need for planning tools that integrate behavioral mechanisms, operational feasibility, and policy objectives. Chapter 2 surveyed the relevant literature on parking policy and governance, behavioral compliance modeling, and optimization approaches for enforcement planning. Chapter 3 situated the three methodological studies within a coherent research program spanning multiple decision horizons. Chapters 4 through 6 then presented the core contributions: (i) an operational planning framework that routes officers under behavior-responsive deterrence; (ii) a single-period, multi-visit operational model that imposes recovery-time separation to capture temporal re-emergence of violations; and (iii) a tactical allocation framework that distributes enforcement resources across heterogeneous regions while enforcing a minimum detection-probability requirement to support equity. Chapter 7 synthesized these results and distilled managerial insights, emphasizing the central roles of credible deterrence, time-aware scheduling, and fair spatial deployment.

Collectively, these developments establish a multi-scale decision-support framework that advances both theory and practice. The results demonstrate that enforcement strategies designed without behavioral feedback or temporal separation can misallocate scarce patrol capacity, undermine deterrence, and yield inequitable outcomes. In contrast, the integrated approach developed here enables agencies to: (a) schedule daily patrols that sustain compliance rather than merely chase citations; (b) time repeat inspections to coincide with violation recurrence; and (c) allocate enforcement presence across neighborhoods in a way that is both effective and fair. Across operational and tactical horizons, the framework reframes parking enforcement as a public service that supports access, safety, and mobility, with success measured by reduced violations and credible compliance rather than by citation volume alone.

8.1 Research Limitations

Several limitations qualify the findings and suggest caution in direct operational deployment. First, the behavioral components rely on equilibrium or steady-state assumptions that may not fully capture short-term shocks, special events, weather anomalies, or abrupt policy changes; their validity depends on context-specific calibration and monitoring. Second, the operational routing models abstract from day-of operations frictions such as unexpected congestion, officer discretion, service interruptions, and heterogeneous inspection times, which can introduce variance between planned and realized patrols. Third, the tactical allocation framework implements equity as a minimum detection-probability requirement; while defensible and transparent, it does not exhaust the broader equity landscape (e.g., distributional impacts across socio-demographic groups, political priorities, or historical disparities). Finally, although the proposed heuristics and dynamic programs are designed for realistic scales, their performance and robustness still depend on data quality (arrival rates, dwell patterns, travel times) and on institutional capacity to maintain parameter calibration over time.

8.2 Future Research Directions

A first avenue is to couple the models with *real-time data and adaptive control*. Integrating feeds from curbside sensors, ALPR video, or mobile reports can enable rolling re-optimization of patrols, learning-based prediction of violation hotspots, and closed-loop adjustment of recovery times. This would shift the framework from plan-and-execute to sense–predict–act, improving resilience to shocks and special events.

A second direction is to *model heterogeneity and multi-actor curb demand*. Commercial delivery fleets, ride-hailing vehicles, service technicians, and residents face different costs and constraints. Extending the behavioral layer to segment responses, incorporate time windows, loading needs, or permit rules, and represent repeated interactions would yield richer, more targeted enforcement and complementary policy levers (e.g., micro-hubs, time-limited loading, or graduated fines).

A third direction is to *expand policy objectives and constraints*. Embedding environmental externalities (e.g., cruising emissions), safety externalities (e.g., blocked crosswalks or bus lanes), and accessibility goals (e.g., around schools and hospitals) would allow multi-criteria planning and explicit trade-off analysis. On the equity front, exploring alternative fairness notions (max–min, Gini bounds, weighted priorities) and auditing distributional impacts would strengthen policy legitimacy.

A fourth avenue is *field validation and organizational integration*. Partnering with municipalities to conduct pilots would test implementation details, data pipelines, calibration workflows, officer interfaces, and exception handling, and quantify benefits in terms of reduced violations, improved turnover, and public satisfaction. Insights from deployment can be fed back into the models to refine parameters and stress-test assumptions.

A fifth direction is *methodological scaling and robustness*. For large networks and frequent re-optimizations, hybrid approaches that combine fast metaheuristics with learned surrogates or decomposition algorithms could provide reliable performance under tight decision cycles. Robust and risk-aware variants (e.g., distributionally robust routing or allocation under demand uncertainty) can protect against parameter drift and rare but impactful conditions.

REFERENCES

[1] G. Marsden, "The evidence base for parking policies—a review," *Transport policy*, vol. 13, no. 6, pp. 447–457, 2006.

[2] F. Kirschner and M. Lanzendorf, "Parking management for promoting sustainable transport in urban neighbourhoods. a review of existing policies and challenges from a german perspective," *Transport Reviews*, vol. 40, no. 1, pp. 54–75, 2020.

[3] E. Inci, "A review of the economics of parking," *Economics of Transportation*, vol. 4, no. 1-2, pp. 50–63, 2015.

[4] C. De Gruyter, L. Davies, and L. T. Truong, "Examining spatial variations in minimum residential parking requirements in melbourne," *Journal of Transport Geography*, vol. 94, p. 103096, 2021.

[5] D. C. Shoup, "The trouble with minimum parking requirements," *Transportation research part A: policy and practice*, vol. 33, no. 7-8, pp. 549–574, 1999.

[6] E. Ferguson, "Zoning for parking as policy process: A historical review," *Transport Reviews*, vol. 24, no. 2, pp. 177–194, 2004.

[7] M. McShane and M. D. Meyer, "Parking policy and urban goals: Linking strategy to needs," *Transportation*, vol. 11, no. 2, pp. 131–152, 1982.

[8] S. Biswas, S. Chandra, and I. Ghosh, "Effects of on-street parking in urban context: A critical review," *Transportation in developing economies*, vol. 3, no. 1, p. 10, 2017.

[9] S. Brooke, S. Ison, and M. Quddus, "On-street parking search: Review and future research direction," *Transportation Research Record*, vol. 2469, no. 1, pp. 65–75, 2014.

[10] K. Cullinane and J. Polak, "Illegal parking and the enforcement of parking regulations: causes, effects and interactions," *Transport Reviews*, vol. 12, no. 1, pp. 49–75, 1992.

[11] M. Morfoulaki, K. M. Kotoula, G. Myrovali, and G. Aifadopoulou, "Calculating the impacts of alternative parking pricing and enforcement policies in urban areas with traffic problems," *Periodica Polytechnica Transportation Engineering*, vol. 45, no. 1, pp. 35–41, 2017.

[12] X. Hu, F. Chen, and H. Sun, "Optimization of urban parking cooperation and pricing: a case of harbin, china," *Transportation Research Interdisciplinary Perspectives*, vol. 33, p. 101619, 2025.

[13] E. Erfina, L. Lukman, H. Hardianti, and B. B. Budai, "Illegal parking control strategy to improve sustainable management of urban areas," *Journal of Governance and Public Policy*, vol. 12, no. 1, pp. 112–124, 2025.

[14] D. Shoup, *High cost of free parking.* Routledge, 2021.

[15] J. Zhou, "An integrated model of parking pricing and cruising," in *CICTP 2014: Safe, Smart, and Sustainable Multimodal Transportation Systems*, 2014, pp. 3441–3449.

[16] C. Lei and Y. Ouyang, "Dynamic pricing and reservation for intelligent urban parking management," *Transportation Research Part C: Emerging Technologies*, vol. 77, pp. 226–244, 2017.

[17] S. H. Bayih and S. L. Tilahun, "Dynamic vehicle parking pricing. a review," *Operations Research and Decisions*, vol. 34, no. 1, pp. 35–59, 2024.

[18] R. I. Xiao and M. Jaller, "Spatial analysis and predictive modeling framework of truck parking and idling impacts on environmental justice communities," *Journal of Transport Geography*, vol. 127, p. 104263, 2025.

[19] M. Shetty, S. Sinha, and J. Chakraborty, "Impact of parking pricing and regulations on user behavior," in *Urban Mobility Research in India (UMI 2022).* Springer, 2022, pp. 111–135.

[20] H. Wang, R. Li, X. C. Wang, and P. Shang, "Effect of on-street parking pricing policies on parking characteristics: A case study of nanning," *Transportation Research Part A: Policy and Practice*, vol. 137, pp. 65–78, 2020.

[21] M. Nourinejad and M. J. Roorda, "Impact of hourly parking pricing on travel demand," *Transportation Research Part A: Policy and Practice*, vol. 98, pp. 28–45, 2017.

[22] F. Zong, M. Zeng, and P. Yu, "A parking pricing scheme considering parking dynamics," *Transportation*, vol. 51, no. 4, pp. 1349–1371, 2024.

[23] N. Khordagui, "Parking prices and the decision to drive to work: Evidence from california," *Transportation Research Part A: Policy and Practice*, vol. 130, pp. 479–495, 2019.

[24] F. Lu, J. Zhang, and J. M. Perloff, “General and specific information in deterring traffic violations: evidence from a randomized experiment,” *Journal of Economic Behavior & Organization*, vol. 123, pp. 97–107, 2016.

[25] D. M. Zaidel, “A modeling perspective on the culture of driving,” *Accident Analysis & Prevention*, vol. 24, no. 6, pp. 585–597, 1992.

[26] M. Karaliopoulos, K. Katsikopoulos, and L. Lambrinos, “Bounded rationality can make parking search more efficient: The power of lexicographic heuristics,” *Transportation Research Part B: Methodological*, vol. 101, pp. 28–50, 2017.

[27] R. G. Thompson and A. J. Richardson, “A parking search model,” *Transportation Research Part A: Policy and Practice*, vol. 32, no. 3, pp. 159–170, 1998.

[28] F. Ghizzawi, A. Galal, and M. J. Roorda, “Modelling parking behaviour of commercial vehicles: a scoping review,” *Transport Reviews*, vol. 44, no. 4, pp. 743–765, 2024.

[29] A. Amer and J. Y. Chow, “A downtown on-street parking model with urban truck delivery behavior,” *Transportation Research Part A: Policy and Practice*, vol. 102, pp. 51–67, 2017.

[30] J. Yan, X. Zhou, and Y. Ji, “What causes curbside illegal parking behavior: A method based on structural equation model,” *KSCE Journal of Civil Engineering*, vol. 27, no. 8, pp. 3581–3590, 2023.

[31] A. Khaliq, P. van der Waerden, and D. Janssens, “Modeling car drivers’ on-street parking decisions using the integrated hierarchical information integration approach,” *Transportation Research Record*, vol. 2672, no. 49, pp. 23–33, 2018.

[32] S. Tartarone and A. H. Levis, “Enforcement policies for parking regulations,” *IEEE Transactions on Systems, Man, and Cybernetics*, no. 6, pp. 604–610, 2007.

[33] A. Brudner, “On the management of residential on-street parking: Policies and repercussions,” *Transport Policy*, vol. 138, pp. 94–107, 2023.

[34] T. Rye and R. Llewellyn, “Parking enforcement activity and public attitudes to enforcement: a case study of edinburgh, scotland (uk) in the context of the eu push and pull project,” in *International Workshop on the Economics of Parking, Barcelona, 28th Nov*, 2016.

[35] E. Fokker, E. Dugundji, and T. Koch, “Analysis of the impact of policy measures on parking behavior using interpretable time series models,” *Journal of Transport and Land Use*, vol. 17, no. 1, pp. 751–780, 2024.

[36] X. Zhou, X. Ding, J. Yan, and Y. Ji, “Spatial heterogeneity of urban illegal parking behavior: A geographically weighted poisson regression approach,” *Journal of Transport Geography*, vol. 110, p. 103636, 2023.

[37] M. Sinning and Y. Zhang, “Social norms or enforcement? a natural field experiment to improve traffic and parking fine compliance,” *Journal of Economic Behavior & Organization*, vol. 210, pp. 43–60, 2023.

[38] X. Deng, T. Wang, L. Xue, and S. Yang, “Norm enforcement on minorities: Evidence from traffic violations,” *Journal of Economic Behavior & Organization*, vol. 231, p. 106930, 2025.

[39] C. Biyik, Z. Allam, G. Pieri, D. Moroni, M. O’Fraifer, E. O’Connell, S. Olariu, and M. Khalid, “Smart parking systems: Reviewing the literature, architecture and ways forward,” *Smart Cities*, vol. 4, no. 2, pp. 623–642, 2021.

[40] C. Georgopoulou and P. Papantoniou, “A literature review of automated roadside parking monitoring using artificial intelligence algorithms,” *Electronics*, vol. 14, no. 20, p. 4119, 2025.

[41] R. Ke, Y. Zhuang, Z. Pu, and Y. Wang, “A smart, efficient, and reliable parking surveillance system with edge artificial intelligence on IoT devices,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 8, pp. 4962–4974, 2020.

[42] D. Chimba and M. Onyango, “Optimization of short-term on-street park-pay license plate surveying,” *Journal of infrastructure systems*, vol. 18, no. 3, pp. 194–201, 2012.

[43] M. A. Ala’anzy, A. Abilakim, R. Zhanuzak, and L. Li, “Real time smart parking system based on IoT and fog computing evaluated through a practical case study,” *Scientific Reports*, vol. 15, no. 1, p. 33483, 2025.

[44] W. Shao, S. Tan, S. Zhao, K. K. Qin, X. Hei, J. Chan, and F. D. Salim, “Incorporating lstm auto-encoders in optimizations to solve parking officer patrolling problem,” *ACM Transactions on Spatial Algorithms and Systems (TSAS)*, vol. 6, no. 3, pp. 1–21, 2020.

[45] Y. Xie, S. Shekhar, and Y. Li, “Statistically-robust clustering techniques for mapping spatial hotspots: A survey,” *ACM Computing Surveys (CSUR)*, vol. 55, no. 2, pp. 1–38, 2022.

[46] M. Adepeju, G. Rosser, and T. Cheng, “Novel evaluation metrics for sparse spatio-temporal point process hotspot predictions–a crime case study,” *International Journal of Geographical Information Science*, vol. 30, no. 11, pp. 2133–2154, 2016.

[47] N. Brazil, “The unequal spatial distribution of city government fines: The case of parking tickets in los angeles,” *Urban Affairs Review*, vol. 56, no. 3, pp. 823–856, 2020.

[48] M. Kladeftiras and C. Antoniou, “Simulation-based assessment of double-parking impacts on traffic and environmental conditions,” *Transportation research record*, vol. 2390, no. 1, pp. 121–130, 2013.

[49] A. M. Said, A. E. Kamal, and H. Afifi, “An intelligent parking sharing system for green and smart cities based IoT,” *Computer Communications*, vol. 172, pp. 10–18, 2021.

[50] T. Adnan, A. Erfani, and Q. Cui, “Paving equity: Unveiling socioeconomic patterns in pavement conditions using data mining,” *Journal of Management in Engineering*, vol. 41, no. 5, p. 04025041, 2025.

[51] H. Ong and T. Goh, “Routing and scheduling of wardens for public car park inspection,” *Computers & industrial engineering*, vol. 12, no. 4, pp. 239–248, 1987.

[52] J.-Y. Wang, “Computer-aided system for planning efficient routes,” Ph.D. dissertation, Purdue University, 1992.

[53] H. Yan and G. L. Thompson, “Finding postal carrier walk paths in mixed graphs,” *Computational Optimization and Applications*, vol. 9, no. 3, pp. 229–247, 1998.

[54] N. S. Summerfield, M. Dror, and M. A. Cohen, “City streets parking enforcement inspection decisions: The chinese postman’s perspective,” *European Journal of Operational Research*, vol. 242, no. 1, pp. 149–160, 2015.

[55] M. Bruglieri, “The parking warden tour problem,” *Networks*, vol. 76, no. 4, pp. 539–554, 2020.

[56] I. Black, K. Cullinane, and C. Wright, “Parking enforcement policy assessment using an economic approach: Part 1: Theoretical background and the development of an economic model,” *Transportation Planning and Technology*, vol. 17, no. 3, pp. 249–257, 1993.

[57] W. Shao, F. D. Salim, T. Gu, N.-T. Dinh, and J. Chan, “Traveling officer problem: Managing car parking violations efficiently using sensor data,” *IEEE Internet of Things Journal*, vol. 5, no. 2, pp. 802–810, 2017.

[58] R. Ferreira, V. Móra, M. C. Mourão, M. Moz, L. S. Pinto, and J. Ribeiro, "Arc routing for parking enforcement officers: Exact and heuristic solutions," *European Journal of Operational Research*, vol. 299, no. 1, pp. 283–301, 2022.

[59] S. Samanta, G. Sen, and S. K. Ghosh, "A literature review on police patrolling problems," *Annals of Operations Research*, vol. 316, no. 2, pp. 1063–1106, 2022.

[60] M. Thabet, M. Messaadia, and A. K. Al-Khalifa, "Police patrol routes optimization: A literature review," in *2023 International Conference On Cyber Management And Engineering (CyMaEn)*. IEEE, 2023, pp. 66–70.

[61] N. Kwak and M. B. Leavitt, "Police patrol beat design: Allocation of effort and evaluation of expected performance," *Decision Sciences*, vol. 15, no. 3, pp. 421–433, 1984.

[62] T. Vlćek, K. Haase, M. Fliedner, and T. Cors, "Police service district planning," *OR Spectrum*, vol. 46, no. 4, pp. 1029–1061, 2024.

[63] H. Chen, T. Cheng, and X. Ye, "Designing efficient and balanced police patrol districts on an urban street network," *International Journal of Geographical Information Science*, vol. 33, no. 2, pp. 269–290, 2019.

[64] K. Chelst, "An algorithm for deploying a crime directed (tactical) patrol force," *Management Science*, vol. 24, no. 12, pp. 1314–1327, 1978.

[65] D. G. Olson and G. P. Wright, "Models for allocating police preventive patrol effort," *Journal of the Operational Research Society*, vol. 26, no. 4, pp. 703–715, 1975.

[66] A. J. de Lucena, S. Bocanegra, and F. L. Gorgônio, "Location-allocation of police vehicles under a multi-period perspective," in *Proceedings of the XV Brazilian Symposium on Information Systems*, 2019, pp. 1–8.

[67] N. Adler, A. S. Hakkert, J. Kornbluth, T. Raviv, and M. Sher, "Location-allocation models for traffic police patrol vehicles on an interurban network," *Annals of operations research*, vol. 221, no. 1, pp. 9–31, 2014.

[68] C. Porras, J. Fajardo, and A. Rosete, "Multi-coverage dynamic maximal covering location problem," *Investigación Operacional*, vol. 40, no. 1, pp. 140–149, 2019.

[69] B. B. Keskin, S. R. Li, D. Steil, and S. Spiller, "Analysis of an integrated maximum covering and patrol routing problem," *Transportation Research Part E: Logistics and Transportation Review*, vol. 48, no. 1, pp. 215–232, 2012.

[70] X. Chen, “Police patrol optimization with security level functions,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 43, no. 5, pp. 1042–1051, 2013.

[71] K. Cullinane, “An aggregate dynamic model of the parking compliance decision,” *International Journal of Transport Economics/Rivista internazionale di economia dei trasporti*, pp. 27–50, 1993.

[72] M. Nourinejad and M. J. Roorda, “Parking enforcement policies for commercial vehicles,” *Transportation Research Part A: Policy and Practice*, vol. 102, pp. 33–50, 2017.

[73] M. Nourinejad, A. Gandomi, and M. J. Roorda, “Illegal parking and optimal enforcement policies with search friction,” *Transportation Research Part E: Logistics and Transportation Review*, vol. 141, p. 102026, 2020.

[74] L. Guo, S. Huang, J. Zhuang, and A. W. Sadek, “Modeling parking behavior under uncertainty: a static game theoretic versus a sequential neo-additive capacity modeling approach,” *Networks and Spatial Economics*, vol. 13, no. 3, pp. 327–350, 2013.

[75] R. Petiot, “Parking enforcement and travel demand management,” *Transport Policy*, vol. 11, no. 4, pp. 399–411, 2004.

[76] Y. J. Gur and E. A. Beimborn, “Analysis of parking in urban centers: equilibrium assignment approach,” *Transportation Research Record*, no. 957, 1984.

[77] M. Duan, D. Wu, and H. Liu, “Bi-level programming model for resource-shared parking lots allocation,” *Transportation Letters*, vol. 12, no. 7, pp. 501–511, 2020.

[78] L. Brotcorne, P. Escalona, B. Fortz, and M. Labbé, “Fare inspection patrols scheduling in transit systems using a stackelberg game approach,” *Transportation Research Part B: Methodological*, vol. 154, pp. 1–20, 2021.

[79] K. K. Trejo, J. B. Clempner, and A. S. Poznyak, “A stackelberg security game with random strategies based on the extraproximal theoretic approach,” *Engineering Applications of Artificial Intelligence*, vol. 37, pp. 145–153, 2015.

[80] T. Garrec, “Continuous patrolling and hiding games,” *European Journal of Operational Research*, vol. 277, no. 1, pp. 42–51, 2019.

[81] R. Naja, N. Mouawad, A. J. Ghandour, and K. Fawaz, “Speed trap optimal patrolling: Stop playing stackelberg security games,” *Wireless Personal Communications*, vol. 98, no. 4, pp. 3563–3582, 2018.

[82] A. Rosenfeld and O. Maksimov, "Optimal cruiser-drone traffic enforcement under energy limitation," *Artificial Intelligence*, vol. 277, p. 103166, 2019.

[83] J. Wang and L. Cui, "Patrolling games with coordination between monitoring devices and patrols," *Reliability Engineering & System Safety*, vol. 233, p. 109109, 2023.

[84] A. Rezazadeh, L. Zhang, G. Reniers, N. Khakzad, and V. Cozzani, "Optimal patrol scheduling of hazardous pipelines using game theory," *Process Safety and Environmental Protection*, vol. 109, pp. 242–256, 2017.

[85] T. Han and A. Levis, "Optimal stochastic enforcement of parking regulations," *IEEE Transactions on Automatic Control*, vol. 21, no. 3, pp. 354–358, 1976.

[86] C. Lei, Q. Zhang, and Y. Ouyang, "Planning of parking enforcement patrol considering drivers' parking payment behavior," *Transportation Research Part B: Methodological*, vol. 106, pp. 375–392, 2017.

[87] V. H. Vidigal Corrêa, H. Dong, M. Iori, A. G. dos Santos, M. Yagiura, and G. Zucchi, "An iterated local search for a multi-period orienteering problem arising in a car patrolling application," *Networks*, vol. 83, no. 1, pp. 153–168, 2024.

[88] (2024) Revenues soar as 2024 sees increase in paid parking in montreal. Accessed: 2025-01-06. [Online]. Available: https://montreal.citynews.ca/2024/11/27/revenues-soar-2024-paid-parking-montreal/

[89] Agence de Mobilité Durable de Montréal, "Parking enforcement: far beyond tickets," 2020, accessed: 2024-12-02.

[90] SARECO. (2020) Evolution of parking compliance through local enforcement. SARECO Panel Discussion. Discussed in panel hosted by AMDM. [Online]. Available: https://www.sareco.com

[91] Barnacle Parking, "What happens when you don't enforce parking," 2020, accessed: 2024-12-02.

[92] K.-W. Song, S.-M. Jeon, K.-J. Park, and Y.-E. Shin, "Study on the psychology of drivers dealing with illegal parking behavior," *KSCE Journal of Civil Engineering*, vol. 26, no. 6, pp. 2919–2925, 2022.

[93] J. Gao, K. Xie, and K. Ozbay, "Exploring the spatial dependence and selection bias of double parking citations data," *Transportation Research Record*, vol. 2672, no. 42, pp. 159–169, 2018.

[94] K. K. Qin, W. Shao, Y. Ren, J. Chan, and F. D. Salim, “Solving multiple travelling officers problem with population-based optimization algorithms,” *Neural Computing and Applications*, vol. 32, pp. 12 033–12 059, 2020.

[95] W. Kim and X. Wang, “Double parking in new york city: a comparison between commercial vehicles and passenger vehicles,” *Transportation*, vol. 49, no. 5, pp. 1315–1337, 2022.

[96] A. Alho, S. Oh, R. Seshadri, G. Dalla Chiara, W. H. Chong, T. Sakai, L. Cheah, and M. Ben-Akiva, “An agent-based simulation assessment of freight parking demand management strategies for large urban freight generators,” *Research in Transportation Business & Management*, vol. 43, p. 100804, 2022.

[97] K. Martens and I. Benenson, “Evaluating urban parking policies with agent-based model of driver parking behavior,” *Transportation Research Record*, vol. 2046, no. 1, pp. 37–44, 2008.

[98] T. Hagen, N. Reinfeld, and S. Saki, “Modeling of parking violations using zero-inflated negative binomial regression: a case study for berlin,” *Transportation Research Record*, vol. 2677, no. 6, pp. 498–512, 2023.

[99] M. Albergaria and L. P. Fávero, “Narrow replication of fisman and miguel’s (2007a) ‘corruption, norms, and legal enforcement: Evidence from diplomatic parking tickets’,” *Journal of Applied Econometrics*, vol. 32, no. 4, pp. 919–922, 2017.

[100] M. Ramezani and A. H. Valadkhani, “Dynamic ride-sourcing systems for city-scale networks—part i: Matching design and model formulation and validation,” *Transportation Research Part C: Emerging Technologies*, vol. 152, p. 104158, 2023.

[101] H. Yang, C. W. Y. Leung, S. C. Wong, and M. G. H. Bell, “Equilibria of bilateral taxi–customer searching and meeting on networks,” *Transportation Research Part B: Methodological*, vol. 44, no. 8-9, pp. 1067–1083, 2010.

[102] M. Darvish, L. C. Coelho, and R. Jans, “Comparison of symmetry breaking and input ordering techniques for routing problems,” CIRRELT, Québec, QC, Canada, Tech. Rep., 2020.

[103] G. Laporte and S. Martello, “The selective travelling salesman problem,” *Discrete Applied Mathematics*, vol. 26, no. 2-3, pp. 193–207, 1990.

[104] M. A. Boschetti and V. Maniezzo, "Contemporary approaches in matheuristics: an updated survey," *Annals of Operations Research*, pp. 1–38, 2024, online first.

[105] E. Angelelli, R. Mansini, and M. G. Speranza, "Kernel search: A general heuristic for the multi-dimensional knapsack problem," *Computers & Operations Research*, vol. 37, no. 11, pp. 2017–2026, 2010.

[106] J.-Y. Luo and J.-H. Yuan, "A kernel search algorithm for virtual machine consolidation problem in cloud computing," *The Journal of Supercomputing*, vol. 79, no. 17, pp. 19 277–19 296, 2023.

[107] F. Carrabs, M. Cerulli, and D. Serra, "Kernel search approach to solve the minimum spanning tree problem with conflicting edge pairs," *arXiv preprint arXiv:2401.02222*, 2024.

[108] R. Mansini, M. Zanella, and R. Zanotti, "Optimizing a complex multi-objective personnel scheduling problem jointly complying with requests from customers and staff," *Omega*, vol. 114, p. 102722, 2023.

[109] C. Filippi, G. Guastaroba, D. L. Huerta-Muñoz, and M. G. Speranza, "A kernel search heuristic for a fair facility location problem," *Computers & Operations Research*, vol. 132, p. 105292, 2021.

[110] S. Hanafi, R. Mansini, and R. Zanotti, "The multi-visit team orienteering problem with precedence constraints," *European Journal of Operational Research*, vol. 282, no. 2, pp. 515–529, 2020.

[111] Y. Zhang, F. Chu, A. Che, and Y. Li, "Closed-loop inventory routing problem for perishable food with returnable transport items selection," *International Journal of Production Research*, vol. 62, no. 1-2, pp. 501–521, 2024.

[112] C. Archetti, G. Guastaroba, D. L. Huerta-Muñoz, and M. G. Speranza, "A kernel search heuristic for the multivehicle inventory routing problem," *International Transactions in Operational Research*, vol. 28, no. 6, pp. 2984–3013, 2021.

[113] Z. Borčinová, "Kernel search for the capacitated vehicle routing problem," *Applied Sciences*, vol. 12, no. 22, p. 11421, 2022.

[114] D. L. Huerta-Muñoz, C. Archetti, E. Fernández, and F. Perea, "The heterogeneous flexible periodic vehicle routing problem: Mathematical formulations and solution algorithms," *Computers & Operations Research*, vol. 141, p. 105662, 2022.

[115] M. Bruglieri, M. Paolucci, and O. Pisacane, “A matheuristic for the electric vehicle routing problem with time windows and a realistic energy consumption model,” *Computers & Operations Research*, vol. 157, p. 106261, 2023.

[116] J. Gao, F. Zuo, K. Ozbay, O. Hammami, and M. L. Barlas, “A new curb lane monitoring and illegal parking impact estimation approach based on queueing theory and computer vision for cameras with low resolution and low frame rate,” *Transportation research part A: policy and practice*, vol. 162, pp. 137–154, 2022.

[117] S. Gössling, A. Humpe, R. Hologa, N. Riach, and T. Freytag, “Parking violations as an economic gamble for public space,” *Transport Policy*, vol. 116, pp. 248–257, 2022.

[118] T. A. Lambe and M. A. Rebman, “Revenue collection and enforcement at parking lots,” *Socio-Economic Planning Sciences*, vol. 20, no. 2, pp. 95–101, 1986.

[119] E. J. Beltrami and L. D. Bodin, “Networks and vehicle routing for municipal waste collection,” *Networks*, vol. 4, no. 1, pp. 65–94, 1974.

[120] P. Francis and K. Smilowitz, “Modeling techniques for periodic vehicle routing problems,” *Transportation Research Part B: Methodological*, vol. 40, no. 10, pp. 872–884, 2006.

[121] M. Yahyaei, M. Gendreau, and N. Van Omme, “Parking enforcement optimization via kernel search and behavioral–operational coupling.” 2025, submitted for publication in *European Journal of Operational Research*.

[122] H. B. Jung, H. I. Kim, and D. H. Lee, “Team orienteering with possible multiple visits: mathematical model and solution algorithms,” *Computers & Industrial Engineering*, vol. 190, p. 110097, 2024.

[123] B. L. Golden, L. Levy, and R. Vohra, “The orienteering problem,” *Naval Research Logistics*, vol. 34, no. 3, pp. 307–318, 1987.

[124] I. Chao, B. L. Golden, and E. A. Wasil, “The team orienteering problem,” *European Journal of Operational Research*, vol. 88, no. 3, pp. 464–474, 1996.

[125] F. Hammami, “An efficient hybrid adaptive large neighborhood search method for the capacitated team orienteering problem,” *Expert Systems with Applications*, vol. 249, p. 123561, 2024.

[126] F. Y. Vincent, N. Y. Salsabila, S. W. Lin, and A. Gunawan, "Simulated annealing with reinforcement learning for the set team orienteering problem with time windows," *Expert Systems with Applications*, vol. 238, p. 121996, 2024.

[127] M. He, Q. Wu, U. Benlic, Y. Lu, and Y. Chen, "An effective multi-level memetic search with neighborhood reduction for the clustered team orienteering problem," *European Journal of Operational Research*, vol. 318, no. 3, pp. 778–801, 2024.

[128] A. Yahiaoui, A. Moukrim, and M. Serairi, "Grasp-ils and set cover hybrid heuristic for the synchronized team orienteering problem with time windows," *International Transactions in Operational Research*, vol. 30, no. 2, pp. 946–969, 2023.

[129] X. Chou, L. M. Gambardella, and R. Montemanni, "A tabu search algorithm for the probabilistic orienteering problem," *Computers & Operations Research*, vol. 126, p. 105107, 2021.

[130] P. Hansen, N. Mladenović, J. Brimberg, and J. A. M. Pérez, "Variable neighborhood search," in *Handbook of Metaheuristics*. Springer, 2018, pp. 57–97.

[131] M. Tagmouti, M. Gendreau, and J. Potvin, "A variable neighborhood descent heuristic for arc routing problems with time-dependent service costs," *Computers & Industrial Engineering*, vol. 59, no. 4, pp. 954–963, 2010.

[132] D. Shiri, V. Akbari, and A. Hassanzadeh, "The capacitated team orienteering problem: An online optimization framework with predictions of unknown accuracy," *Transportation Research Part B: Methodological*, vol. 185, p. 102984, 2024.

[133] E. A. Ok, *Real Analysis with Economic Applications*. Princeton, NJ: Princeton University Press, 2011.

[134] A. Gunawan, H. C. Lau, P. Vansteenwegen, and K. Lu, "Well-tuned algorithms for the team orienteering problem with time windows," *Journal of the Operational Research Society*, vol. 68, no. 8, pp. 861–876, 2017.

[135] N. Labadie, J. Melechovský, and R. Wolfler Calvo, "Hybridized evolutionary local search algorithm for the team orienteering problem with time windows," *Journal of Heuristics*, vol. 17, no. 6, pp. 729–753, 2011.

[136] A. Gunawan, K. M. Ng, G. Kendall, and J. Lai, "An iterated local search algorithm for the team orienteering problem with variable profits," *Engineering Optimization*, vol. 50, no. 7, pp. 1148–1163, 2018.

[137] B. Jarboui, H. Derbel, S. Hanafi, and N. Mladenović, “Variable neighborhood search for location routing,” *Computers & Operations Research*, vol. 40, no. 1, pp. 47–57, 2013.

[138] S. Toumi, M. Cheikh, and B. Jarboui, “0–1 quadratic knapsack problem solved with VNS algorithm,” *Electronic Notes in Discrete Mathematics*, vol. 47, pp. 269–276, 2015.

[139] T. Wang, H. Lei, J. Wang, H. Liu, and Z. Zhu, “Adaptive variable neighborhood search algorithm with Metropolis rule and tabu list for satellite range scheduling problem,” *Computers & Operations Research*, vol. 170, p. 106757, 2024.

[140] S. N. Bezerra, M. J. F. Souza, and S. R. Souza, “A variable neighborhood search-based algorithm with adaptive local search for the vehicle routing problem with time windows and multi-depots aiming for vehicle fleet reduction,” *Computers & Operations Research*, vol. 149, p. 106016, 2023.

[141] Á. Costa, C. Rocha, and S. Melo, “Parking management policies and the effectiveness of public policy solutions,” *Procedia-Social and Behavioral Sciences*, vol. 111, pp. 965–973, 2014.

[142] M. Manville and M. Pinski, “The causes and consequences of curb parking management,” *Transportation Research Part A: Policy and Practice*, vol. 152, pp. 295–307, 2021.

[143] NYC DOF, “Annual report of New York City parking tickets and camera violations fiscal year 2024,” New York City Department of Finance, Report, 2024. [Online]. Available: https://www.nyc.gov/assets/finance/downloads/pdf/24pdf/2024-local-law-6-report.pdf

[144] R. Arnott and J. Rowse, “Modeling parking,” *Journal of urban economics*, vol. 45, no. 1, pp. 97–124, 1999.

[145] J. Amaya, T. Encarnación, and M. Delgado-Lindeman, “Understanding delivery drivers’ parking preferences in urban freight operations,” *Transportation Research Part A: Policy and Practice*, vol. 176, p. 103823, 2023.

[146] G. D. Chiara, L. Cheah, C. L. Azevedo, and M. E. Ben-Akiva, “A policy-sensitive model of parking choice for commercial vehicles in urban areas,” *Transportation Science*, vol. 54, no. 3, pp. 606–630, 2020.

[147] Z. Xu and X. Sun, "Economic analysis of on-street parking with urban delivery," *Transportation Science*, vol. 58, no. 6, pp. 1300–1318, 2024.

[148] N. Brazil, B. Vang, and H. Abdelnur, "Neighborhood inequality in government fines: The case of parking tickets in 16 us cities," *Cities*, vol. 152, p. 105229, 2024.

[149] X. Sui, Z. Feng, and S. Zhang, "Spatio-temporal heterogeneity in street illegal parking: A case study in new york," *Journal of Transport Geography*, vol. 127, p. 104262, 2025.

[150] X. Peng, R. Song, Q. Cao, Y. Li, D. Cui, X. Jia, Z. Lin, and G.-B. Huang, "Real-time illegal parking detection algorithm in urban environments," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 11, pp. 20 572–20 587, 2022.

[151] Z. Chen, H. Xu, J. Zhao, and H. Liu, "Curbside parking monitoring with roadside lidar," *Transportation Research Record*, vol. 2677, no. 10, pp. 824–838, 2023.

[152] K. L. Judd, *Numerical methods in economics.* MIT press, 1998.

[153] M. Nourinejad and M. J. Roorda, "Parking enforcement: How to optimally choose the citation fine and level of enforcement," University of Toronto Transportation Research Institute, Technical Report, 2017.

[154] T. S. Ferguson and C. Melolidakis, "On the inspection game," *Naval Research Logistics (NRL)*, vol. 45, no. 3, pp. 327–334, 1998.

[155] R. Hohzaki, "An inspection game with multiple inspectees," *European Journal of Operational Research*, vol. 178, no. 3, pp. 894–906, 2007.

[156] P. Hernández and Z. Neeman, "How bayesian persuasion can help reduce illegal parking and other socially undesirable behavior," *American Economic Journal: Microeconomics*, vol. 14, no. 1, pp. 186–215, 2022.

[157] E. Kamenica and M. Gentzkow, "Bayesian persuasion," *American Economic Review*, vol. 101, no. 6, pp. 2590–2615, 2011.

[158] Y. Zhao, Z. Zhou, Q. Pan, and T. Zhou, "G/m/n queuing model-based research on the parking spaces for primary and secondary school," *Discrete Dynamics in Nature and Society*, vol. 2020, no. 1, p. 8870862, 2020.

[159] J. Schmid, X. C. Wang, and A. Conway, "Commercial vehicle parking duration in new york city and its implications for planning," *Transportation Research Part A: Policy and Practice*, vol. 116, pp. 580–590, 2018.

[160] H. Kellerer, U. Pferschy, and D. Pisinger, *Knapsack problems.* Springer, 2004.

[161] C. Morency and M. Trépanier, "Characterizing parking spaces using travel survey data. cirrelt," *CIRRELT, https://www. cirrelt. ca/DocumentsTravail/CIRRELT-2008-15. pdf*, 2008.

[162] Agence de mobilité durable de Montréal. (2025) Paid on-street parking. Accessed: 2025-10-10. [Online]. Available: https://www.agencemobilitedurable.ca/en/parking/paid-on-street-parking

[163] Ville de Montréal, "Plan de déplacement urbain: Plan d'action 2009–2010," Ville de Montréal, Montréal, Québec, Report, 2009, accessed 2025-10-10. [Online]. Available: https://ville.montreal.qc.ca/pls/portal/docs/page/arrond_pmr_fr/media/documents/pdu_plan_action_2009-2010.pdf