UNIVERSITÉ DU LUXEMBOURG

PhD-FDEF-2025-036
The Faculty of Law, Economics and Finance

DISSERTATION

Defence held on 4 December 2025 in Luxembourg

to obtain the degree of

DOCTEUR DE L'UNIVERSITÉ DU LUXEMBOURG

EN SCIENCES ECONOMIQUES

by

Poulad MORADI SHAHMANSOURI
Born on 21 March 1985 in Shahin Shahr, Iran

ASYMPTOTICS AND SUSTAINABILITY IN OPERATIONS MANAGEMENT

Dissertation defence committee:

Prof. Dr. Joachim Arts, Supervisor
Professor, Université du Luxembourg

Prof. Dr. Çağıl Koçyiğit, Chair
Associate Professor, Université du Luxembourg

Prof. Dr. Melvin Drent
Associate Professor, Tilburg University

Prof. Dr. Geert-Jan van Houtum
Professor, Eindhoven University of Technology

Prof. Dr. Michael Jong Kim
Associate Professor, University of British Columbia

To my dearest Aida and our little one, Lena, soon to join us

Acknowledgments

First, I owe my deepest gratitude to my advisor, Joachim Arts, for his invaluable support, mentorship, and guidance throughout the past five years. It is primarily thanks to him that I decided to pursue Operations Research, specifically Applied Probability. His exceptional insight into impactful research directions, wide-ranging interests, and high academic standards have significantly strengthened my abilities as a researcher. He has shaped my academic identity in a profound way, through countless discussions, invaluable guidance, and constant encouragement to think critically and independently. He pushed me to engage with the broader scientific community, encouraging me to present my work, collaborate with others, and develop a clear research vision. He also supported my endeavors outside of research, including my social and community activities. I am deeply grateful for the trust he placed in me and for the freedom he gave me to explore ideas while always providing guidance when needed. I hope to have the opportunity to collaborate with him in the future.

Next, I am deeply thankful to my co-advisor, Melvin Drent, for his support and friendship since my master's studies. His belief in my potential early on played a decisive role in starting my academic development. I am especially grateful for his availability, patience, and interest in my progress. It is a privilege to continue working with him as a postdoctoral researcher and I eagerly look forward to it.

I sincerely thankful for the members of my doctoral committee and for their time, guidance, and thoughtful feedback. Special thanks to Çağıl Koçyiğit for chairing the committee, supporting me throughout my doctoral studies, and for inspiring me to explore data-driven optimization at the intersection of optimization and machine learning. I am grateful to Michael Kim for traveling from Canada to attend my defense, for his valuable feedback, and for our stimulating discussions during his sabbatical at the University of Luxembourg, which led to the development of the final chapter. I also want to thank Geert-Jan Van Houtum for attending my defense in person and for his valuable comments which have strengthened my work.

I extend my sincere gratitude to Josué Velázquez Martínez, my coauthor for Chapter 3, for his guidance, thoughtful insight, and inspiring collaboration. His knowledge and professionalism have profoundly enriched my work. I enjoy our collaboration and hope we can continue working together in the future. Chapter 2 builds upon my master's thesis

conducted with Ranit Sinha, and I am grateful for his contributions.

My heartfelt appreciation goes to colleagues and friends at LCL and the University of Luxembourg. The many coffee breaks, meals, ping pong matches, gatherings, and other interactions made my PhD journey enjoyable. In particular, I thank Benny Mantin, Nils Löhndorf, Anne Lange, Francesco Viti, Steffen Klosterhalfen, Jackie Brown, Daria Pawlowska, Raquel Rubio, Carla Rosen-Vacher, Roozbeh Qorbanian, Tiffany Nguyen, Elshan Sarkarfarshi, Rishikesh Parma, Saber Mousavi, Neeraj Podichetty, Junxia He, Kristof Horvat, Tianxiong Yu, Laura Palacios, Matteo Cosmi, Nihat Oner, Daniel Dobos, Sarah van der Auweraer, Bikey Seranilla, You Wu, Nicole Perez-Becker, Laurens Deprez, Farid Alavi, Saswat Patra, Sandria Weisshuhn, Ganesh Balasubramanian, Ayush Gupta, Mohammad Namakshenas, Tom Rauber, Qiaoke Zhang, Yang Wang, Celine Barlier, Sybille Barvaux, Meztli Matadamas, Izabela Silva, Menglin Zheng, and Kartikeya Singh (apologies if I have missed anyone). Their support was essential to completing this thesis.

My visit to TU Eindhoven in Fall 2024 was inspiring. I thank Melvin Drent for hosting me and the OPAC group members for making the experience memorable. This visit led to a later collaboration with Collin Drent on dual-sourcing inventory systems. I am grateful to Collin and Melvin for including me in the project and later offering a postdoc position in the OPAC group. Although we paused the work while I focused on this thesis, I look forward to resuming it when I join the OPAC group in early 2026. I also thank Alireza Yazdani, Sina Shahri, and Aran Nasiri for making my visit so enjoyable.

Throughout my PhD journey, engaging with distinguished scholars has been instrumental in shaping this thesis. I am particularly thankful to Willem van Jaarsveld and Xiting Gong for their stimulating and insightful discussions on lost-sales inventory systems.

Some unfinished projects during my PhD, though not included in this thesis, form the foundation of future research. I am particularly grateful to Nan Zhang, Tahmineh Dastan, and Thao Nguyen and hope to continue these collaborations.

On a personal note, I am thankful for my family and friends. Their love, patience, and encouragement sustained me throughout this journey. I am deeply thankful to my mom, Nikoo, for her sacrifices and unwavering support. She encouraged me to pursue a PhD. I also thank my brother, Babak, for his constant support.

Above all, I am deeply grateful to my wife, Aida, for her love, understanding, and support. She shared every uncertainty, long hours, and emotional challenge of this journey, offering patience, encouragement, and faith at every step. This achievement would not have been possible without her. We are also joyfully expecting a baby girl, Lena, whose imminent arrival shortly after completing my PhD makes this milestone even more special and deepens my gratitude for Aida's unwavering support.

Contents

Chapter 1

Introduction

Organizations today operate in increasingly complex, uncertain, and competitive environments. Operations management equips them with analytical tools to guide informed and effective operational decisions. For example, retailers must effectively manage lost sales to maintain competitiveness. Producers should dynamically select supply modes, and logistics companies must incorporate new features into transportation models to ensure resilience, cost efficiency, and sustainability. Advanced critical systems must remain operational at low maintenance cost by utilizing abundant real-time sensor data.

Although these problems arise in different contexts, they can be expressed as mathematical models that reveal their underlying similarities. These mathematical models, however, can become highly complex in representing real-world problems, which makes the identification of optimal decisions difficult. For instance, calculating the optimal ordering decision for a product is straightforward under the assumption that customers wait during a stockout, but dropping this assumption generally transforms the problem into a far greater challenge. Some simplifying assumptions, previously regarded as insignificant, can now pose major risks to organizations. This situation requires research to develop effective and practical solutions for more complex and realistic settings.

Economic efficiency, reliability, and sustainability are among the core performance metrics that define excellence in operations. While the first two metrics have long been considered in nearly all operational decisions, sustainability only became an important dimension in the 21st century. Early research in the 1980s, such as Hansen (Hansen and Lebedeff, 1987), provided evidence that CO_2 emissions lead to a rise in global temperatures and emphasized their environmental consequences. Across operational domains, transportation and industry each represent about 22% of the European Union's CO_2 emissions (European Environment Agency, 2021). Sustainability in operations also includes resource efficiency, which covers the reduction of energy, water, and material consumption in production and logistics processes. Improved resource efficiency reduces waste, lowers operational costs, supports environmental objectives, and strengthens approaches that extend asset life. These observations highlight the critical need to incorporate sustain-

ability into decision making processes. However, sustainability goals introduce additional costs that organizations must take into account in their operational decisions. Balancing these costs against efficiency, reliability, and other performance objectives requires careful consideration of their trade-offs. Optimization models provide a natural framework to evaluate such trade-offs and help decision makers identify solutions that achieve environmental targets while preserving operational and economic performance.

Operational systems are inherently subject to uncertainty in some decision parameters such as demand, supply, and equipment status. Stochastic models capture these uncertainties and allow researchers to model variability and to evaluate the performance of different decisions under more realistic conditions. Many stochastic models become highly complex as a key problem parameter grows, which makes exact solution methods infeasible. Asymptotic analysis provides a powerful approach to address this challenge by characterizing the behavior of the system under limiting regimes, such as the high cost of losing a sale in inventory models, the high payload in routing models, and the long lifetime of components in maintenance models. This method allows for the derivation of asymptotically optimal policies that maintain strong performance in non-asymptotic settings, without the computational expense of solving a high-dimensional stochastic model.

This thesis addresses two topics: integrating sustainability into certain operations management problems and developing tractable policies through asymptotic analysis of high-dimensional stochastic systems.

1.1 Contributions of the thesis

This thesis consists of four independent chapters, which can be categorized according to the relevant literature as follows: Chapters 2 and 4 focus on stochastic inventory models, with Chapter 2 on inventory models with two supply modes and Chapter 4 on lost sales inventory systems. Chapter 3 studies the green routing problem, and Chapter 5 covers condition-based maintenance of stochastically deteriorating systems. A brief overview of each chapter is provided below.

1.1.1 Efficient emission reduction through dynamic supply mode selection

Chapter 2 studies the inbound supply and inventory decisions of a company sourcing multiple products through a third-party logistics (3PL) provider. Each product can be shipped using one of two transport modes that vary in cost, lead time, and emissions. Demand is stochastic for all products. The company makes periodic decisions on order quantities for both transport modes to minimize long-run average holding, backorder, and transportation costs while ensuring that total greenhouse gas (GHG) emissions remain

below a specified target. Assortment-wide constraints are likely to become increasingly common, either through voluntary measures or government regulation. To address this problem, we adopt the dual-index policy of Veeraraghavan and Scheller-Wolf (2008) that dynamically orders shipments to both modes based on on-hand and outstanding inventory. The decision problem is formulated as a non-linear non-convex integer program, which we reformulate into a large integer programming problem and solve through the Dantzig-Wolfe decomposition. A comprehensive numerical study based on data from multiple industries demonstrates the effectiveness of this approach. Our results indicate that dynamic selection of transport modes leads to significant cost reductions compared to static policies that rely on a single mode. We also show that applying a single, aggregate emission target across the entire product assortment yields markedly better economic and environmental outcomes than assigning separate emission targets to individual products. Furthermore, our analysis shows that the difference in emissions relative to the difference in costs between transport modes is the key factor that determines the potential for emission reduction. Chapter 2 is based on Drent et al. (2023b).

1.1.2 Load asymptotics and dynamic speed optimization for the greenest path problem: a comprehensive analysis

Chapter 3 investigates the influence of high-resolution elevation data on identifying the most fuel-efficient, or greenest, paths for trucks that operate in urban environments with diverse topographies. Using a variant of the Comprehensive Modal Emission Model (CMEM), we show that both the greenest path and the optimal driving speed depend on the slope of the road and the payload of the vehicle. On flat networks, the shortest path (or the fastest in congestion) with a constant speed minimizes fuel consumption, but in networks with varying elevations, the shortest path may not be the greenest one, and the speed must be adjusted dynamically along the route. This result highlights that the traditional practice of selecting paths between origin and destination as a preprocessing step for pollution routing problems (PRPs) can lead to suboptimal solutions, and that the path selection should instead be integrated directly within the routing model. Nonetheless, the integration of path selection adds further complexity to the PRP, which is inherently NP-hard. We show that the greenest path converges to an asymptotic route as payload increases, and this limiting path is attained at finite loads. This property can be used to approximate the greenest paths as input to PRP models.

We carry out comprehensive numerical experiments under free-flow conditions using geo-spatial data from 25 cities across six continents. We evaluate the advantages of the greenest path selection and dynamic speed optimization. Our results show that the greenest paths quickly diverge from the shortest ones, and that dynamic speed optimization significantly reduces CO_2 emissions. We perform an analysis of variation (ANOVA) on our results and find that the potential for emission reduction is primarily driven by vari-

ations in road gradients and the relative elevation between origin and destination. We conduct additional numerical experiments using traffic data for New York City during rush-hour periods to evaluate the CO_2 reduction potential of the greenest path and dynamic speed optimization, in traffic congestion, compared to the fastest path with the traffic speed. Our results indicate that the selection of the greenest paths can substantially reduce CO_2 emissions when traffic congestion occurs. Moreover, while adjusting speeds on uphill segments provides notable emission savings, opportunities to exploit downhill acceleration are constrained by traffic congestion. This last result highlights the advantages of scheduling truck deliveries during off-peak periods to reduce emissions. Chapter 3 is based on Moradi et al. (2024).

1.1.3 Asymptotic optimality of projected inventory level policies for lost sales inventory systems with large leadtime and penalty cost

Chapter 4 investigates the performance of the projected inventory level (PIL) policy in single-item, single-echelon, periodic-review lost sales inventory systems with positive leadtime and independent and identically distributed (i.i.d.) demand under the long-run average cost criterion (canonical problem). In these systems the unmet demand is lost and generates penalty costs. The canonical lost sales inventory problem is a classical challenge in inventory theory, since the optimal ordering policy can generally be obtained by solving a high-dimensional dynamic programming problem but this approach suffers from the "curse of dimensionality". Consequently, a key area of research in inventory theory focuses on designing tractable and high-performing control policies for the canonical problem that are asymptotically optimal in certain regimes.

Prior literature has established that base stock policies are asymptotically optimal for the canonical problem when the cost of losing a sale grows large (Huh et al., 2009; Bijvank et al., 2014), whereas constant order policies achieve asymptotic optimality as the leadtime grows large (Goldberg et al., 2016; Xin and Goldberg, 2016). However, neither policy is simultaneously optimal across both regimes. The capped base stock policy (Xin, 2021) is a hybrid policy that reconciles these limitations by using a base stock policy but individual orders cannot a exceed a cap that is the second policy parameter. More recently, the PIL policy has been proposed (van Jaarsveld and Arts, 2024), which dynamically adjusts orders to maintain the expected inventory level at the time of replenishement order receipt, leveraging the probabilistic information available at each decision epoch. van Jaarsveld and Arts (2024) show that the PIL policy outperforms the base stock policy for a general demand process under mild technical assumptions and outperforms the constant order policy for exponential demand. Their empirical results indicate that the PIL policy outperforms both base stock and constant order policies, as well as capped base stock polcies across a range of demand distributions.

In Chapter 4, we analyze the asymptotic behavior of the PIL policy under general demand as both the leadtime and lost sales penalty costs become large. Under mild assumptions on the demand distribution, we show that the difference in long-run average cost between the PIL policy and the constant order policy remains bounded. This finding implies that the PIL policy attains asymptotic optimality as the leadtime increases with sufficiently high penalty costs, which is new to the literature on lost sales inventory theory. Our analysis relies on demonstrating that the relative value function under the constant order policy satisfies the Wiener-Hopf equation. We solve this equation using ladder process techniques applied to the random walk induced by the constant order policy. Then, we apply a one-step policy improvement argument to establish our main convergence result.

1.1.4 Risk or replace: efficient asymptotics for data-driven maintenance

Condition-based maintenance (CBM) is an approach that plans interventions for deteriorating systems according to their observed operational state. CBM reduces unplanned downtime, extends usable lifetime, avoids unnecessary replacement, and mitigates the environmental impact associated with manufacturing and transporting new equipment. Consequently, it has gained widespread adoption in recent decades. In Chapter 5, we study a heterogeneous population of components that degrade over time. We consider a general class of stochastic degradation processes with non-negative and i.i.d. increments that are characterized by component-specific parameters that remain unobservable to the decision maker. We rely on degradation data to estimate these parameters and determine replacement actions at equidistant epochs. The goal is to minimize the long-run average cost, which incorporates fixed replacement costs, failure costs, and operating costs that increase as components degrade. This problem can be formulated as a high-dimensional partially observable Markov decision process (POMDP), which is generally intractable due to the "curse of dimensionality." Drent et al. (2023a) develop a tractable POMDP model for a degradation process that belongs to the one-parameter exponential family, called the Integrated Bayes Policy.

We develop a tractable, data-driven CBM policy that estimates the optimal policy of a hypothetical Oracle that has full information of the underlying degradation parameters. We call this policy the Estimated Oracle's Optimal Policy (EOP). The tractability of the EOP directly depends on the tractability of the underlying parameter estimation method. We introduce a scaling regime where both the failure thresholds and cost parameters increase proportionally, reflecting practical settings in which component lifetimes and maintenance costs are large relative to the observation frequency of modern sensors. In this regime, POMDP-based methods rapidly lose tractability. We show that the regret of the EOP, defined as the difference between its long-run average cost and that of the

Oracle, converges to zero in the scaling regime when the parameter estimator is consistent. We benchmark our EOP against the Integrated Bayes Policy using real degradation data; our policy performs excellently. Through extensive numerical experiments on both discrete and continuous state spaces, we show that the EOP achieves remarkably low regret across all settings. In instances where the optimal policy can be computed using the POMDP method, the optimality gap of the EOP is statistically indistinguishable from zero.

1.2 Organization of the thesis

Table 1.1 summarizes the main literature streams associated with each chapter, while Table 1.2 outlines the key methodologies used. Each chapter is self-contained and can be read independently.

Table 1.1: Navigating the thesis based on related literature stream and scope of analysis.

	Topic					Scope	
Chapter	Sustainability	Lost sales	Dual sourcing	Maintenance	Transportation	Single item	Multi item
2	x		x		x		x
3	x				x		
4		x				x	
5	x			x		x	

Table 1.2: Navigating the thesis based on main methodology.

Chapter	Asymptotics	Random walk	Renewal theory	Queueing theory	Markov decision theory	Bayesian learning	Column generation	Routing	Simulation
2							x		x
3	x							x	x
4	x	x	x	x	x				
5	x		x		x	x			x

Chapter 2

Efficient emission reduction through dynamic supply mode selection

2.1 Introduction

The transportation sector has consistently been one of the most polluting European sectors for more than a decade now, and it is projected to remain so for the foreseeable future (European Environment Agency, 2020). This, unfortunately, appears to be a trend that stretches beyond Europe. Recent analysis indicates that the G20 countries, currently responsible for 80% of the global greenhouse gas (GHG) emissions, will see an increase of 60% in their transportation sector emissions by 2050 (Vieweg et al., 2018). Prominent global climate targets, such as the ones outlined in the Paris agreement, will soon become unattainable (European Environment Agency, 2020; United Nations Environment Programme, 2020).

In light of the above, the European Union (EU) announced the Green Deal in 2019, a framework containing climate targets and policy initiatives that sets the EU on a path to reach carbon neutrality by 2050. The Green Deal is legally enshrined in the European Climate Law, which states that member states are legally committed to meet the targets, and face penalties in case they do not meet these targets. Being among the most polluting sectors, a key part of the Green Deal relates to policy initiatives that impact the transportation sector. For instance, the EU plans to extend the European emissions trading scheme (EU ETS) to include both road and maritime transport (Abnett, 2020; Commission, 2020). Under the ETS, which until now includes only air transport, the EU enforces a cap on the total amount of GHG emissions from sectors covered by the scheme. The EU also investigates whether to increase fossil fuel taxation, thereby effectively raising the price of GHG emissions. More and more companies are reducing their emissions voluntarily as part of their corporate social responsibility. If not penalized by governments, companies that excessively pollute might still lose revenues as environmentally conscious customers take their business elsewhere (Dong et al., 2019).

The developments described above highlight the urgency for companies to explicitly incorporate GHG emissions in their supply chain decision making. In this chapter, we study the inbound supply mode and inventory decision making of a company that sells an assortment of products which are sourced from outside suppliers. The company wishes to keep the total GHG emissions associated with using different supply modes across the assortment below a certain target level in the most economically viable manner. As is often the case in practice, the company may rely on a third party logistics (3PL) provider for the inbound transport of the products. 3PLs typically offer several transport modes for the transportation of products – these may differ in terms of transportation costs, transit times, and GHG emissions. Alternatively, there may be different suppliers (e.g. a near and offshore supplier) for a product that naturally have different costs and transportation emissions. In the remainder of the chapter, we will use the terminology of a 3PL provider that offers multiple transportation modes.

The company can utilize the heterogeneity in the fleet of the 3PL to its advantage. While some transport modes are low emitting but slow, others may be fast but result in more emissions. Fast transport modes also typically come at the expense of a cost premium, and yet they are often relied upon when responsiveness is required (e.g. in case of imminent stock outs). Thus the company should rely dynamically on both transport modes. Implementing this holistically across the entire assortment of products allows the company to reduce emissions significantly for products for which it is relatively cost-efficient to do so and less for products for which this is more expensive. It additionally enables the company to reduce the total inventory and transportation costs by shipping the majority of products with a relatively cheap but slow transport modes while simultaneously resorting to faster but more expensive and often more polluting transport modes whenever expedited shipments are needed. While the advantages of dynamically selecting different transport modes are evident, two important and interrelated questions remain:

1. When should the company ship how many units of which product with what transport mode?

2. What is the value of dynamically shipping products with different transport modes?

These questions are interesting but also intricate when one wishes to answer them for an entire assortment of products where the combined total of GHG emissions from transportation must not exceed a certain target level.

To tractably answer the questions above, we focus on the setting where the 3PL offers two distinct transport modes for the transport of each product (or, equivalently, the setting where the company has already decided on the two transport modes for each product). These transport modes need not be the same for every product; they will depend on the characteristics of the suppliers as well as the 3PL (e.g., some products can be transported using aircraft or rolling stock while others can be shipped via inland

waterways or ocean shipping). We consider long distance and/or high volume transport lanes where the impact of the transport mode decisions of any particular individual shipper on the actual shipping and carbon footprint is negligible. The company decides periodically how many units it wishes to transport with what transport mode and incurs mode specific unit transportation costs. Shipments arrive at the company after a deterministic transit time that depends on the transport mode that is used. Demand for each product in every period is stochastic and independent and identically distributed across periods. Any demand in excess of on-hand inventory is backlogged and satisfied in later periods. The company incurs per unit holding and backorder penalty costs, and the specific cost parameters may vary from product to product. The company seeks to minimize the long-run average holding, backorder, and transportation costs while keeping the total long-run average GHG emissions from transportation of the entire assortment below a certain target level.

It is well-known that the optimal policy for the inventory system described above is complex, even in the simplest case of a single product and absent of the emission constraint (Whittemore and Saunders, 1977; Feng et al., 2006). For the control of each product, we therefore use a heuristic policy that is originally due to Veeraraghavan and Scheller-Wolf (2008). They show that their so-called dual-index policy performs quite well compared to the optimal policy. The dual-index policy tracks two inventory positions for each product: The slow inventory position, which equals the on-hand inventory plus all in-transit products minus backlog, and the fast inventory position, which is defined similarly but includes only those in-transit products that are due to arrive within the transit time of the fastest transport mode. Under the dual-index policy, we place orders with both modes such that these inventory positions are kept at (or above) certain target levels, also referred to as base-stock levels. As such, the dual-index policy dynamically prescribes shipment quantities for both transport modes based on the net inventory level and the number of products that are still in-transit. To find the optimal base-stock levels for the entire assortment of products, we formulate the decision problem as a non-linear non-convex integer programming problem. A partition reformulation of this problem allows us to use column generation techniques to solve the decision problem. These techniques enable us to decompose the complex multi-product decision problem into simpler sub-problems per product. Leveraging a separability result of Veeraraghavan and Scheller-Wolf (2008), we show that this sub-problem constitutes a special Newsvendor problem that can be solved efficiently through a simulation-based optimization procedure.

The main contributions of this chapter are:

1. We are the first to study dynamic mode selection for an assortment of products with stochastic demand where the total average GHG emissions from the inbound transport of those products must be kept below a certain target level.

2. We provide a tractable optimization model that finds a tight lower bound on the

optimal solution as well as near-optimal feasible solutions within reasonable time. We show that our mathematical formulation of the decision problem allows us to decompose the non-linear non-convex integer programming problem into sub-problems per product. We leverage results from Veeraraghavan and Scheller-Wolf (2008) to show that the sub-problems can be solved efficiently through a one dimensional search procedure in which each instance constitutes a Newsvendor type problem that is readily solved through simulation.

3. We perform an extensive computational experiment based on data from different industries. Through these experiments:

   i. We establish the value of dynamic mode selection by comparing our model with a model in which only one transport mode per assortment product can be used. This value can go up to 15 percent in cost savings;

   ii. We show that decomposing an aggregate carbon emission reduction target into targets for each product in the assortment individually is financially detrimental. Our holistic approach can lead to cost savings of over 40 percent relative to the approach with reduction targets per individual product;

   iii. We find that the emission differences between transport modes relative to the cost difference between modes is the main determinant of emission reduction potential for a given assortment. In our experiments we find that, 20 percent of the products for which this ratio is highest contribute between 59 and 94 percent of the emission reduction.

The remainder of this chapter is organized as follows. In Section 2.2, we review the existing literature and position our work within the literature. Section 2.3 contains the model description as well as the mathematical formulation of the decision problem. A column generation procedure to solve the decision problem is provided in Section 2.4. We subsequently report on an extensive computational experiment in Section 2.5, and we provide concluding remarks in Section 2.6.

2.2 Related literature

This chapter integrates carbon emissions from inbound transportation into an inventory system with two supply modes. As such, our work contributes to the large stream of literature that studies multi-mode or multi-supplier inventory systems. For an excellent overview of such systems, we refer the reader to the review papers of Thomas and Tyworth (2006), Engebrethsen and Dauzére-Pérés (2019) and Svoboda et al. (2021), and references therein. We also contribute to the extensive body of literature that revolves around the integration of environmental aspects into supply chain decision making; see Dekker et al. (2012), Brandenburg et al. (2014), Barbosa-Póvoa et al. (2018), and references therein,

for an overview of this field. In what follows, we focus on contributions that are most relevant to the present chapter.

The decision how many products to order from which supplier is considered a canonical problem in the inventory management literature. It has been studied extensively since the sixties, mostly under the assumption that lead times are deterministic, that unmet demand is backlogged, and that only two distinct suppliers are at the disposal of the decision maker; the fastest being more expensive than the slowest. Fukuda (1964) and Whittemore and Saunders (1977) were the first to study this system. Assuming periodic review, they show that its optimal policy is a simple base-stock rule only under the assumption that the difference between the lead times of both suppliers is one period. For general lead time differences, the optimal policy is complex and can only be computed through dynamic programming for small instances. Since then, most researchers have focused on developing well-performing heuristic polices for which the best control policy parameters can be tractably obtained.

In this chapter, we rely on the so-called dual-index policy to decide upon the shipment sizes for both transport modes for each product. Under this policy, which is originally due to Veeraraghavan and Scheller-Wolf (2008), two different inventory positions are kept track off: One that includes all outstanding shipments and one that includes only those outstanding shipments that are due to arrive within the lead time time of the fastest mode. Veeraraghavan and Scheller-Wolf (2008) show numerically that the dual-index policy performs well compared to the optimal policy. In fact, Drent and Arts (2022) show that the dual-index policy is asymptotically optimal as the cost of the fastest transport mode and the backorder penalty cost become large simultaneously. The policy has received quite some attention in recent years (see, e.g., Sheopuri et al., 2010; Arts et al., 2011; Sun and Van Mieghem, 2019). We employ the dual-index policy because it is intuitive, has good performance, and can be optimized efficiently. Unlike the present work, the dual-index policy has so far been studied exclusively in single product settings under conventional cost criteria absent of any emission considerations.

Within the transportation literature, inventory systems with multiple transport modes have received considerable attention too. To properly embed the present work in the existing literature, we group contributions to this field into two categories depending on the modelling assumptions regarding the usage of the available transport modes (c.f. Engebrethsen and Dauzére-Pérés, 2019). The first category, which we refer to as *dynamic* mode selection, is concerned with inventory systems in which multiple transport modes are used simultaneously over a given (possibly infinite) planning horizon. Since we study an infinite horizon periodic review inventory model in which products can be transported with two distinct modes in each period, our work falls into this category – as do all the inventory papers with two suppliers described so far. Only few papers exist in this category that explicitly account for carbon emissions, and the few that do differ substantially from the present work in terms of modelling choices as well as analysis. They either assume

deterministic demand and a finite horizon (Palak et al., 2014) or study the closely related yet different problem of splitting an order among several transport modes (Konur et al., 2017). While not explicitly modeling carbon emissions, Dong et al. (2018) and Lemmens et al. (2019) also study the benefit of dynamically switching between multiple transport modes in the context of multi-modal transport. They show that this can lead to more usage of less polluting transport modes without compromising on costs or responsiveness. Different from our work, all papers mentioned above consider the inventory control and transport mode decisions for a single product only.

The second category concerns inventory systems in which a single transport mode is selected a priori at the start of a planning horizon; all replenishment orders until the end of that planning horizon are then shipped with this mode. We refer to this category as *static* mode selection. Two papers belonging to this category are particularly relevant to our work. Hoen et al. (2014b) study a periodic review inventory system under backlogging where inbound transport is outsourced to a 3PL that offers multiple transport modes. Assuming base-stock control for each mode, they are interested in selecting the transportation mode that leads to the lowest long-run average total cost consisting of holding, backlogging, ordering, and emission costs. For calculating transportation emissions, they rely on the well-known NTM methodology (we discuss this methodology in more detail later in Section 2.3.1 and Appendix 2.A). We extend Hoen et al. (2014b) in two important directions. First, we move from static to dynamic mode selection, thereby incorporating the flexibility to dynamically switch between different transport modes for each product. Second, we consider an assortment of products under a single constraint on the total average transportation emissions from those products. Hoen et al. (2014a) consider a similar constraint in a multi-product variant of the setting of Hoen et al. (2014b) under the assumption that demand is deterministic and inversely related to the price set by the decision maker. They show that because of the portfolio effect of such an assortment-wide emission constraint, carbon emissions from transportation can be reduced substantially at hardly any additional cost.

The dual-index policy studied in this chapter has the appealing feature that it can mimic static mode selection. This is useful in our computational experiment where we establish the added value of dynamic mode selection over static mode selection. A closely related paper in that respect is Berling and Martínez de Albéniz (2016) who study dynamic speed optimization of a single transport mode in a single-product stochastic inventory problem. They show that the value of dynamically controlling the speed of outstanding shipments, as opposed to a static speed policy, can be significant, both financially and from a carbon emission perspective.

Our review so far has almost exclusively revolved around papers on multi-period inventory systems. We note that there is also a stream of literature that integrates carbon emissions into single period multi-supplier models, see e.g., Rosič and Jammernegg (2013), Arıkan and Jammernegg (2014), and Chen and Wang (2016). Similar to the majority of

the papers discussed so far, these papers focus on single-product settings.

2.3 Model description

In this section, we first provide a description of the inventory system under consideration and introduce the notation that we use throughout this chapter. We then describe the policy we propose to dynamically ship products with two transport modes. We conclude with providing a mathematical formulation of the decision problem.

2.3.1 Description and notation

We consider a company that sells an assortment of products. The inventories for these products are replenished from external suppliers through a third party logistics provider (3PL). A 3PL often offers several transport modes. We focus on the setting where the company has already decided upon two distinct transport modes that it would like to use for the transport of each product. These two transport modes will differ in terms of costs, lead times, emissions, or a combination thereof. Given these two transport modes for each product, the operational question that remains is how many units of each product the company should transport using which transport mode at what time so that costs –holding, backlog, and ordering– are minimized and an overall emission constraint is met. Companies will increasingly impose such constraints, either voluntarily or due to government regulation.

The inventory system runs in discrete time with $t \in \mathbb{N}_0$ denoting the period index. Without loss of generality, we assume that the period is of unit length and coincides with the review epoch. Let $J = \{1, 2, \ldots, |J|\}$ denote the nonempty set of products that the company offers for sale. Demand for product $j \in J$ across periods is a sequence of non-negative independent and identically distributed (i.i.d.) random variables $\{D_j^t\}$. Any demand in excess of on-hand inventory is backlogged. Let I_j^t denote the net inventory level (on-hand inventory minus backlog) of product j at the beginning of period t after any outstanding orders have arrived. Each unit of product j in on-hand inventory $(I_j^t - D_j^t)^+$ carries over to the next period and incurs a holding cost $h_j > 0$. Similarly, each unit of product j in backlog $(D_j^t - I_j^t)^+$ incurs a penalty cost $p_j > 0$. Here we use the standard notation $x^+ = \max(0, x)$.

Each product can be shipped using two distinct transport modes from one supplier (or, equivalently, using one or two distinct transport modes from two distinct suppliers). Let $M = \{f, s\}$ denote the set of available transport modes, where we use f and s to refer to the faster and slower transport mode, respectively. Associated with the transport of one unit of product $j \in J$ with mode $m \in M$ is a cost $c_{j,m} \geq 0$, a deterministic lead time $l_{j,m} \in \mathbb{N}_0$, a distance traveled from the supplier to the company $d_{j,m} > 0$, and a certain

number of units CO2 emission $e_{j,m} \geq 0$. The weight of one unit of product j is denoted $w_j > 0$. Recall that the company outsources its transport to a 3PL provider and hence has no control over the actual shipping. We therefore consider variable emissions that depend only on product and transport mode specific characteristics as well as on distance traveled, and we refrain from incorporating a fixed emission factor per actual shipment. This is a reasonable assumption for long distance and/or high volume transport lanes where the impact of the decisions of any particular individual shipper on the carbon footprint are negligible. In line with previous literature that models transportation emissions in the context of mode selection (e.g., Hoen et al., 2014a,b), we endow $e_{j,m}$ with the following structure which is based on the NTM methodology:

$$e_{j,m} = w_j(a_m + d_{j,m}b_m), \tag{2.1}$$

where $a_m \geq 0$ and $b_m > 0$ are a fixed and variable transport mode specific emission constant, respectively. Consistent with the NTM methodology, we assume that each product is shipped with an averagely loaded transport mode. We define the lead-time difference between the fast and slow mode as $l_j = l_{j,s} - l_{j,f} \geq 0$ for each product $j \in J$. Conventional literature on dual model problems (e.g., Sheopuri et al., 2010) imposes the assumption that the cost premium of using the fast mode does not exceed the lead-time difference multiplied by the penalty costs, i.e., $(c_{j,f} - c_{j,s}) < l_j p_j$ to ensure that using the fast mode is attractive. We do not impose this assumption as whether using the fast supply mode will also depend on the target carbon reduction. When the fast supply mode is less polluting than the slow mode, the fast mode may become attractive even when $(c_{j,f} - c_{j,s}) < l_j p_j$. Conversely, the fast supply mode may become unattractive when the fast supply mode is more polluting, even when $(c_{j,f} - c_{j,s}) < l_j p_j$. Thus our model allows for situations where, e.g., the expensive transport mode is either the fastest and most polluting or the fastest and least polluting. Finally, the amount of items of product $j \in J$ to be shipped with transport mode $m \in M$ in period t is denoted by $Q^t_{j,m}$. With this notation, observe that shipments $Q^{t-l_{j,f}}_{j,f}$ and $Q^{t-l_{j,s}}_{j,s}$ arrive in period t so that we can write the following recursion for the inventory level I^t_j of each product j:

$$I^t_j = I^{t-1}_j - D^{t-1}_j + Q^{t-l_{j,f}}_{j,f} + Q^{t-l_{j,s}}_{j,s}.$$

All notation introduced so far as well as notation that we will introduce later is summarized in Table 3.1.

2.3.2 Control policy

It is well-known that even for the simplest case where $|J| = 1$ and absent of the emission constraint, the policy that prescribes the optimal shipment quantities is complex and can only be computed through dynamic programming for very small instances that are arguably not representative for practice (Whittemore and Saunders, 1977; Feng et al.,

Table 2.1: Overview of notation.

Notation	Description
Sets	
J	Assortment; Set of all products.
M	Set of available transport modes, i.e. $M = \{f, s\}$.
Input	
D_j^t	Random demand for product $j \in J$ in period $t \in \mathbb{N}_0$.
p_j	Penalty cost for one unit of product $j \in J$ in backlog carried over to the next period.
h_j	Holding cost for one unit of product $j \in J$ in on-hand inventory carried over to the next period.
$c_{j,m}$	Cost of shipping one unit of product $j \in J$ with transport mode $m \in M$.
$l_{j,m}$	Transportation lead time for product $j \in J$ by transport mode $m \in M$.
l_j	Transportation lead time difference between the fast and slow mode for product $j \in J$, i.e. $l_j = l_{j,s} - l_{j,f}$.
$d_{j,m}$	Distance for the transport of product $j \in J$ with transport mode $m \in M$.
a_m	Fixed emission constant corresponding with transport mode $m \in M$.
b_m	Variable emission constant corresponding with transport mode $m \in M$.
w_j	Unit weight of product $j \in J$.
$e_{j,m}$	Total units CO2 emission associated with shipping one unit of product $j \in J$ with transport mode $m \in M$.
$\mathcal{E}^{max}$	The maximally allowable carbon emissions for the transport of the entire assortment of products.
Decision variables	
$S_{j,m}$	Base-stock level for product $j \in J$ and transport mode $m \in M$.
Δ_j	Difference between the slow and fast base-stock level for product $j \in J$, i.e. $S_{j,s} - S_{j,f}$.
$\mathbf{S}_f$	The vector $(S_{1,f}, S_{2,f}, \ldots, S_{\|J\|,f})$.
$\boldsymbol{\Delta}$	The vector $(\Delta_1, \Delta_2, \ldots, \Delta_{\|J\|})$.
State variables	
I_j^t	Inventory level of product $j \in J$ at the beginning of period $t \in \mathbb{N}_0$ after orders have arrived.
$IP_{j,f}^t$	Fast inventory position of product $j \in J$ in period $t \in \mathbb{N}_0$ before shipping orders.
$IP_{j,s}^t$	Slow inventory position of product $j \in J$ in period $t \in \mathbb{N}_0$ after shipping orders with the fast transport mode.
$Q_{j,m}^t$	Amount of product $j \in J$ shipped with transport mode $m \in M$ in period $t \in \mathbb{N}_0$.
O_j^t	The overshoot of product $j \in J$ in period $t \in \mathbb{N}_0$, i.e. $(IP_{j,f}^t - S_{j,f})^+$.
Output of model	
$C(\mathbf{S}_f, \boldsymbol{\Delta})$	Total long-run average holding, backlog, and ordering costs under a given control policy $(\mathbf{S}_f, \boldsymbol{\Delta})$.
$E(\mathbf{S}_f, \boldsymbol{\Delta})$	Total emissions from transportation under a given control policy $(\mathbf{S}_f, \boldsymbol{\Delta})$.
$C_P^{UB}(C_P^{LB})$	Upper (lower) bound for the optimal solution to Problem (P).

2006). For the control of this inventory system, we therefore use a heuristic policy that is originally due to Veeraraghavan and Scheller-Wolf (2008). They show numerically that their so-called dual-index policy performs quite well compared to the optimal policy. The dual index policy tracks two indices: One that contains all orders that are still in-transit and one that contains only those in-transit orders that are due to arrive within the lead time of the fast mode. Based on these outstanding orders, the policy dynamically ships orders with both modes to keep these indices at certain target levels. In line with standard inventory management nomenclature, we also refer to these target levels as base-stock levels. More specifically, the policy operates as follows. At the beginning of every period t after orders $Q_{j,f}^{t-l_{j,f}}$ and $Q_{j,s}^{t-l_{j,s}}$ have arrived, we review the *fast* inventory position, which includes all in-transit orders – i.e. shipped with both the slow and the fast transport mode – that will arrive within the lead time of the fast transport mode:

$$IP_{j,f}^t = I_j^t + \sum_{k=t-l_{j,f}+1}^{t-1} Q_{j,f}^k + \sum_{k=t-l_{j,s}+1}^{t-l_j} Q_{j,s}^k.$$

Then, if necessary, we place order $Q_{j,f}^t$ with the fast transport mode to raise the fast inventory position to its target level $S_{j,f}$. That is, the amount of product j shipped in

period t with the fast transport mode equals:

$$Q_{j,f}^t = (S_{j,f} - IP_{j,f}^t)^+.$$

After placing the fast shipment order, we inspect the *slow* inventory position, which includes the fast order just placed

$$IP_{j,s}^t = I_j^t + \sum_{k=t-l_{j,f}+1}^{t} Q_{j,f}^k + \sum_{k=t-l_{j,s}+1}^{t-1} Q_{j,s}^k,$$

and ship an order with the slow transport mode such that this inventory position is raised to its target level $S_{j,s}$, with $S_{j,s} \geq S_{j,f}$ since the fast inventory position is contained in the slow inventory position. Thus the amount of product j shipped with the slow transport mode in period t equals:

$$Q_{j,s}^t = S_{j,s} - IP_{j,s}^t.$$

Note that, contrary to the fast inventory position, the slow inventory position can never be larger than its base stock level $S_{j,s}$. After shipping both orders, demand D_j^t is satisfied or backlogged, depending on whether there is sufficient inventory available or not. The period then concludes with charging holding or backlog costs.

The order of events in a period t for each product j is thus as follows:

1. Orders $Q_{j,f}^{t-l_{j,f}}$ and $Q_{j,s}^{t-l_{j,s}}$ arrive with the fast and slow transport mode, respectively, and are added to the on-hand inventory I_j^t.
2. Review the fast inventory position and ship order $Q_{j,f}^t$ with the fast transport mode at unit cost $c_{j,f}$.
3. Review the slow inventory position and ship order $Q_{j,s}^t$ with the slow transport mode at unit cost $c_{j,s}$.
4. Demand D_j^t occurs and is satisfied from on-hand inventory if possible, and otherwise backlogged.
5. Incur a cost h_j for any unit in on-hand inventory $(I_j^t - D_j^t)^+$ and a cost p_j for any unit in backlog $(D_j^t - I_j^t)^+$.

Observe that under a dual-index policy, slow orders entering the information horizon of the fast transport mode may cause the fast inventory position to exceed its target level. The amount by which the fast inventory position exceeds its target level is referred to as the overshoot. The fast inventory position in period t after placing orders with both modes thus equals $S_{j,f} + O_j^t$, where O_j^t denotes the overshoot for product j in period t:

$$O_j^t = IP_{j,f}^t + Q_{j,f}^t - S_{j,f} = (IP_{j,f}^t - S_{j,f})^+.$$

Later, in Section 2.4.2, we shall see that computing the steady state distribution of the overshoot is crucial for determining the performance of a given control policy for a single product.

We furthermore define $\Delta_j = S_{j,s} - S_{j,f}$, $j \in J$, so that the control policy for a product can be specified in terms of its base-stock levels $S_{j,s}$ and $S_{j,f}$ or in terms of its base-stock level for the fast transport mode $S_{j,f}$ and the difference Δ_j. We mostly use the latter specification in our subsequent analysis. A control policy $(\mathbf{S}_f, \Delta)$ for the entire assortment of products consists of the vectors $\mathbf{S}_f = (S_{1,f}, S_{2,f}, \ldots, S_{|J|,f})$ and $\mathbf{\Delta} = (\Delta_1, \Delta_2, \ldots, \Delta_{|J|})$.

In what follows, for all sequences of random variables X^t, we define their stationary expectation as $\mathbb{E}[X] = \lim_{T\to\infty}(1/T)\sum_{t=0}^{T} X^t$ and their distribution as $\mathbb{P}(X \leq x) = \lim_{T\to\infty}(1/T)\sum_{t=0}^{T} \mathbb{1}\{X^t \leq x\}$, where $\mathbb{1}\{A\}$ is the indicator function for the event A. Whenever we drop the period index t we refer to the generic stationary random variable X with expectation and distribution as defined above.

2.3.3 Decision problem

For a given control policy $(\mathbf{S}_f, \mathbf{\Delta})$, we define the total long-run average holding, backlog, and ordering costs per period for the entire assortment of products as

$$\begin{aligned} C(\mathbf{S}_f, \mathbf{\Delta}) &= \sum_{j\in J} C_j(S_{j,f}, \Delta_j) \\ &= \sum_{j\in J} \left(h_j \mathbb{E}[(I_j - D_j)^+] + p_j \mathbb{E}[(D_j - I_j)^+] + \sum_{m\in M} c_{j,m}\mathbb{E}[Q_{j,m}] \right), \end{aligned} \tag{2.2}$$

and the total emissions as

$$E(\mathbf{S}_f, \mathbf{\Delta}) = \sum_{j\in J} E_j(S_{j,f}, \Delta_j) = \sum_{j\in J}\sum_{m\in M} e_{j,m}\mathbb{E}[Q_{j,m}],$$

where it is understood that the expectation operators are conditional on the control policy $(\mathbf{S}_f, \mathbf{\Delta})$. Veeraraghavan and Scheller-Wolf (2008) show that $C(\mathbf{S}_f, \mathbf{\Delta})$ is well-defined for any control policy $(\mathbf{S}_f, \mathbf{\Delta})$ as long as $\mathbb{E}[D_j] < \infty$ for all products $j \in J$.

The objective of our decision problem is to minimize the total long-run average costs while keeping the total emissions below a target level $\mathcal{E}^{max}$. Combining the above-mentioned leads to the following mathematical formulation of our decision problem which we refer to as problem (P):

$$\begin{aligned} (P) \qquad & \min && C(\mathbf{S}_f, \mathbf{\Delta}) \\ & \text{subject to} && E(\mathbf{S}_f, \mathbf{\Delta}) \leq \mathcal{E}^{max}, \\ & && \mathbf{S}_f \in \mathbb{R}^{|J|}, \quad \mathbf{\Delta} \in \mathbb{R}_0^{|J|}. \end{aligned}$$

Let $(\mathbf{S}_f^*, \mathbf{\Delta}^*)$ denote an optimal solution to problem (P) and let C_P be the corresponding optimal cost. Note that Problem (P) is a non-linear non-convex knapsack problem

where more than one copy of each item can be selected. It is well-known that even the simplest types of such knapsack problems are $\mathcal{NP}$-hard (e.g. Kellerer et al., 2004). Since our knapsack is more complex, we conclude that Problem (P) also falls in that same complexity class; it is hence likely that also for our problem no exact polynomial time solution algorithm exists.

We remark that Problem (P) enables companies to reduce carbon emissions from their inbound logistics by imposing a constraint on the maximally allowable carbon emissions. This is particularly useful for companies that seek to reduce carbon emissions proactively. However, companies may also take a reactive position and make supply mode decisions based only on inventory and transport costs. This cost will also include a carbon emission price component in regions where emissions are subject to carbon pricing mechanisms such as carbon crediting or taxing. We now briefly show that our model and analysis also apply to that setting. To that end, let c^e denote the price of one unit of carbon emissions. This price can also depend on the transport mode m and/or product j, but for ease of exposition we omit those dependencies. The long-run average costs per period in (2.2) can now be redefined as follows:

$$\tilde{C}(\mathbf{S}_f, \mathbf{\Delta}) = \sum_{j \in J} \left(h_j \mathbb{E}[(I_j - D_j)^+] + p_j \mathbb{E}[(D_j - I_j)^+] + \sum_{m \in M} (c_{j,m} + c^e e_{j,m}) \mathbb{E}[Q_{j,m}] \right).$$

The decision problem is now to minimize the long-run average costs per period:

$$\begin{array}{lll} (\tilde{P}) & \min & \tilde{C}(\mathbf{S}_f, \mathbf{\Delta}) \\ & \text{subject to} & \mathbf{S}_f \in \mathbb{R}^{|J|}, \quad \mathbf{\Delta} \in \mathbb{R}_0^{|J|}. \end{array}$$

Problem $(\tilde{P})$ is less complex than the original decision problem since it does not involve a constraint that links the individual products. As such, Problem $(\tilde{P})$ can be decomposed in $|J|$ product specific problems, each of which can be solved individually. The column generation sub-problem that we will discuss in Section 2.4.2 has a similar structure as the product specific problems of $(\tilde{P})$, and the solution method we discuss there thus also applies to $(\tilde{P})$.

2.4 Analysis

This section focuses on finding the optimal control policy for Problem (P). Our approach relies on the technique of column generation – also named Dantzig-Wolfe decomposition after its pioneers (Dantzig and Wolfe, 1960). This technique enables a natural decomposition of the original multi-product decision problem into smaller single-product problems that have more structure. We refer the interested reader to Lübbecke and Desrosiers (2005) for a comprehensive survey on column generation. Below, we first explain how we apply column generation to Problem (P), and we then describe a simulation-based optimization method for solving the sub-problem of this column generation procedure.

2.4.1 The column generation procedure

We first reformulate decision problem (P) as an integer linear program in which each binary decision variable corresponds to a certain combination of values for the decision variables of our original decision problem. We subsequently relax the integrality constraint and we call this problem the master problem (MP). Formally, let K_j be the set of all possible dual index policies for product $j \in J$. Each policy $k \in K_j$ is determined by its policy parameters $S^k_{j,f}$ and Δ^k_j. Let $x^k_j \in \{0,1\}$ denote the decision variable that indicates whether policy $k \in K_j$ is selected ($x^k_j = 1$) for product $j \in J$ or not ($x^k_j = 0$). By relaxing the integrality constraint on this binary decision variable, we arrive at the mathematical formulation of the master problem (MP):

$$
\begin{aligned}
(MP) \qquad \min \quad & \sum_{j\in J}\sum_{k\in K_j} C_j(S^k_{j,f},\Delta^k_j)x^k_j && && (2.3)\\
\text{subject to} \quad & \sum_{j\in J}\sum_{k\in K_j} E_j(S^k_{j,f},\Delta^k_j)x^k_j \le \mathcal{E}^{max}, && && (2.4)\\
& \sum_{k\in K_j} x^k_j = 1, && \forall j \in J && (2.5)\\
& x^k_j \ge 0, && \forall j \in J, \forall k \in K_j.
\end{aligned}
$$

Let C^{LB}_P denote the optimal cost for master problem (MP). Due to the linear relaxation of the integrality constraint on x^k_j, an optimal cost C^{LB}_P is also a lower bound on the optimal cost for Problem (P), C_P.

Due to its large number of decision variables, master problem (MP) is solved using column generation. To this end, we first restrict master problem (MP) to a small subset $\tilde{K}_j \subseteq K_j$ of trivial policies per product $j \in J$ (i.e. columns) that are feasible for Problem (P) (and thus also for Problem (MP)). Such a trivial policy prescribes, for instance, to ship orders exclusively with the least polluting mode. This restricted problem is referred to as the restricted master problem (RMP). We then solve (RMP) to optimality, and we are interested in new policies $k \in K_j \setminus \tilde{K}_j$, $j \in J$, that will improve the objective value of (RMP) if they are added to $\tilde{K}_j$. Such policies $k \in K_j \setminus \tilde{K}_j$ are identified through solving a column generation sub-problem for each product $j \in J$. The objective function of such a sub-problem is the reduced cost as a function of the policy with respect to the current dual variables obtained through solving (RMP) to optimality. If a policy $k \in K_j \setminus \tilde{K}_j$ has a negative reduced cost, then adding that policy as a column to $\tilde{K}_j$ in (RMP) will reduce the objective value of (RMP). More formally, the column generation sub-problem for product $j \in J$ has the following form:

$$
\begin{aligned}
(SP(j)) \quad \min \; & h_j\mathbb{E}[(I_j - D_j)^+] + p_j\mathbb{E}[(D_j - I_j)^+] + \sum_{m\in M}(c_{j,m} - \eta e_{j,m})\mathbb{E}[Q_{j,m}] - \upsilon_j,\\
\text{subject to} \quad & S_{j,f} \in \mathbb{R}, \quad \Delta_j \in \mathbb{R}_0,
\end{aligned}
$$

where η denotes the dual variable of (RMP) that corresponds with emission constraint (2.4) and υ_j denotes the dual variable of (RMP) that corresponds with constraint (2.5)

that assures that for each product $j \in J$ a convex combination of policies is chosen. Note that these dual variables can also be interpreted as the Lagrange multipliers of relaxing the corresponding constraints (Lübbecke and Desrosiers, 2005). If for product $j \in J$ a feasible solution for ($SP(j)$) exists with a negative objective value (i.e. a negative reduced cost), then this policy is added to $\tilde{K}_j$ since the objective value of (RMP) can be improved when solved with the enlarged set $\tilde{K}_j$.

We continue with iterating between optimizing (RMP) and finding new policies through solving ($SP(j)$), $j \in J$, until no product for which there is a policy with a negative reduced cost exists. An optimal solution for (RMP) is then also an optimal solution for (MP). If this optimal solution contains integer values only, then it is also an optimal solution for (P). If this is not the case, then we solve (RMP) one last time as an integer program to find an integer solution for (RMP), which is then also a feasible solution for (P). Recent inventory literature has shown that solving the restricted master program as an integer program to arrive at an integer solution leads to good performance in terms of optimality gaps (e.g., Drent and Arts, 2021; Haubitz and Thonemann, 2021), and often outperforms alternative approaches such as local searches or rounding procedures (Alvarez et al., 2013). The corresponding cost of the resulting feasible solution is also an upper bound, denoted C^P_{UB}, for C^P.

In the next section, we provide a simulation-based optimization procedure to solve the column generation sub-problem ($SP(j)$).

2.4.2 Solving the column generation sub-problem

The column generation sub-problem ($SP(j)$) has the same structure as the problem studied by Veeraraghavan and Scheller-Wolf (2008). We follow their simulation-based optimization procedure to solve ($SP(j)$). This procedure is grounded in the following separability result that allows us to find the optimal $S_{j,f}$ for a given Δ_j as the solution to a special Newsvendor problem.

Lemma 2.1. *(Veeraraghavan and Scheller-Wolf, 2008, Proposition 4.1) The distributions of the overshoot O_j, the fast transport mode shipment size $Q_{j,f}$, and the slow transport mode shipment size $Q_{j,s}$ are functions of Δ_j only, independent of $S_{j,f}$.*

Let $O_j^t(\Delta_j)$ denote the overshoot of product j in period t for a given Δ_j. Recall that the fast inventory position of product j in period t after shipping equals $S_{j,f} + O_j^t(\Delta_j)$. Consequently, for the net inventory level of product j in each period t, we can also write

$$I_j^t = S_{j,f} - \left(\sum_{k=t-l_{j,f}}^{t-1} D_j^k - O_j^{t-l_{j,f}}(\Delta_j) \right). \tag{2.6}$$

By plugging (2.6) in the objective function of ($SP(j)$), and using Lemma 2.1 as well as the fact that in each period the overshoot is independent of the demand and $S_{j,f}$, we

readily recognize that for given Δ_j the objective function is convex in $S_{j,f}$. This implies the following result.

Lemma 2.2. *(Veeraraghavan and Scheller-Wolf, 2008, Theorem 4.1) The optimal base-stock level $S^*_{j,f}$ for a given Δ_j, denoted $S^*_{j,f}(\Delta_j)$, equals*

$$S^*_{j,f}(\Delta_j) = \inf\left\{S_{j,f} \in \mathbb{R} : \mathbb{P}\left(\sum_{i=1}^{l_{j,f}+1} D_i - O_j(\Delta_j) \leq S_{j,f}\right) \geq \frac{p_j}{p_j + h_j}\right\}.$$

It now remains to calculate the objective value of $SUB(j)$ for given Δ_j and corresponding $S^*_{j,f}(\Delta_j)$. Observe that in each period immediately after shipping orders with both modes, the slow inventory position equals the fast inventory position plus the overshoot and all remaining outstanding slow orders. Since these inventory positions are equal to their respective base-stock levels following order placement, we have for each product j in each period t:

$$S_{j,s} = S_{j,f} + O^t_j + \sum_{k=0}^{l_j-1} Q^{t-k}_{j,s}. \tag{2.7}$$

From (2.7) it follows that $\mathbb{E}[Q_{j,s}] = (\Delta_j - \mathbb{E}[O_j])/l_j$. Since under backlogging the sum of both orders must on average be equal to the period demand, we finally find $\mathbb{E}[Q_{j,f}] = \mathbb{E}[D_j] - \mathbb{E}[Q_{j,s}]$.

To solve $(SP(j))$, $j \in J$, to optimality, it thus suffices to perform a one-dimensional search over Δ_j. For each Δ_j, we compute the stationary distribution of the overshoot. With this stationary distribution we readily find the optimal base-stock level $S^*_{j,f}(\Delta_j)$ through Lemma 2.2 and the total reduced cost through the identities following Equation (2.7). As there is in general no closed-form expression for the stationary distribution of the overshoot, we follow Veeraraghavan and Scheller-Wolf (2008) and rely on simulation to compute this distribution.

Note that our optimization model and analysis readily extends to settings where the slow transport modes of all (or some) products are operated according to any other rule that depends only on the current overshoot as well as all in-transit orders that are not yet included in the fast inventory position. That is, any other rule that depends only on the information state $(O^t_j, Q^{t-1}_{j,s}, Q^{t-2}_{j,s}, \ldots, Q^{t-l_j+1}_{j,s})$, $j \in J$, $t \in \mathbb{N}_0$. Most well-performing control policies satisfy this condition, e.g., the Capped Dual-Index policy (Sun and Van Mieghem, 2019), the Tailored Base-Surge policy (Allon and Van Mieghem, 2010), and the Projected Expedited Inventory Position policy (Drent and Arts, 2022). Sheopuri et al. (2010) show that for such control policies, the stationary distribution of the overshoot is a function of only the parameter(s) for operating the slow transport mode, and that consequently a Newsvendor result similar to Lemma 2.2 holds for all such policies.

2.5 Computational experiment

Our test-bed has three different types of assortments of products, each representing a different type of industry. The first assortment type consists solely of products for which emissions from the slowest transport mode are less than the emissions from the faster transport mode. This assortment is inspired by apparel goods that are delivered from Vietnam to Europe by sea transport as the slowest mode and by air transport as the fastest mode. In this example, the fast supply mode has a higher carbon footprint. The opposite holds for the second assortment type. Here we are inspired by industrial goods that are delivered from China to Europe by sea transport as the slowest mode and from Germany by truck as the fastest mode. In this case, the slow supply mode is associated with more emissions as goods are transported over a longer distance. The third assortment has products of both types.

We perform a parametric computational experiment. The base case is set up as follows. We consider 100 products for each assortment type, i.e. $|J| = 100$. The input parameters in the base case are identical for all three assortments, except for the carbon emissions from transportation. For each product $j \in J$, the period demand D_j follows a negative binomial distribution. To create heterogeneous assortments, the parameters of this negative binomial distribution are randomly drawn from two separate distributions for each product j. The mean $\mu_{D_j} := \mathbb{E}[D_j]$ is randomly drawn from a gamma distribution with mean 100 and coefficient of variation of 0.5. The coefficient of variation $CV_{D_j} := \sqrt{\mathbf{Var}[D_j]}/\mu_{D_j}$ is randomly drawn from a shifted beta distribution with mean 0.9, standard deviation of 0.25, and shifted to the right by 0.3. Since low demand products typically have higher holding cost, the holding cost h_j is negatively correlated with the mean demand μ_{D_j} of each product j through a Gaussian copula with a fitting covariance matrix. In particular, h_j is drawn from a gamma distribution with mean 1 and coefficient of variation equal to CV_{D_j}, with a Pearson correlation coefficient of -0.5. Details regarding our approach to generate correlated random numbers are relegated to Appendix 2.B.

We set $l_{j,f}$ and $c_{j,s}$ to 0 for all products, and focus on $l_{j,s}$ and $c_{j,f}$, which now coincide with the lead time difference and the cost premium of product $j \in J$, respectively. We set $l_{j,f} = 3$ for all products. The back-order penalty cost p_j for product j is a function of its holding cost h_j. The ratio between the holding and the penalty cost is an important determinant of the service level in an inventory system. Therefore we set $p_j = \psi_j \chi_j^p h_j$ where ψ_j is a parameter we use to control the ratio between p_j and h, and χ_j^p is a random perturbation. That is, χ_j^p has a shifted beta distribution with mean 0.98, standard deviation 0.1, and shifted 0.02 to the right.

The cost premium $c_{j,f}$ of the fast transport mode of product j equals $\chi_j^c p_j l_{j,f}$, where χ_j^c has a beta distribution with mean 0.25 and standard deviation 0.1. Table 2.2 provides a summary of how we randomly generated the products of the base case; we shortly

explain how we randomly generated the emission units for these products for all the three assortment types. In Table 2.2, $NB(\mu, cv)$ denotes a negative binomial random variable with mean μ and coefficient of variation cv, $\Gamma(\mu, cv)$ denotes a gamma random variable with mean μ and coefficient of variation cv, and $B(\mu, \sigma, s)$ denotes a beta random variable with mean μ and standard deviation σ that is shifted to the right by s; if we drop s then this beta random variable is not shifted, i.e. $B(\mu, \sigma, s) =_d B(\mu, \sigma) + s$ where $=_d$ denotes equality in distribution.

Table 2.2: Generating the base case input parameters.

Input parameter	Generation
$\|J\|$	100
D_j	$NB(\mu_{D_j}, \mathrm{CV}_{D_j})$, with $\mu_{D_j} \sim \Gamma(100, 0.5)$ and $\mathrm{CV}_{D_j} \sim B(0.9, 0.25, 0.3)$
$l_{j,f}$	0
$l_{j,s}$	3
h_j	$\Gamma(1, 0.5)$, with $\rho_{\mu_{D_j}, h_j} = -0.5$
p_j	$\psi_p \chi_j^p h_j$, with $\psi_p = 9$ and $\chi_j^p \sim B(0.98, 0.1, 0.02)$
$c_{j,s}$	0
$c_{j,f}$	$\chi_j^c p_j l_{j,f}$, with $\chi_j^c \sim B(0.25, 0.1)$

As explained in Section 2.3, we rely on the NTM framework (NTM, 2015) to set emissions based on the structure of equation (2.1). We apply the NTM framework to data from the UN Comtrade Database (Database, 2020) to obtain sample emission units; details regarding this methodology are relegated to appendix 2.A. We then apply maximum likelihood estimation on these sample unit emissions to obtain three distinct sets of two distribution functions; two for each assortment type of products. The number of carbon emission units $e_{j,m}$ for product $j \in J$ with transport mode $m \in M$ are then randomly drawn from these distributions. These distributions are presented in Table 2.3. In this table, $LN(\mu, \sigma)$ denotes a random variable whose logarithm is normally distributed with mean μ and standard deviation σ, and $WB(\lambda, k)$ denotes a Weibull random variable with scale λ and shape k.

Table 2.3: Generating the emission units for the base case.

Assortment type	Emission parameter	Generation
1	$e_{j,s}$	$\Gamma(0.35, 0.21)$
	$e_{j,f}$	$e_{j,s} + LN(1.52, 0.21)$
2	$e_{j,f}$	$\Gamma(0.19, 1.27)$
	$e_{j,s}$	$e_{j,f} + \Gamma(2.19, 1.27)$
3	$e_{j,f}$	$WB(0.87, 0.77)$
	$e_{j,s}$	$\Gamma(3.31, 1.34)$

The total allowable carbon emissions from transportation $\mathcal{E}^{\max}$ is set as a percentage of the total *reducible* carbon emissions. For each instance of the test-bed, the reducible carbon emissions is defined as the difference between the total amount of carbon emissions of the control policy that is optimal for Problem (P) absent of the emission constraint and the total amount of carbon emissions of the control policy that leads to the lowest possible total carbon emissions. The latter implies that each product is only shipped

with its least polluting transport mode. Under the dual-index policy, setting Δ_j to zero implies that all orders for product $j \in J$ are shipped with its fastest transport mode. Alternatively, letting $S_{j,f}$ go to $-\infty$ implies that all orders for product j are shipped with its slowest transport mode (Veeraraghavan and Scheller-Wolf, 2008).

To evaluate the effectiveness of the column generation procedure in solving Problem (P), we compute for each instance of the test-bed the relative difference between the total average cost under a feasible solution and the corresponding lower bound. That is,

$$\%GAP = 100 \cdot \frac{C_P^{UB} - C_P^{LB}}{C_P^{LB}},$$

where C_P^{LB} and C_P^{UB} are obtained using the methods described in Section 2.4. In what follows, we also refer to this feasible solution as dynamic mode selection (DMS). Hence the long run average cost of dynamic mode selection equals C_P^{UB}.

To quantify the benefit of using two transport modes dynamically rather than relying statically on one transport mode, we define for each instance of the test-bed a benchmark instance in which we can only select one transport mode for each product. As described above, the dual-index policy can mimic static mode selection (SMS). Hence, to find a feasible static mode selection solution to this benchmark instance of Problem (P), we apply our column generation procedure of Section 2.4 in which we restrict the solution space for each product $j \in J$ such that all orders are shipped with either the fastest or the slowest transport mode. The mathematical formulation for the static mode selection approach as well as the benchmark approach described in the next paragraph can be found in Appendix 2.C. (Note that in the Master Problem (MP) for this approach the set of possible policies K_j for each product $i \in J$ contains only two single transport mode policies.) The long run average cost of this solution is denoted C_P^{SMS}. To quantify the value of dynamic mode selection, we compare for each instance of the test-bed the long run average cost of static mode selection with the long run average cost of dynamic mode selection. That is,

$$\%SMS = 100 \cdot \frac{C_P^{SMS} - C_P^{UB}}{C_P^{UB}},$$

where $\%SMS$ indicates the relative increase in the long run average cost when the company chooses to rely on only one transport mode for each product in meeting an assortment wide constraint on total emissions.

To quantify the portfolio effect, we define for each instance of the test-bed an additional benchmark in which we enforce an emission constraint for each product $j \in J$. This emission constraint is set as a percentage of the total reducible emissions of the individual product rather than of the entire assortment of products. We compute a feasible solution for this benchmark instance with the column generation procedure of Section 2.4, which can readily be modified so that it can be applied to settings where we have additional emission constraints. (An alternate solution procedure is to use a simple search procedure per product.) This approach, which we refer to as blanket mode selection (BMS),

seems a plausible approach for most practitioners. Indeed, they can consider each product individually absent of the complicating linking emission constraint and yet they are guaranteed that the total emissions of the entire assortment is kept below the target level. The long run average cost of the blanket mode selection approach is denoted C_P^{BMS}. To quantify the portfolio effect, we compare for each instance of the test-bed the long run average cost of blanket mode selection with the long run average cost of dynamic mode selection. That is,

$$\%BMS = 100 \cdot \frac{C_P^{BMS} - C_P^{UB}}{C_P^{UB}},$$

where $\%BMS$ indicates the relative increase in the long run average cost when the company enforces an emission constraint on each individual product rather than one single constraint for the emissions from the entire assortment.

In solving the column generation sub-problem for each product, we simulated 10 samples of 9500 time periods following a 5000 time periods warm-up . The width of the 95 percent confidence interval of the long run average cost per period for each product was no larger than 3 percent of its corresponding point estimate for each instance of the column generation sub problem that we solved. The average computational time of our column generation procedure was 23 minutes, 5 minutes for blanket dynamic mode selection, and less than 5 sec for static mode selection.

2.5.1 Results for the base case

Figure 2.1 presents the normalized optimal average costs of each approach for each assortment group under emission targets that range from from 0 to 100 percent of the total reducible carbon emissions. Observe that for each assortment group, all approaches have the same performance when the emission target is set at 100 percent of the total reducible emissions. In this case, all approaches solely utilize the least polluting transport mode. Alternatively, when we impose no target on the emissions from transportation, then both dual mode approaches perform equally well while the static mode selection approach seems to perform the poorest over all possible emission targets. Indeed, static mode selection is around 15 percent more expensive than both dual dual mode approaches for all assortment types when transportation emissions are not constrained. Based on Figure 2.1, we conclude that dynamic mode selection, as opposed to static and blanket mode selection, has great potential to efficiently curb carbon emissions from transportation at relatively little additional costs across all assortment types.

We explicitly compare our dynamic mode selection with the benchmark approaches in Figure 2.2, which presents the $\%SMS$ and $\%BMS$ percentages for each assortment group under emission targets that range from from 0 to 100 percent of the total reducible carbon emissions. The figure indicates that the performance of static mode selection over dynamic mode selection is consistent across all assortments. The relative increase

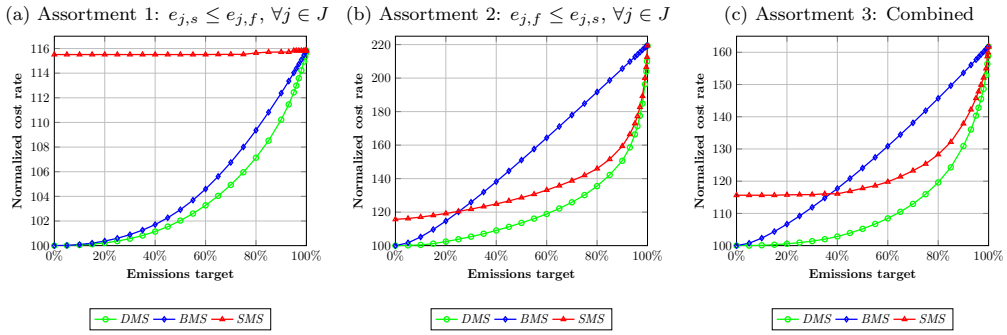

Figure 2.1: Optimal normalized absolute costs of each approach for different targets on the reducible emissions.

in its total cost over dynamic mode selection is the largest when there is no emission target, and gradually decreases as the emission constraint tightens. At moderate carbon emission targets, around 40 to 60 percent of the total reducible emissions, static mode selection still leads to increases in the total average cost per period of around 10 to 15 percent for all assortment types.

The performance of the blanket mode selection approach depends on the specific assortment type. Figure 2.2(a) illustrates that when the unit emissions from the fast transport mode are more than those from the slow transport mode, the performance of the blanket mode selection approach seems to be quite reasonable. This can be explained as in this setting, the cost of the fastest transport mode is larger than the cost of the slowest transport mode. The most polluting transport mode is thus also the most expensive transport mode. The portfolio effect is limited for this assortment type.

For the other two assortment types, however, the cheapest transport mode is not necessarily also the least polluting transport mode, and the portfolio effect is more prevalent. Figures 2.2(b) and 2.2(c) show that $\%BMS$ can be more than 35 and 20 percent in assortment type 2 and 3, respectively, under carbon emission reduction targets of 50 percent. The static mode selection approach, which also takes advantage of the portfolio effect, even outperforms the blanket mode selection approach for quite some emission targets.

Figure 2.3 shows the usage of the fast supply mode as a function of the carbon emission reduction target under dynamic mode selection as measured by

$$\%F = 100 \cdot \frac{1}{|J|} \sum_{j \in J} \frac{\mathbb{E}[Q_{j,f}]}{\mathbb{E}[Q_{j,f}] + \mathbb{E}[Q_{j,s}]}.$$

Dynamic mode selection is economically attractive regardless of emissions targets. Carbon reduction targets make dynamic mode selection even more attractive for assortment type 2, and mixed assortments, but not for assortments of type 1.

To recapitulate, the value of dynamically shipping products with two transport modes simultaneously rather than statically selecting one transport mode a priori is quite large.

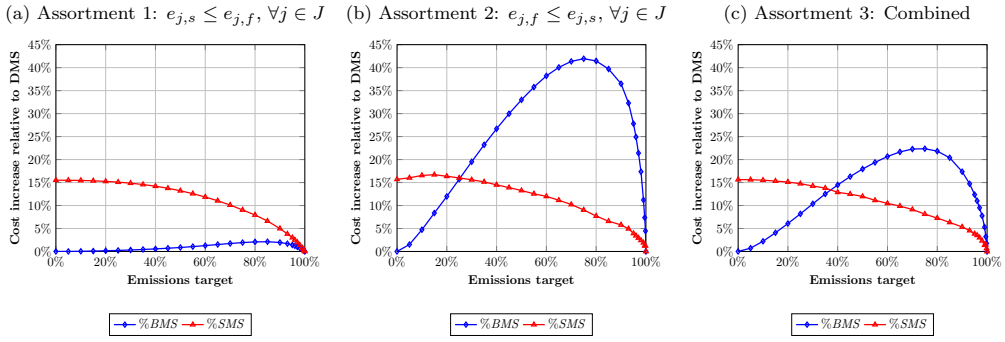

Figure 2.2: Relative surplus of the optimal cost of the alternative approaches (BMS and SMS) compared to the DMS approach for different targets on the reducible emissions.

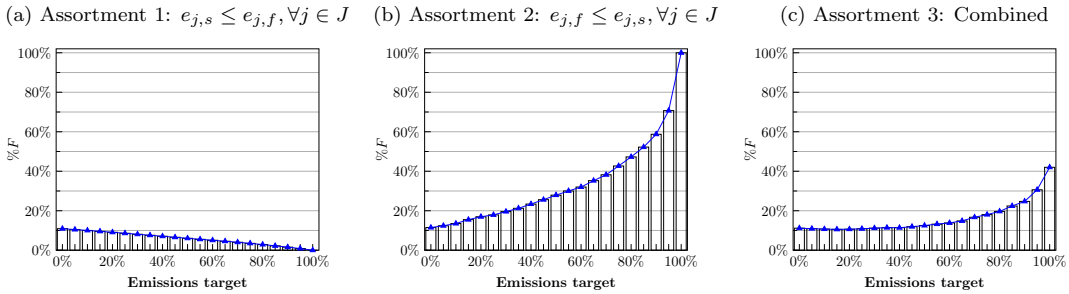

Figure 2.3: Percentage of products shipped with the fast mode for different carbon emission reduction targets.

Regardless of the assortment type, $\%SMS$ is in between 5 and 15 percent for emission reduction targets up to 90 percent. The portfolio effect depends on the specific assortment type. If the least polluting transport mode of each product is also its cheapest transport mode, then the fastest and most polluting transport modes are typically only relied upon in case of imminent backorders. This behavior remains in case of an assortment-wide emission target, and the portfolio effect is consequently rather limited. If the least polluting transport modes are not necessarily the cheapest transport modes, then there is substantial value to be reaped in optimizing the assortment of products under a single emission constraint rather than under separate emission constraints for each individual product. Indeed, $\%BMS$ can go up to 40 and 20 percent for assortment type 2 and 3, respectively.

Table 2.4 below presents the average relative slack in the emission constraints for each assortment over the different emission targets considered in the base case analysis. The table shows that due to the binary nature of the static mode selection approach, the total average emissions under this approach are often substantially lower than the target level. This leads to particularly poor performance for assortment type 1. We observed in our computational experiments that for this assortment type, the static mode selection approach selects the cheapest and thus least polluting transport mode for almost all products under each emission target.

Table 2.4: Average slack in emission constraints for the base case analysis.

	Assortment type		
Approach	1	2	3
DMS	0.00%	0.00%	0.00%
BMS	0.08%	0.61%	0.12%
SMS	12.16%	0.17%	0.97%

The average $\%GAP$ of the base case over all emission targets is less than 0.01 percent, indicating that the column generation procedure finds feasible solutions that are close to optimal. Such a low average $\%GAP$ occurs because there can be at most 1 product for which the optimal solution to Problem (MP) is fractional. Indeed, Problem (MP) has $|J|+1$ constraints and an optimal solution for this problem has the same number of basic variables. Constraint (2.5) assures that for each product $j \in J$ a convex combination of policies is chosen. As such, there is at least one basic variable for each product j. This implies that there is at most 1 product for which the optimal solution to Problem (MP) is fractional.

2.5.2 Determinant of emission reduction potential

The base case analysis of our DMS approach indicates that the emission reductions are not evenly distributed across the products. Products can be ordered by their contribution to emission reduction following the DMS optimization. In this manner we can construct the cumulative reduction in emissions as shown in the Lorenz curves (Figure 2.4) with the dashed line. Figure 2.4 shows that 20% of the products in assortments 1 through 3 account for 61.22%, 94.19%, and 91.88% of the emission reduction, respectively. This suggests that most of the emission reduction can be achieved by using dynamic mode selection for a limited subset of a given assortment. Although it is possible to determine the limited subset of products that account for most of the emission reduction after performing the DMS optimization, it would be convenient to know which products to focus on without having to solve a sophisticated optimization problem. Suppose we order products in increasing order of $\frac{|e_{j,f}-e_{j,s}|}{c_{j,f}-c_{j,s}}$, i.e., we sort products according to the how much emission can be saved by using the least polluting transport mode relative to the additional cost of the faster transport mode. Figure 2.4 shows the cumulative emission saving by products ordered this way with the solid line. Here we see that focusing on the 20% of products in assortments 1 through 3 for which $\frac{|e_{j,f}-e_{j,s}|}{c_{j,f}-c_{j,s}}$ is highest, already achieves 58.93%, 93.63%, and 82.64% of the potential emission reduction, respectively. Thus firms seeking to minimize the carbon footprint of their inbound logistics should focus their attention on products for which the difference in emission in different transport modes is large relative to the additional cost of fast transportation modes. That is emission differences between modes relative to cost difference between modes is the main determinant of emission reduction potential for a given assortment.

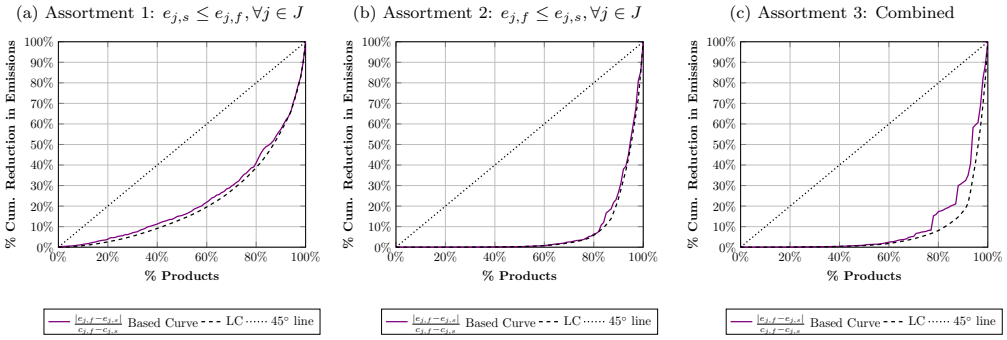

Figure 2.4: The cumulative emissions reduction share of items arranged based on different criteria.

2.5.3 Comparative statics

In this section, we study how changes in the input parameters with respect to the base case affect the performance of the blanket mode selection and the static mode selection approach. In what follows, we keep the emission target fixed at a 50 percent reduction of the reducible emissions, and we study the effects of changing a certain input parameter while generating the other input parameters as in the base case, i.e. as in Table 2.2. We also investigate the effect of scaling the emission differences between the most polluting and least polluting transport modes when the target on the total emissions is kept fixed. To achieve this, we first generate emission units as in the base case. We subsequently change the emission units of the most polluting transport mode through scaling $\mid e_{j,s} - e_{j,f} \mid$ by a constant δ_e while keeping the emission units from the least polluting mode fixed at its base level. The changes in the parameters we investigate are summarized in Table 2.5.

Table 2.5: Changes in the base case input parameters.

Parameter	Generation	Base case	Changes
$\|J\|$		100	$\|J\| \in \{40, 60, 80\}$
μ_{D_j}	$\Gamma(100, CV_{\mu_{D_j}})$	0.5	$CV_{\mu_{D_j}} \in \{0.3, 0.4, 0.6, 0.7\}$
CV_{D_j}	$B(0.9, 0.25, s_{D_j})$	0.3	$s_{D_j} \in \{0.2, 0.25, 0.35, 0.4\}$
$\rho_{\mu_{D_j}, h_j}$		-0.5	$\rho_{\mu_{D_j}, h_j} \in \{-0.3, -0.4, -0.6, -0.7\}$
ψ_p		9	$\psi_p \in \{3, 4, 5, 19, 99\}$
χ_j^c	$B(0.25, \sigma_c)$	0.1	$\sigma_c \in \{0.15, 0.2, 0.3, 0.35\}$
$l_{j,s}$		3	$l_{j,s} \in \{2, 4\}$
δ_e		1	$\delta_e \in \{0.8, 0.9, 1.1, 1.2\}$

Figure 2.5 shows the effect of changing the coefficient of variation of the gamma distribution from which we sample the mean demands per period for each product. The figure indicates that this effect is relatively limited. With respect to the base case, both $\%BMS$ and $\%SMS$ only change up to 1 percent point for all three assortment types. We can draw a similar conclusion for the effect of changing the Pearson correlation coefficient between the holding cost and the mean demand per period of each product. Figure 2.6 shows that both $\%BMS$ and $\%SMS$ change at most 1 percent point with respect to the

base case for all three assortment types.

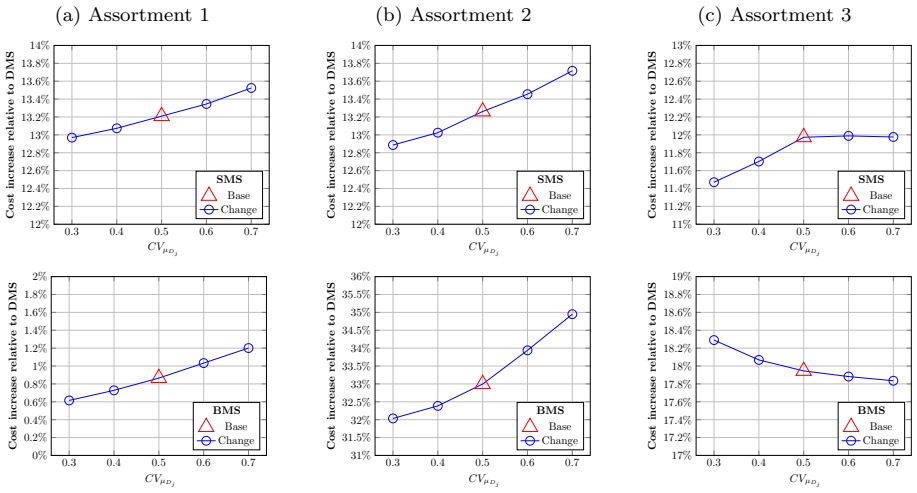

Figure 2.5: Effect of changing $CV_{\mu_{D_j}}$ while keeping the rest of the parameters as in the base case.

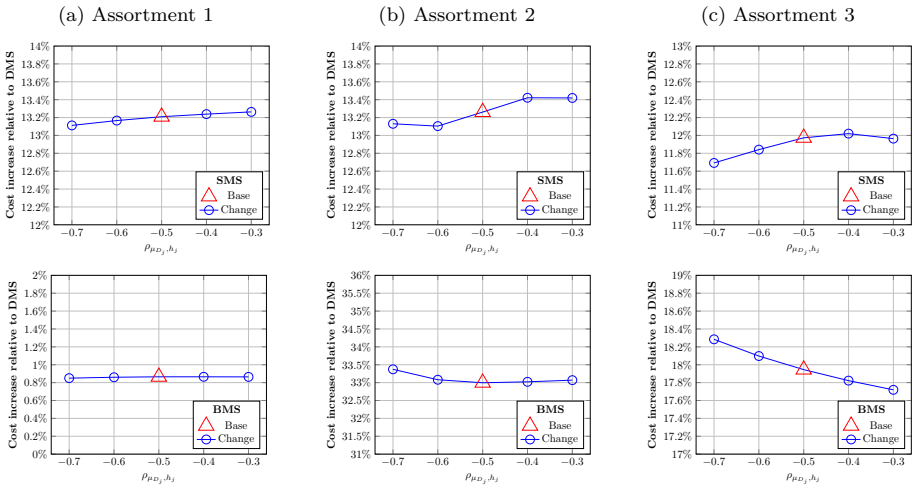

Figure 2.6: Effect of changing $\rho_{\mu_{D_j},h_j}$ while keeping the rest of the parameters as in the base case.

Alterations in the shift parameter of the beta distribution from which we sample the coefficient of variation of the demand per period for each product has a relatively moderate effect on both $\%BMS$ and $\%SMS$. Figure 2.7 illustrates that for all assortment types, the $\%SMS$ tends to increase in the variability of the demand while the $\%BMS$ decreases. This indicates that the flexibility to dynamically ship products with two transport modes has particular merit in highly variable demand settings.

Figure 2.8 indicates that for all assortment types, an increase (decrease) in the lead time difference between the fastest and the slowest transport modes of each product leads to an increase (decrease) in both $\%SMS$ and $\%BMS$. The blanket mode selection approach seems to be more susceptible to changes in the lead time difference than the

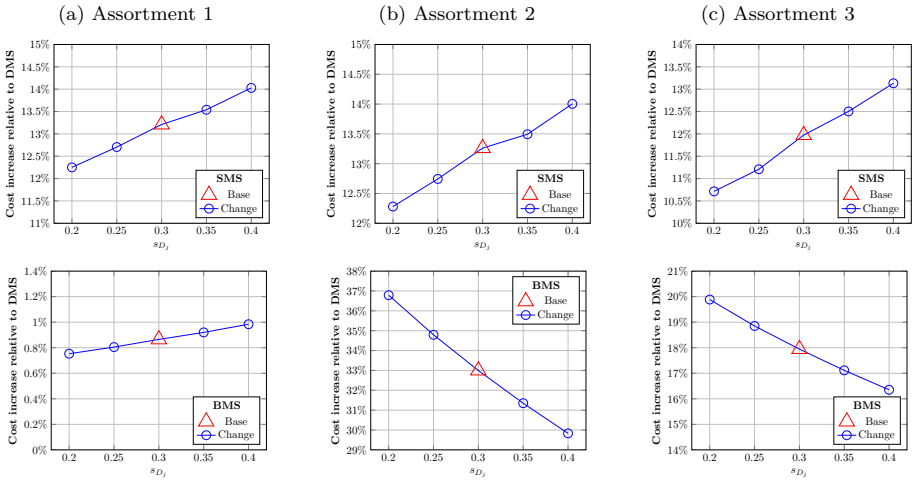

Figure 2.7: Effect of changing s_{D_j} while keeping the rest of the parameters as in the base case.

static mode selection approach. For assortment type 2, for instance, an increase in the lead time difference to 4 leads to an increase in $\%BMS$ of 7 percent points with respect to the base case. By contrast, $\%SMS$ increases only slightly by 0.5 percent points. This can be attributed to the fact that the blanket mode selection approach imposes constraints on the emissions of each individual product while the static mode selection approach imposes a single constraint on the entire assortment of products.

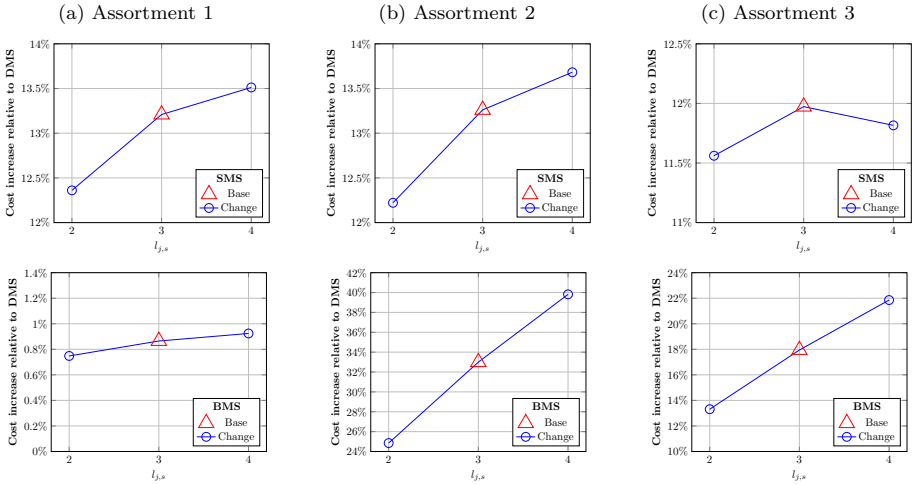

Figure 2.8: Effect of changing $l_{j,s}$ while keeping the rest of the parameters as in the base case.

Figure 2.9 illustrates the effect of changing the critical ratio for all products through varying ψ_p. We conclude that this effect is quite large. For assortment type 1, for instance, $\%SMS$ varies from 20 percent to 2 percent. While $\%SMS$ seems to decrease in the critical ratio for all products, $\%BMS$ tends to increase. For assortment type 2 and 3, for instance, $\%BMS$ increases from 15 percent to over 80 percent. These effects can be explained by the fact that as the critical ratios of all products approach 1, our

dynamic mode selection approach will mimic the static mode selection approach in which the assortment wide emission constraint is met by relying on less polluting transport mode for product for which this is relatively cheap to do so. The blanket mode selection approach, however, must meet emission targets for each product individually which leads to poor performance if we increase the critical ratios for all products.

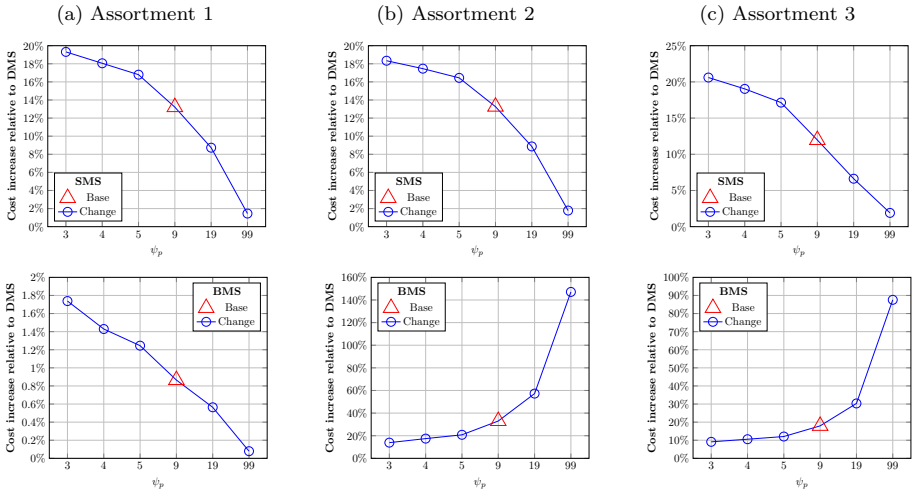

Figure 2.9: Effect of changing ψ_p while keeping the rest of the parameters as in the base case.

Figure 2.10 indicates that $\%SMS$ decreases in the cost of the fast transport mode. This can again be explained by the fact that our dynamic mode selection approach will also rely more on the cheaper transport mode as the cost premium for the fast transport mode increases, and that consequently the gap with the static mode selection approach decreases. By contrast, $\%BMS$ increases in the cost of the fast transport mode for assortment type 2 and 3. This can be attributed to the fact that the blanket mode selection approach, contrary to the other two approaches, imposes itemized emission constraints and relying on the most expensive but least polluting transport mode is therefore inevitable. Note that this is not true for assortment type 1 because there the fast, expensive transport mode is also the most polluting transport mode. We can draw similar conclusions for the effects scaling the emission units of the most polluting transport mode, see Figure 2.11. Finally, Figure 2.12 indicates that the impact of the assortment size on both $\%SMS$ and $\%BMS$ is relatively limited. With respect to the base case, both $\%BMS$ and $\%SMS$ only change up to 2 percent point for all three assortment types.

2.6 Concluding remarks

As carbon emissions from the transportation sector are projected to increase over the next decades, it is important for companies to rethink their supply chain strategies and ex-

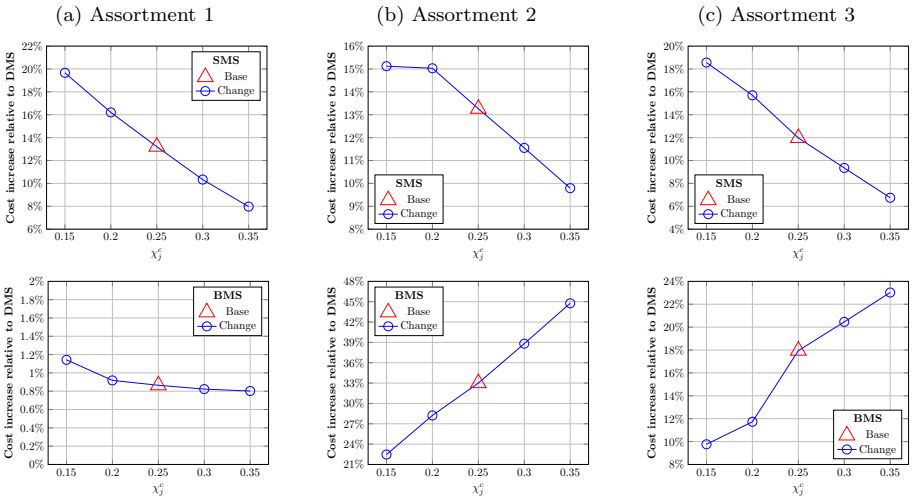

Figure 2.10: Effect of changing χ_j^c while keeping the rest of the parameters as in the base case.

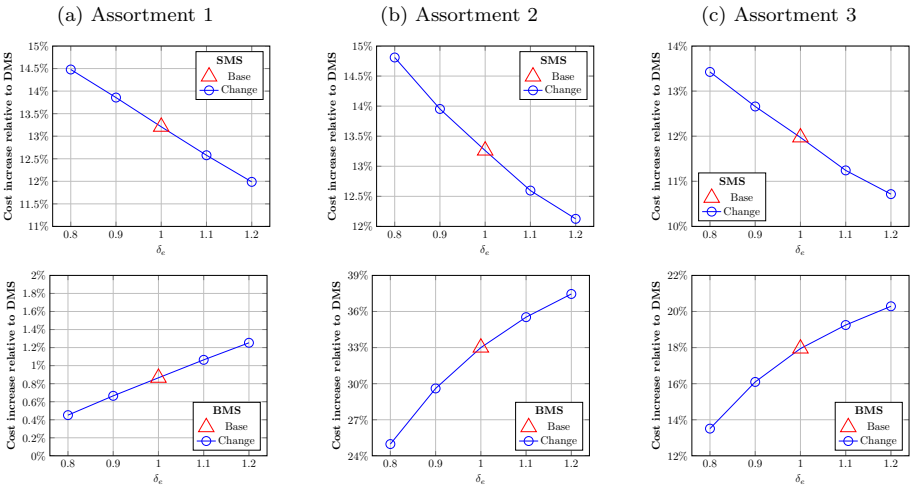

Figure 2.11: Effect of changing δ_e while keeping the rest of the parameters as in the base case.

plicitly incorporate carbon emissions into their decision making. In this chapter, we have studied the inbound transport and inventory management decision making for a company that sells an assortment of products. The company wishes to minimize inventory costs while keeping the total emissions from the inbound transport of the entire assortment below a certain target level. Each product can be shipped using two distinct transport modes. As each mode has its own merits, we have proposed a dynamic mode selection model that allows the company to ship products with either mode depending on when one mode is more favorable than the other. Since the optimal policy for dual transport mode problems are known to be complex, we have assumed that shipment shipment quantities for each product are governed by a dual-index policy. We have formulated the resulting decision problem as a mixed integer linear program that we have solved through a column generation solution procedure. This column generation procedure decomposes the com-

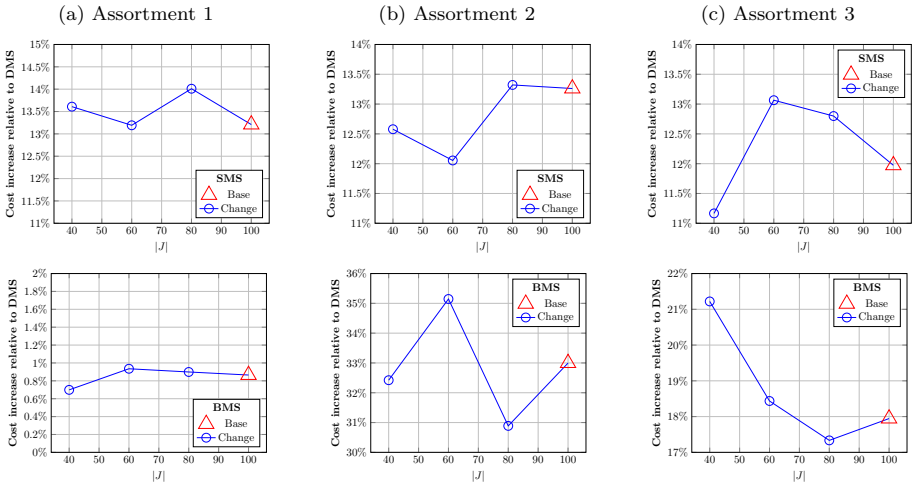

Figure 2.12: Effect of changing $|J|$ while keeping the rest of the parameters as in the base case.

plex multi-product problem into smaller sub-problems per product. These sub-problems are readily solved through a simple bisection search over Newsvendor type problems.

In an extensive computational experiment, we have compared the performance of our dynamic mode selection approach with two alternative approaches that are considered state of the art. The first benchmark, static mode selection, lacks the flexibility to dynamically ship products with two transport modes; it rather selects one transport mode for each product a priori. The second benchmark, blanket mode selection, does have the flexibility to rely on two transport modes simultaneously but it makes transport decisions for each product individually rather than holistically for the entire assortment. Our computational experiments indicate that the value of our dynamic mode selection approach over the blanket mode selection approach is particularly high for assortments of products for which the fastest transport modes are not necessarily the most polluting transport modes. For such settings, our dynamic mode selection approach can reduce the long run average costs by 40 percent under moderate carbon emission targets. These huge savings can be attributed to the portfolio effect inherent to our approach. The computational experiments further indicate that dynamic mode selection can significantly outperform static mode selection. Under moderate emission targets, dynamically relying on two transport modes rather than a single transport mode can lead to cost savings of up to 15 percent.

Future studies can extend the current model by studying other settings with multiple transport modes such as multiple echelons in a serial system (e.g. Lawson and Porteus, 2000; Arts and Kiesmüller, 2013) or assembly systems (e.g. Angelus and Özer, 2016). Alternatively one may consider more sophisticated dual mode heuristic policies such as the projected expedited inventory position policy (Drent and Arts, 2022), capped base-stock policy (Sun and Van Mieghem, 2019), or vector base-stock policy (Sheopuri et al., 2010).

2.A Carbon accounting

In this section, we briefly explain how we determine the distribution functions that we use for pseudo-random generation of the unit emissions in our computational experiments. We utilize the Database (2020) to calculate the average unit weights for 122 groups of products imported by The Netherlands in 2020. These product groups consist of two categories: (i) apparel goods that are imported from Vietnam and (ii) industrial goods that are imported from China and Germany. We consider air transport (from Tan Son Nhat international airport) and sea transport (from Haiphong port) as the fast and slow transport mode for the apparel category, respectively. For the industrial goods category, we assume sea transport from Shanghai, China, as the slow mode and road transportation from Stuttgart, Germany, as the fast mode.

We rely on the Network for Transport Measures methodology (NTM, 2015) to model and measure transportation emissions based on equation 2.1. This model has been widely used in literature (e.g., Hoen et al., 2014a,b). Following the NTM methodology, we first compute the overall carbon emissions generated by a single vehicle and then allocate a proportion of those emissions to each freight unit carried by the vehicle.

Sea transportation. We assume that all products are shipped via container. The average age of the container fleet worldwide is around 12 years and the average vessel size (dwt) of container ships with age 10-14 is 43,993 ton (United Nations Conference on Trade and Development, 2020). Based on section 7 of the NTM framework and resolutions of the Marine Environment Protection Committee (International Maritime Organization, 2011), we approximate sea transportation emissions in kilograms of CO_2 of one unit of a certain product with weight w (in kilogram) for a certain trip with distance d (in kilometers) using the following relation,

$$e_{sea} = w \cdot EI_{ship} \cdot 10^{-3} \cdot d$$

where EI_{ship} is kilograms of CO_2 emissions per kilogram weight per kilometer. Furthermore EI_{ship} is computed through,

$$EI_{ship} = \frac{(a \cdot dwt^{-c})/(PDR_{ship} \cdot LCU)}{1.852}$$

where a and c are constants, dwt is the deadweight tonnage of the ship, LCU is average load capacity utilization, $F(LCU)$ is fuel consumption as a function of load, and PDR_{ship} is the payload of the ship. 1.852 is the nautical mile to km conversion coefficient. For a container ship, NTM methodology states: $a = 0.17422, c = 0.201, LCU = 0.70, F(LCU) = 1$, and $PDR = 0.8$. Succinctly, we have for the total emissions in kilograms of CO_2 of one unit of a product with weight w (in kilogram) for a sea trip with distance d (in kilometers)

$$e_{sea} = w(1.996 \cdot 10^{-5} \cdot d). \tag{2.8}$$

Air transportation. Our calculations for the emissions of air transportation are based on section 8 of the NTM Framework. We consider an Airbus A310-300 F as the aircraft. Based on the May 2021 Air Cargo Market Analysis of The International Air Transport Association, we assume an average international cargo load factor of 65%. Following the NTM Framework, we have the following relation for air transportation emissions

$$e_{air} = \frac{w}{c_{max}}(CEF + VEF \cdot d),$$

where c_{max} is the maximum freight load, CEF is the constant emissions factor, and VEF is the variable emissions factor. CEF and VEF are the outcomes of applying a linear regression on real data provided by the NTM. We obtain the CEF and VEF parameters via interpolation over the associated tables provided by the NTM. We furthermore assume $c_{max} = 39,000kg$ as per section 8.3.1 and perform the interpolation on table 4.1 of section 8.2.1. Succinctly, we have for the total emissions in kilograms of CO_2 of one unit of a product with weight w (in kilogram) for an air trip with distance d (in kilometers)

$$e_{air} = w(1.525 \cdot 10^{-1} + 4.938 \cdot 10^{-4} \cdot d). \tag{2.9}$$

Road transportation. We rely on Hoen et al. (2014a) to obtain the emission units of road transportation. They too rely on the NTM framework to estimate CO_2 emissions from road transportation in Europe. In particular, they approximate the total emissions in kilograms of CO_2 of one unit of a product with weight w (in kilogram) for a truck trip with distance d (in kilometers) as

$$e_{road} = w(3.214 \cdot 10^{-4} + 4.836 \cdot 10^{-5} \cdot d). \tag{2.10}$$

Distances. We use NTMCalc Basic 4.0 (NTM, n.d.), which is an online tool provided by the NTM for approximating emissions, to calculate travel distances between the origin destination pairs as described at the beginning of this section. Based on this tool we find that the sea distance between Haiphong and Rotterdam is 9,610 nautical miles (17,798 km) and the sea distance between Shanghai and Rotterdam is 10,525 nautical miles (19,492 km). The distance traveled by aircraft between Tan Son Nhat international airport and Rotterdam The Hague Airport is 10,073 km, and the road distance between Stuttgart and Rotterdam is 633 km. With these distances, we compute the total kilogram CO_2 emissions of one unit of product with weight w for each mode-trip category using Equations (2.8)-(2.10). We call these emission coefficients and they are presented in Table 2.6 below.

Table 2.6: Emissions coefficients for each mode-trip category.

Industry	Slow Mode	Fast Mode
Apparel	$3.552 \cdot 10^{-1}$	5.127
Industrial	$3.891 \cdot 10^{-1}$	$3.093 \cdot 10^{-2}$

Fitting distribution functions. We use the emission coefficients from Table 2.6 to calculate the unit emissions for the 122 groups of products mentioned at the beginning of this section. We subsequently use maximum likelihood estimation on the resulting emission units to find distribution functions from which we can sample the emission units of the fast and slow transport modes for all three assortment types in our computational experiment. The emission unit distributions for assortment type 1 are based on the apparel category, the emission unit distributions for assortment type 2 are based on the industrial category, and the emission unit distributions for assortment type 3 are based on both categories. The final distribution functions are provided in Table 2.3.

2.B Generating correlated random numbers

Suppose X and Y are two real random variables with marginal distribution functions F and G, respectively. Suppose their joint distribution is bi-variate standard normal $\mathcal{N}_\rho$ with Pearson's correlation coefficient $\rho = \mathbf{Cov}(X,Y)/(\sqrt{\mathbf{Var}[X]}\sqrt{\mathbf{Var}[Y]})$ (Nelsen, 2006). Let Z be a vector of size two with independent random elements that have standard normal distributions Φ, and let $W = AZ$ be a linear combination of Z with

$$A = \begin{bmatrix} 1 & 0 \\ \rho & \sqrt{1-\rho^2} \end{bmatrix}.$$

It can be shown that W has a bivariate normal distribution $\mathcal{N}_\rho$ with covariance matrix $\Sigma = AA^T$ (see, e.g., Gut, 2009b). We use this result to sample from X and Y as follows:

1. Generate the vector $Z = \begin{bmatrix} Z_1 \\ Z_2 \end{bmatrix}$ by independently sampling from a standard normal distribution function,
2. Calculate the bivariate normal sample $W = \begin{bmatrix} W_1 \\ W_2 \end{bmatrix} = AZ$,
3. Generate the required samples by inversion $\begin{bmatrix} X \\ Y \end{bmatrix} = \begin{bmatrix} F^{-1}(\Phi(W_1)) \\ G^{-1}(\Phi(W_2)) \end{bmatrix}$.

2.C Benchmark approaches

The mathematical formulation for the blanket mode selection approach, which enforces emission constraints $\mathcal{E}_j^{max}$ per item $j \in J$, is called Problem (BMS) and is given as follows:

$$\begin{array}{lllll} (BMS) & \min & C(\mathbf{S}_f, \mathbf{\Delta}) & & \\ & \text{subject to} & E_j(S_{j,f}, \Delta_j) \leq \mathcal{E}_j^{max}, & & \forall j \in J, \\ & & (S_{j,f}, \Delta_j) \in (\mathbb{R} \times \mathbb{R}_0), & & \forall j \in J. \end{array}$$

The mathematical formulation for the static mode selection approach, which selects one transportation mode per item, is called Problem (*SMS*) and is given as follows:

$$
\begin{aligned}
(SMS) \qquad & \min && \sum_{j\in J} C_{j,s}(S^*_{j,s})x_{j,s} + \sum_{j\in J} C_{j,f}(S^*_{j,f})x_{j,f} \\
& \text{subject to} && \sum_{j\in J} E_{j,s}x_{j,s} + \sum_{j\in J} E_{j,f}x_{j,f} \leq \mathcal{E}^{max}, \\
& && x_{j,s}, x_{j,f} \in \{0,1\}, \qquad \forall j \in J,
\end{aligned}
$$

where $C_{j,s}(S_{j,s}) := h_j\mathbb{E}\left[\left(S_{j,s} - \sum_{t=0}^{l_{j,s}} D_j^t\right)^+\right] + p_j\mathbb{E}\left[\left(\sum_{t=0}^{l_{j,s}} D_j^t - S_{j,s}\right)^+\right] + c_{j,s}\mathbb{E}[D_j]$, and $C_{j,f}(S_{j,f}) := h_j\mathbb{E}\left[\left(S_{j,f} - \sum_{t=0}^{l_{j,f}} D_j^t\right)^+\right] + p_j\mathbb{E}\left[\left(\sum_{t=0}^{l_{j,f}} D_j^t - S_{j,f}\right)^+\right] + c_{j,f}\mathbb{E}[D_j]$ are the average cost-rate for using the slow and fast transport mode respectively, $S^*_{j,s} := \text{argmin}_{S_{j,s}} C_{j,s}(S_{j,s})$, and $S^*_{j,f} := \text{argmin}_{S_{j,f}} C_{j,f}(S_{j,f})$ are the base-stock levels that minimize inventory related costs when exclusively using the slow and fast transport mode respectively, and $E_{j,s} := e_{j,s}\mathbb{E}[D_j]$ and $E_{j,f} := e_{j,f}\mathbb{E}[D_j]$ are the emissions per time unit of shipping exclusively with the slow and fast transport mode respectively.

Chapter 3

Load asymptotics and dynamic speed optimization for the greenest path problem

3.1 Introduction

The transportation sector is one of the largest sources of anthropogenic CO_2 emissions, as attested by the Intergovernmental Panel on Climate Change (2021), US Environmental Protection Agency (2022), and the European Environment Agency (2021). In 2020, 36.3% of U.S. CO_2 emissions from fossil fuel combustion came from the transportation sector, of which 45.2% was generated by heavy-, medium-, and light-duty trucks (US Environmental Protection Agency, 2022). Similarly, the transportation sector accounted for 22% of the EU's CO_2 emissions in 2020 (European Environment Agency, 2021). Accordingly, there has been considerable attention on reducing CO_2 emissions through "green routing"; see e.g. Demir et al. (2012); Scora et al. (2015); Raeesi and Zografos (2019). The objective to reduce CO_2 emissions in transportation aligns with efforts to reduce fuel expenditure. The reduction of fuel consumption has become imperative as fuel increases in price and volatility due to recent geopolitical events, namely the Russian invasion of Ukraine (Goldfarb and Patterson, 2022).

Road gradient and vehicle speed are two major factors that influence the carbon footprint of a diesel truck (Demir et al., 2014). Demir et al. (2011) demonstrate through numerical analysis that a medium-duty truck may consume an additional six liters of diesel per 100 kilometer while traveling up a hill with a 1% gradient. The same study also shows that increasing the speed of an empty medium-duty truck from 50 km/h to 100 km/h can raise fuel consumption by more than 3% on a level path. Gravity is an important factor in finding the most efficient path between two points. Johann Bernoulli posed such a problem as early as 1696, in which a path was deemed efficient if the travel time was minimized and only gravity could be used to accelerate. The solution to this

problem gave rise to so-called brachistochrone curves, which differ from the shortest path between two points. Gravity and speed interact when finding the greenest (or most fuel-efficient) route between two points in a road network. The aim of the present chapter is to provide a thorough analysis of the difference between shortest paths and greenest paths as a function of speed and vehicle type for a large variety of geographic settings.

In principle, empirical methods are the most precise way of measuring carbon emissions associated with traversing a road with a certain vehicle at a certain speed. Unfortunately, it is not practical to empirically find carbon emissions for all roads, speeds, and vehicle types as well as many other parameters (e.g. road surface type) that affect fuel efficiency. Hence, several CO_2 emissions models for trucks have been proposed in literature. Demir et al. (2014) offers a summary of these models. The Comprehensive Modal Emission Model (CMEM) is an instantaneous emissions modeling approach that was introduced by Barth et al. (2005); Scora and Barth (2006); Boriboonsomsin and Barth (2009). Bektaş and Laporte (2011) and Demir et al. (2012) present a simplified variant of CMEM that is differentiable with respect to speed. This model is convenient in practical applications. Rao et al. (2016) and Brunner et al. (2021) show that this model can be made more realistic for cases where a vehicle travels downhill. Their modification of the CMEM, unfortunately, renders it no longer differentiable at all speeds. Over the past decade, CMEM has been the prevalent emissions model utilized in green/pollution vehicle routing problems (e.g. Bektaş and Laporte, 2011; Franceschetti et al., 2013; Huang et al., 2017; Xiao et al., 2020).

We call an optimization problem that seeks a path between an origin and destination a *path selection* problem. In this chapter, we focus on the selection of the greenest (most fuel-efficient) path. The greenest path is the path with the least CO_2 emissions. Some authors also call this the eco-friendly path (e.g. Scora et al., 2015; Andersen et al., 2013; Boriboonsomsin and Barth, 2009; Schröder and Cabral, 2019). Path selection is the backbone of a multitude of transport-based supply chain problems, from strategic supply chain network design to operational vehicle routing problems. The complexity of transportation problems forces many solution approaches to use path selection as a pre-processing activity. It is common to use either the shortest or the fastest path in this pre-rocessing step. The implicit assumption is that these paths are also the greenest. In this chapter, we show that the actual topology of urban road networks requires that we consider the greenest path selection as a part of the main optimization problem, e.g. vehicle routing problem (VRP) or supply chain network design (SCND).

The development of Geographic Information Systems (GIS) have made high-resolution geospatial data available at low cost. It is not sufficient to only consider the elevation of the origin and destination of a path. Rather, for any path, elevation along different sections of a path determine whether gravity increases or decreases the amount of fuel needed for travel. Thus, detailed elevation data of each segment of a possible path is required to find the greenest path. Furthermore, the slope along different segments of a

path also determines the most fuel efficient speed along each segment of a path.

In this chapter, we show that the most fuel efficient speed will change along different segments of any path. Thus, dynamic speed optimization is important to find the greenest path between any origin and destination. The greenest path also depends on the payload of a vehicle. We prove that the greenest path converges to an asymptotic greenest path as the payload approaches infinity and that this limiting path is attained for a finite payload. Our results are illustrated through numerical experiments. These experiments consider a setting wherein a logistics service provider seeks to reduce the CO_2 emissions of their transport operations. The company's fleet consists of heavy-, medium-, and light-duty trucks that operate in an urban environment. We use the modified CMEM proposed by Brunner et al. (2021) and focus our analysis on the effects that road gradient, speed, payload, and truck type have on CO_2 emissions.

We use an extensive numerical study to provide statistical answers to empirical research questions listed at the end of this section. We utilize the real road network and elevation data of 25 cities across six continents. It is worth noting that the closest paper to our work in terms of the CO_2 emissions model is Brunner et al. (2021). Apart from the differences of our objective functions, the main differences between our study and Brunner et al. (2021) are twofold. Firstly, Brunner et al. (2021) base their analysis on the static speed policy along different segments of a path. We show that a static speed policy can be suboptimal in terms of CO_2 emissions for traversing a path in a city with uneven topography. In addition, we demonstrate that the speed policy influences which path is the greenest. Secondly, Brunner et al. (2021) solve the path selection problem as a pre-processing activity for their main VRP problem. Moving the path selection to a pre-processing step forces them to consider fixed loads and speeds. By contrast, we consider dynamic speed optimization, and study asymptotics greenest paths as payloads increase. We utilize estimated traffic speeds during rush hour for a large subgraph of New York City's road network to study the potential CO_2 reduction by choosing the greenest paths and optimizing speed instead of taking the fastest paths. We also examine the increased travel duration on the greenest paths, as well as the convergence to the asymptotic greenest path when traffic congestion occurs.

The main contributions of this chapter are listed below:

1. We show that the greenest path is speed and payload dependent for accurate emission models. We provide a tractable algorithm to optimize the path and the speed jointly, where the speed varies along the path.

2. We show that the greenest path converges to an asymptotic greenest path when the load becomes large and that this path is attained for a finite load in Section 3.3.5. In Sections 3.4.3 and 3.4.4, we show that this convergence happens relatively quickly in practice. We also show, in Section 3.4.3, that the greenest path for a slope-dependent optimal speed policy is quite similar to the one associated with

the static speed policy of Demir et al. (2012), yet significantly different from the shortest path.

3. We conduct an extensive numerical study with data from 25 cities over 6 continents and over 3 million origin destination pairs. We use detailed elevation data from U.S. Geological Survey (2000)'s SRTM 1 Arc-Second Global data set. This thorough study allows us to answer the following research questions:

   (a) How much CO_2 emissions can dynamic speed optimization and green path selection reduce jointly? What are the effects of truck type, payload, and city on the carbon reduction potential?

   (b) What is the marginal contribution of speed optimization and path optimization in the reduction of CO_2 emissions?

   (c) How different is the slope-dependent greenest path from the shortest path (which is the slope-disregarding path)?

   (d) What is the impact of the speed policy on the greenest path?

   (e) In what settings are the integration of elevation data in path selection most valuable?

4. We conduct an extensive numerical study with road network, elevation, and speed data of New York City for more than 20 thousands origin destination pairs. The traffic speed estimates are collected from Google's Distance Matrix API. With this study we answer to the questions 3a through 3d when shortest path is replaced by the fastest path. We also study the increased time in traffic when the greenest path and speed optimization are decided.

The rest of this chapter is organized as follows. We review the related literature in Section 3.2. Section 3.3 describes the mathematical model used in this study and the policies that can minimize CO_2 emissions. Section 3.4 provides the setting and results of the extensive numerical studies under free-flow conditions. We present the setting and results of our numerical study under traffic congestion in Section 3.5. Finally, we offer the conclusions and final remarks in Section 3.7.

3.2 Literature review

Green transportation has been studied extensively over various decision-making settings. Asghari and Alehashem (2021); Moghdani et al. (2021); Demir et al. (2014) give reviews on the most important recent literature on the green VRP. Additionally, Waltho et al. (2019) reviews pivotal studies in the field of green SCND from 2010 to 2017. In most of the main stream green VRP and SCND, the path between every two nodes of interest

is computed as a pre-processing step (e.g. Demir et al., 2012). This has been partially relaxed for the VRP by Behnke and Kirschstein (2017). In other words, road networks are reduced to distances between origin and destination pairs to simplify later computations. The implicit assumption is that distances or travel times are the main drivers of costs and/or emissions. This paper extensively studies to which extent this implicit assumption is tenable.

A large body of work in the field of green transportation relies on macroscopic (average aggregate), microscopic (instantaneous) fuel consumption models, or a combination of both. Demir et al. (2014) and Zhou et al. (2016) provide an extensive review of fuel consumption models. A number of studies, including Boriboonsomsin et al. (2012); Scora et al. (2015), and Ericsson et al. (2006), estimate the CO_2 emissions of a specific vehicle based on the measurement of that vehicle. Demir et al. (2014) explain the main factors that influence fuel consumption in road freight transportation. Among the pertinent determinants for the case of the greenest path are road gradient, speed, truck type, and payload.

The path optimization under environmental consideration (the greenest path) has been explored over the past two decades. This problem can be formulated based on a variant of the shortest path algorithm of Dijkstra (1959) to minimize the total fuel consumption of a vehicle between two nodes. Ericsson et al. (2006) studies the CO_2 emissions of light-duty cars by using a navigation system that computes the greenest path based on in-vehicle data and traffic information in Lund, Sweden. They conclude that selecting the greenest path can reduce fuel consumption by 4% on average in Lund. Boriboonsomsin et al. (2012) presents an Eco-Routing Navigation System (EFNav) as a framework to integrate GIS and traffic data with emissions model estimates to compute eco-friendly paths for light vehicles. Scora et al. (2015) extend the EFNav model to heavy-duty trucks (EFNav-HDT) and conduct a numerical study to test the benefits of EFNav-HDT across different vehicle weights in Southern California. Scora et al. (2015) provide excellent insights into the specifications of the greenest path for trucks. Both Boriboonsomsin et al. (2012) and Scora et al. (2015) base their studies on the CMEM model and estimate the energy/emissions model using linear regression over data from actual measurements. Boriboonsomsin et al. (2012) take advantage of a logarithmic transformation and Scora et al. (2015) use a minimum fuel cutoff point to avoid negative fuel consumption results. Andersen et al. (2013) take advantage of free road network data, such as OpenStreetMap, and use Controller Area Network (CAN bus) data to compute the greenest path by assigning weights to the different segments of the network. Since this work does not rely on a fuel consumption model, it is very accurate for the paths and vehicles for which fuel consumption data is available, but it does not transfer to other settings without the collection of a large amount of data in that setting. Pamučar et al. (2016) utilize a similar approach and include other negative externalities associated with transportation, such as noise, land use, and pollutants other than CO_2 . Schröder and Cabral (2019) consider a Digital Elevation Model and Copert III emissions model to

compute the greenest path. Dündar et al. (2022) propose an approach to increase the resolution of the road network and compute the fuel consumption over along a path more accurately.

Speed optimization as a means to reduce the emissions and driving costs was first introduced by Demir et al. (2012). Franceschetti et al. (2013) present a speed optimization technique that can also be used for traffic congestion. Both of these works, as well as many other well-cited papers, such as Lai et al. (2024), rely on the CMEM model of Demir et al. (2011), which results in negative fuel consumption over many downhill paths (Brunner et al., 2021). Brunner et al. (2021) modify the fuel consumption model, yet only consider a constant travel speed. Some papers consider the fastest path or the emissions minimizing path under dynamic speeds induced by congestion (e.g. Ehmke et al., 2016a,b; Huang et al., 2017; Ehmke et al., 2018).

In this chapter, we consider the modified CMEM (Brunner et al., 2021). We explore the individual and combined effects of elevation, speed optimization, truck type, payload, and characteristic city topography on CO_2 emissions reduction and the greenest path policies. This chapter, is the first work to provide asymptotic results for a path selection problem and the the greenest path problem in particular.

3.3 Model Description

In this section, we introduce the notations (Section 3.3.1) and mathematical foundations of our research, including the CO_2 emissions models (Section 3.3.2) together with the optimal speed policies (Section 3.3.3). We formally introduce the greenest paths between two locations in a city road network and discuss how optimal speed policies complicate the computation of the greenest path (Section 3.3.4). We study the asymptotic behavior of the greenest path when the payload increases (Section 3.3.5). In Section 3.3, we only consider speed, payload, and/or path (or a single arc) as the explicit arguments of functions, since these three factors are the focus of our analysis in Sections 3.3.3, 3.3.4, and 3.3.5.

3.3.1 City Network and Notations

Let a directed graph $\mathcal{G} = (V, A)$ represent the road network of a city, where $V = \{1, \ldots, m\}$ is the set of m vertices, the points of interest along the roads (e.g. road intersections), and $A \subseteq V \times V$ is the set of arcs (road segments) that connects the vertices. Any arc $a \in A$ has the following features: the length $\delta : A \to \mathbb{R}_{++}$, the angle $\theta : A \to \mathbb{R}$, the maximum allowable speed by $v^{\max} : A \to \mathbb{R}_{++}$, and the minimum allowable speed by $v^{\min} : A \to \mathbb{R}_{++}$, where $\mathbb{R}_{++} = \{x \in \mathbb{R} : x > 0\}$. We consider an internal combustion engine truck that traverses an arc $a \in A$ with speed $v \in [v^{\min}(a), v^{\max}(a)]$.

v is constant along arc a, but the speed of the truck can vary on other arcs. The truck consumes f_a liters of diesel fuel and produces e_a kilograms of CO_2 to traverse arc $a \in A$ (v will be selected to minimize f_a and e_a according to different emission models). Notation, including those of truck properties, are listed in Table 3.1.

3.3.2 Emission Models

We discuss two emission models. The first of these models is most commonly used in recent papers on the green/pollution routing problem (e.g. Bektaş and Laporte, 2011; Demir et al., 2012; Franceschetti et al., 2013; Dabia et al., 2017). We will call this the standard model. The second model is a small improvement on the standard model to disallow negative fuel consumption on downward sloping road segments.

Standard Emissions Model

The CMEM (Barth et al., 2005; Scora and Barth, 2006; Boriboonsomsin and Barth, 2009) is a microscopic truck fuel consumption model that has been widely used in literature for pollution/green vehicle routing problems. The Standard model is an instantiation of the CMEM approach. Suppose a truck with the parameters given in Table 3.1 and payload l travels along arc $a \in V$ with speed v. In the standard emission model introduced by Bektaş and Laporte (2011) and Demir et al. (2012), the truck's fuel consumption is given by:

$$\tilde{f}_a(v,l) = \frac{P\delta(a)}{v} + Q\delta(a)\left(g\sin\theta(a) + C_r g\cos\theta(a)\right)(w+l) + R\delta(a)v^2 \tag{3.1}$$

$$\text{with } P = \frac{\xi k N D}{\kappa\psi}, \tag{3.2}$$

$$Q = \frac{\xi}{1000\eta\eta_{tf}\kappa\psi}, \tag{3.3}$$

$$R = \frac{\xi C_d \rho S}{2000\eta\eta_{tf}\kappa\psi}. \tag{3.4}$$

The main assumption behind Equation (3.1) is that the truck parameters remain constant along each arc $a \in V$. This model sets aside a number of minor sources of fuel consumption, such as air conditioning and compressed air systems. Burning one liter of diesel in a combustion engine produces $c_e = 2.67$ kg/L of CO_2 (Agency, 2005). Thus we find that the CO_2 emissions associated with traversing an arc a with load l at speed v is given by,

$$\tilde{e}_a(v,l) = c_e\tilde{f}_a(v,l) = 2.67\tilde{f}_a(v,l)$$

under the standard model.

Table 3.1: Overview of notation.

Notation	Description
Sets	
$\mathcal{G}$	Directed multigraph representing the urban road network, $\mathcal{G} = (V, A)$.
V	Set of vertices of $\mathcal{G}$, $V = \{1, \dots, m\}$, where m is the number of vertices.
A	Set of arcs of $\mathcal{G}$, $A \subseteq V \times V$.
Π	Set of all paths between a pair of nodes $n_s, n_t \in A$.
Π_d	Subset of Π such that for all $a \in \Pi_d$, $tan\theta(a) < -C_r$.
Network Features	
$\delta(a)$	$\delta : A \to \mathbb{R}_{++}$, length of arc $a \in A$.
$\theta(a)$	$\theta : A \to \mathbb{R}$, angle of arc $a \in A$.
$h'(a)$	$h' : A \to \mathbb{R}$, augmented ascent of arc $a \in A$, i.e. $h'(a) = \delta(a) \sin(\theta(a) + \arctan C_r)^+$.
$v^{\max}(a)$	$v^{\max} : A \to \mathbb{R}_{++}$, maximum allowable speed for traversing arcs $a \in A$.
$v^{\min}(a)$	$v^{\min} : A \to \mathbb{R}_{++}$, minimum allowable speed for traversing arcs $a \in A$.
Parameters	
ξ	Fuel-to-air mass ratio.
g	Gravitational constant (m/s^2).
ρ	Air density (kg/m^3).
C_r	Coefficient of rolling resistance.
η	Efficiency parameter for diesel engines
η_{tf}	Vehicle drivetrain efficiency.
κ	Heating value of a typical diesel fuel (kJ/g).
ψ	Conversion factor (g/s to L/s).
w	Curb weight (kg).
L	Maximum payload (kg).
k	Engine friction factor (kJ/rev/L).
N	Engine speed (rps).
D	Engine displacement (L).
C_d	Coefficient of aerodynamic drag.
S	Frontal surface area (m^2).
c_e	Fuel's CO_2 Emissions Coefficient.
Variables	
l	$l \in \mathbb{R}_+$, payload (kg).
v	$v \in \mathbb{R}_{++}$, speed of a truck to traverse arc $a \in A$.
c_v	Constant speed (Equation (3.7)) that minimizes the standard emissions model (Equation (3.1)).
$v(a)$	$v : A \to \mathbb{R}_{++}$, speed policy for a truck to traverse arc $a \in V$.
$v^d(a, l)$	$v^d : A \times \mathbb{R}_+ \to \mathbb{R}_{++}$, dynamic speed policy on arc $a \in A$ with payload l as per Proposition 3.1.
$v^s(a)$	$v^s : A \to \mathbb{R}_{++}$, static speed policy on arc $a \in A$ as per Equation (3.7).
$v^t(a, l)$	$v^t : A \times \mathbb{R}_+ \to \mathbb{R}_{++}$, terminal velocity on arc $a \in A$ with payload l as per Proposition 3.1.
$t_a(v)$	$t_a : \mathbb{R}_{++} \to \mathbb{R}_{++}$, traveling time on arc $a \in A$ with speed $v \in \mathbb{R}_{++}$.
π	$\pi \in \Pi$, path between a pair of nodes $n_s, n_t \in A$.
π^{sp}	$\pi^{sp} \in \Pi$, the shortest path between a pair of nodes (Equation (3.14)).
$\pi^g(v, l)$	$\pi^g(v, l) \in \Pi$, the most fuel-efficient (greenest) path between a pair of nodes under the speed policy v and payload l (Equation (3.13)).
$\pi^\infty(v)$	$\pi^\infty(v) \in \Pi$, the asymptotic greenest path between a pair of nodes under the speed policy v, i.e. the greenest path when the payload is arbitrarily large (Proposition 3.4).
$\tilde{f}_a(v, l)$	$\tilde{f}_a : \mathbb{R}_{++} \times \mathbb{R}_+ \to R_{++}$, amount of fuel (liter) that a truck consumes for traversing arc $a \in V$ with speed v and payload l under the standard emissions model (Section 3.3.2).
$f_a(v, l)$	$f_a : \mathbb{R}_{++} \times \mathbb{R}_+ \to R_{++}$, amount of fuel (liter) that a truck consumes for traversing arc $a \in V$ with speed v and payload l under the improved emissions model (Section 3.3.2).
$\tilde{e}_a(v, l)$	$\tilde{e}_a : \mathbb{R}_{++} \times \mathbb{R}_+ \to R_{++}$, amount of CO_2 (kg) that a truck emits for traversing arc $a \in V$ with speed v and payload l under the standard emissions model (Section 3.3.2).
$e_a(v, l)$	$e_a : \mathbb{R}_{++} \times \mathbb{R}_+ \to R_{++}$, amount of CO_2 (kg) that a truck emits for traversing arc $a \in V$ with speed v and payload l under the improved emissions model (Section 3.3.2).
$\mathcal{E}(\pi, v, l)$	Total amount of CO_2 emitted by a truck along a path π under the speed policy v and payload l, i.e. $\mathcal{E}(\pi, v, l) = \sum_{a \in \pi} e_a(v(a), l)$.

Improved Emissions Model

Rao et al. (2016) and Brunner et al. (2021) establish that the standard emission model gives rise to negative fuel consumption on some negative road angles that are not realistic

for internal combustion engine vehicles. Thus, Brunner et al. (2021) propose the following adjustment to Equation (3.1):

$$f_a(v,l) = \frac{P\delta(a)}{v} + \Big(Q\delta(a)\left(g\sin\theta(a) + C_r g\cos\theta(a)\right)(w+l) + R\delta(a)v^2\Big)^+, \tag{3.5}$$

where $(\cdot)^+ = \max\{\cdot, 0\}$. Equation (3.5) shows that gravity works in favor of the vehicle over downhills arcs to compensate the energy loss caused by drag and rolling resistance force. This equation assumes that any engine-powered brakes consume a negligible amount of fuel. The standard and improved emission models (3.1) and (3.5) are identical on a flat network ($\theta(a) = 0$ for all $a \in A$). We note that a slight modification of the above models can also allow for electric vehicles with regenerative braking; see Larminie and Lowry (2012). As before we now find that the CO_2 emissions associated with traversing arc a with load l at speed v is given by,

$$e_a(v,l) = c_e f_a(v,l) = 2.67 f_a(v,l). \tag{3.6}$$

3.3.3 Optimal Speed

The most fuel efficient speed to traverse an arc depends on the emission model that is used. We will show below that there is one optimal speed for all arcs in a network under the standard emission model, but that the optimal speed may differ by arc for the improved emission model.

Static Speed Optimization

The *speed optimization problem (SO)* is to compute the speed policy which minimizes the carbon emissions when a vehicle travels across an arc $a \in A$. Under the standard emissions model, SO can be formulated as,

$$\tilde{e}_a^* = \min_{v\in[v^{\min}(a),v^{\max}(a)]} \tilde{e}_a(v,l) \quad \text{and} \quad v^s(a) = \arg\min_{v\in[v^{\min}(a),v^{\max}(a)]} \tilde{e}_a(v,l).$$

This implies that the most fuel efficient speed is the same along any arc $a \in A$ and is given by $v^s : A \to \mathbb{R}_{++}$ that is defined by,

$$v^s(a) := \begin{cases} v^{\min}(a) & \text{if } c_v \le v^{\min}(a) \\ c_v & \text{if } v^{\min}(a) < c_v \le v^{\max}(a) \\ v^{\max}(a) & \text{if } v^{\max}(a) < c_v, \end{cases} \tag{3.7}$$

where c_v,

$$c_v = \sqrt[3]{\frac{P}{2R}}. \tag{3.8}$$

is the optimal speed without any speedlimits. Expressions for P and R are given by Equations (3.2) and (3.4). Equation (3.8) is obtained by solving the first order conditions to minimize (3.1) with respect to v. Since c_v is constant along all arcs, we use the term *static speed* policy to denote a policy that will have a vehicle traverse every arc at the speed v^s. We note that for practically meaningful values of $v^{\min}(a)$ and $v^{\max}(a)$ the optimal speed is usually given by (3.8), or the second case in (3.7).

Dynamic Speed Optimization

In the improved emissions model, the most fuel-efficient speed to traverse an arc depends on its slope, if the slope is sufficiently negative. Under the improved emissions model (3.5), the speed optimization problem is formulated as,

$$e_a^* = \min_{v\in[v^{\min}(a),v^{\max}(a)]} e_a(v,l) \quad \text{and} \quad v^d(a,l) = \arg\min_{v\in[v^{\min}(a),v^{\max}(a)]} e_a(v,l). \tag{3.9}$$

Note that the derivative of

$$e_a(v,l) = \frac{c_e P\delta(a)}{v} + c_e\left(Q\delta(a)(g\sin\theta(a) + C_r g\cos\theta(a))(w+l) + R\delta(a)v^2\right)^+$$

with respect to v is given by

$$\frac{\partial e_a(v,l)}{\partial v} = \begin{cases} -\dfrac{c_e P\delta(a)}{v^2} & \text{if } 0 \le v < v^t(a,l) \\ -\dfrac{c_e P\delta(a)}{v^2} + 2c_e R\delta(a)v & \text{if } v^t(a,l) < v(a), \end{cases}$$

where $v^t : A \times \mathbb{R}_+ \to \mathbb{R}_+$ is defined by,

$$v^t(a,l) := \begin{cases} \sqrt{\dfrac{-Q(g\sin\theta(a) + C_r g\cos\theta(a))(w+l)}{R}}, & \text{if } \tan\theta(a) < -C_r \\ 0, & \text{if } \tan\theta(a) \ge -C_r. \end{cases} \tag{3.10}$$

This derivation shows that the CO_2 emissions of an arc $e_a(v,l)$ is not differentiable with respect to v at the point $v^t(a,l)$. The speed v^t has a physical interpretation as the terminal velocity of a vehicle on a slope with inclination θ. It is the speed at which the gravitational force along the slope equals the sum of the drag and rolling resistance forces (see e.g. Fox et al. (2020)). A vehicle reaches a non-zero terminal velocity on an arc if the angle falls below $-\arctan C_r$. The optimal solution to the speed optimization problem in (3.9) is slightly more involved as it needs to account for the terminal velocity. The solution is given in Proposition 3.1.

Proposition 3.1. *The optimal solution to the speeds optimization problem in* (3.9) *is*

given by $v^d : A \times \mathbb{R}_+ \to \mathbb{R}_{++}$ *that is defined by,*

$$v^d(a,l) := \arg \min_{v \in [v^{\min}(a), v^{\max}(a)]} e_a(v,l)$$
$$= \begin{cases} v^{\min}(a), & \text{if } \max\{c_v, v^t(a,l)\} \le v^{\min}(a) \\ \max\{c_v, v^t(a,l)\}, & \text{if } v^{\min}(a) < \max\{c_v, v^t(a,l)\} \le v^{\max}(a) \\ v^{\max}(a), & \text{if } v^{\max}(a) < \max\{c_v, v^t(a,l)\}. \end{cases} \tag{3.11}$$

Proof of Proposition 3.1. We consider the case where the terminal velocity is zero, and where it is strictly positive separately.

Case 1 ($\tan\theta(a) \ge -C_r$; $v^t(a,l) = 0$): Equation (3.5) reduces to Equation (3.1) so that one may verify that

$$v^d(a,l) = c_v > v^t(a,l) = 0.$$

Case 2 ($\tan\theta(a) < -C_r$; $v^t(a,l) > 0$): In this case,

$$e_a(v,l) =$$
$$\begin{cases} e_a^1 = \dfrac{c_e P \delta(a)}{v}, & \text{if } 0 \le v < v^t(a,l) \\ e_a^2 = c_e \left(\dfrac{P\delta(a)}{v} + Q\delta(a)(g\sin\theta(a) + C_r g\cos\theta(a))(w+l) + R\delta(a)v^2 \right), & \text{if } v^t(a,l) \le v. \end{cases}$$

It is straightforward to verify that $e_a(v,l)$ is continuous, e_a^1 is convex and non-increasing in v, and e_a^2 is convex in v with a minimum at c_v. Consequently, the optimal speed exceeds the terminal velocity, i.e. $v^d(a,l) \ge v^t(a,l)$.

As $e_a(v,l)$ is convex in v on $[v^t(a,l), \infty)$, it has an extremum at c_v if $v^t(a,l) \le c_v$, or at $v^t(a,l)$ if $v^t(a,l) > c_v$. It follows that the optimal speed is $\max\{c_v, v^t(a,l)\}$ if it lies within the allowable speed interval, $[v^{\min}(a), v^{\max}(a)]$. In case $\max\{c_v, v^t(a,l)\} < v^{\min}(a)$, then $e_a(v,l)$ is non-decreasing in $v \in [v^{\min}(a), v^{\max}(a)]$ and the optimal speed is $v^{\min}(a)$. If $\max\{c_v, v^t(a,l)\} \ge v^{\max}(a)$, then $e_a(v,l)$ is non-increasing in $v \in [v^{\min}(a), v^{\max}(a)]$ and the optimal speed is $v^{\max}(a)$. □

The main insight from Proposition 3.1 is that it is efficient to use gravity to reduce the required engine power and emission. Proposition 3.1 indicates that a static speed policy is not optimal on a path that contains downhill arcs. Thus, a *dynamic speed* policy (v^d), as per Proposition 3.1, reduces a truck's fuel consumption, CO_2 emissions, and travel time since it requires higher speeds on downhills.

3.3.4 The Greenest Path

Let n_s and n_t be two different vertices of $\mathcal{G}$ such that n_t is reachable from n_s. Let Π be the set of all possible paths between n_s and n_t. Under a given speed policy $v : A \to \mathbb{R}_{++}$

and a constant payload l, the total CO_2 emissions of a truck to travel between n_s and n_t along a path $\pi \in \Pi$, $\mathcal{E}(\pi, v, l)$, is defined as,

$$\mathcal{E}(\pi, v, l) = \sum_{a \in \pi} e_a(v(a), l). \tag{3.12}$$

Based on this definition, the *greenest path problem (GPP)* is to compute the path with the least CO_2 emissions, π^g, between n_s and n_t, i.e.

$$\begin{aligned} \mathcal{E}^*(v, l) &= \min_{\pi \in \Pi} \mathcal{E}(\pi, v, l) = \min_{\pi \in \Pi} \sum_{a \in \pi} e_a(v(a), l), \quad \text{and,} \\ \pi^g(v, l) &= \arg\min_{\pi \in \Pi} \mathcal{E}(\pi, v, l) = \arg\min_{\pi \in \Pi} \sum_{a \in \pi} e_a(v(a), l). \end{aligned} \tag{3.13}$$

We define the *shortest path problem (SPP)* as the computation of the minimum-distance path (π^{sp}) between n_s and n_t, i.e.

$$\delta^{sp} = \min_{\pi \in \Pi} \sum_{a \in \pi} \delta(a) \quad \text{and} \quad \pi^{sp} = \arg\min_{\pi \in \Pi} \sum_{a \in \pi} \delta(a). \tag{3.14}$$

The following proposition shows that if the elevation data is ignored and the speeds are identical along all arcs then the shortest path, π^g, is an optimal solution to GPP.

Proposition 3.2. *If the road gradient $\theta(a) = 0$ and the speeds $v(a)$ are identical for all arcs $a \in A$, then the Greenest Path ($\pi^g(v, l)$) is the Shortest Path (π^{sp}) .*

Proof of Proposition 3.2. Let angle $\theta(a) = 0$ for all $a \in A$, and the payload l and speed policy $v(a)$ be identical, i.e. $v(a) = v^*$, where $v^* \in \mathbb{R}$ is constant. Taking into account that $\sin\theta(a) = 0$ and $\cos\theta(a) = 1$ for all $a \in A$, the GPP implies that,

$$\begin{aligned} \mathcal{E}^* &= c_e \min_{\pi \in \Pi} \sum_{a \in \pi} \frac{P\delta(a)}{v(a)} + Q\delta(a)(g\sin\theta(a) + C_r g\cos\theta(a))(w + l) + R\delta(a)v(a)^2 \\ &= c_e \min_{\pi \in \Pi} \sum_{a \in \pi} \frac{P\delta(a)}{v^*} + Q\delta(a)C_r g(w + l) + R\delta(a)v^{*2} \\ &= c_e \left(\frac{P}{v^*} + QC_r g(w + l) + Rv^{*2}\right) \min_{\pi \in \Pi} \sum_{a \in \pi} \delta(a) \\ &= c_e \left(\frac{P}{v^*} + QC_r g(w + l) + Rv^{*2}\right) \delta^*_{n_s, n_t}. \end{aligned}$$

Thus, the π^{sp} satisfies this problem that proves the proposition. □

When the speeds are bounded by traffic or variable speed limits, then the analogous result holds for the fastest path. It is straightforward to verify that the greenest path is the fastest path when all road gradients are zero and the speeds are constant ; see Proposition 3.2. Further notice that by Proposition 3.1, the speed c_v in (3.7) is optimal for all arcs when $\theta(a) = 0$ for all $a \in A$. This implies that a decision maker will believe the shortest path is the greenest path when she ignores elevation data.

Nonetheless, the improved emissions model and Proposition 3.1 show that if the elevation data is considered, the speed along each segment of a path can change. Even under the static speed policy, the greenest path is not necessarily the shortest due to the non-linearity of emission along an arc when in the gradient. We note that the emission model does not explicitly account for acceleration and deceleration of a vehicle and so the estimates emissions $\mathcal{E}^*(v^d, l)$ are a lower-bound for the CO_2 emissions of a truck traveling from n_s to n_t.

3.3.5 The Asymptotic Greenest Path

In this section, we explore the greenest path as the payload becomes arbitrarily large. Let $e'_a(v)$ be the CO_2 emissions per unit payload when a truck traverses arc $a \in A$ with speed v, that is to say,

$$e'_a(v) = \frac{e_a(v,l)}{l} \tag{3.15}$$

$$= \frac{c_e P \delta(a)}{vl} + c_e \left(Q\delta(a)(g \sin\theta(a) + C_r g \cos\theta(a)) \left(1 + \frac{w}{l}\right) + \frac{R}{l}\delta(a)v^2 \right)^+. \tag{3.16}$$

Consider two distinct connected vertices n_s and n_t. Observe that under a speed policy $v : A \to \mathbb{R}^+$ and a constant load $l \in \mathbb{R}_+$, the greenest path, i.e. $\pi^g(v, l)$, minimizes the total CO_2 emissions and the total CO_2 emissions per unit payload between n_s and n_t. Thus, we can interchangeably use the total CO_2 emissions and the total CO_2 emissions per unit payload to compute the greenest path.

Let Π be the set of all paths from n_s to n_t. Let $\Pi_d \subseteq \Pi$ be the subset of paths Π that are entirely downhill with a slope below $\arctan(-C_r)$, i.e. $\tan\theta(a) < -C_r$ for all $a \in \pi$ with $\pi \in \Pi_d$. We can now state the definition of the asymptotic greenest path:

Definition 3.3. The asymptotic greenest path satisfies

$$\pi^\infty(v) \in \begin{cases} \arg\min_{\pi\in\Pi} \lim_{l\to\infty} \sum_{a\in\pi} e_a(v(a)) & \text{if } \Pi_d \neq \emptyset \\ \arg\min_{\pi\in\Pi} \lim_{l\to\infty} \sum_{a\in\pi} e'_a(v(a), l) & \text{if } \Pi_d = \emptyset. \end{cases} \tag{3.17}$$

Note that the set Π_d plays an important role in this definition. The emission per load vanishes for any sufficiently steep down downhill path ($\Pi_d \neq \emptyset$) because gravity will get the vehicle to its destination. Among all sufficiently steep downhill paths ($\pi \in \Pi_d$), the one with the least absolute emission is given by the second case in (3.17). When gravity does not suffice to move a vehicle from its origin to its destination ($\Pi_d = \emptyset$) then the asymptotic greenest path is the one that minimizes emissions per load; see case 1 in (3.17). The following proposition demonstrates that $\pi^\infty(v)$ exists and provides an explicit form to compute it.

Proposition 3.4. *$\pi^\infty(v)$ exists and can be computed as follows.*

$$\pi^\infty(v) \in \begin{cases} \arg\min_{\pi\in\Pi_d} \sum_{a\in\pi} t_a(v(a)) & \text{if } \Pi_d \neq \emptyset \\ \arg\min_{\pi\in\Pi} \sum_{a\in\pi} h'(a) & \text{if } \Pi_d = \emptyset, \end{cases} \tag{3.18}$$

where $t_a : \mathbb{R}_{++} \to \mathbb{R}_{++}$, is defined by

$$t_a(v) = \frac{\delta(a)}{v},$$

and $h' : A \to \mathbb{R}_+$, is defined by

$$h'(a) = \delta(a) \sin\left(\theta(a) + \arctan C_r\right)^+,$$

if $-90^\circ < \theta(a) + \arctan C_r < 90^\circ$ for all $a \in A$.

We call $\pi^\infty(v)$ the asymptotic greenest path. Proposition 3.4 explains that $\pi^\infty(v)$ is the fastest downward path $\pi \in \Pi_d$, if Π_d is non-empty. Otherwise, it is the path with the minimum total augmented ascents, h'. Evidently, $\pi^\infty(v)$ can be computed using the algorithms offered to solve the shortest path problem (e.g. Dijkstra (1959)). The requirement that $-90^\circ < \theta(a) + \arctan C_r < 90^\circ$ for all $a \in A$ is completely benign.

Proof of Proposition 3.4. For all payloads $l \in \mathbb{R}_+$, and any speed policy v, $\pi^g(v, l)$ exists from n_s to n_t, since by Equations (3.6) and (3.15) there are no negative emissions cycles between the vertices. Suppose that the payload l satisfies,

$$l \geq \max_{a\in A} \left\{ \frac{R(v^{\max}(a))^2}{-Q(g\sin\theta(a) + C_r g\cos\theta(a))} - w \right\}. \tag{3.19}$$

Then for arc $a \in A$,

$$e_a(v(a), l) = \begin{cases} \dfrac{c_e P \delta(a)}{v(a)} & \text{if } \tan\theta(a) < -C_r \\ \dfrac{c_e P \delta(a)}{v(a)} + c_e \left(Q\delta(a)(g\sin\theta(a) + C_r g\cos\theta(a))(w+l) + R\delta(a)v(a)^2 \right) & \text{if } \tan\theta(a) \geq -C_r, \end{cases} \tag{3.20a}$$

$$e'_a(v(a)) = \begin{cases} \dfrac{c_e P \delta(a)}{l v(a)} & \text{if } \tan\theta(a) < -C_r \\ \dfrac{c_e P \delta(a)}{l v(a)} + c_e \left(Q\delta(a)(g\sin\theta(a) + C_r g\cos\theta(a))(1 + \dfrac{w}{l}) + \dfrac{R\delta(a)v(a)^2}{l} \right) & \text{if } \tan\theta(a) \geq -C_r, \end{cases} \tag{3.20b}$$

by Equations (3.6) and (3.15).

Suppose that Π_d is a non-empty set. For this case, we use the total CO_2 emissions to compute the $\pi^\infty(v)$. Thus, by Equations (3.20a),

$$\lim_{l\to\infty}\sum_{a\in\pi} e_a(v(a),l) = \begin{cases}\sum_{a\in\pi}\frac{c_eP}{v(a)}\delta(a) & \text{if } \pi\in\Pi_d\\ \infty & \text{if } \pi\in\Pi\setminus\Pi_d.\end{cases}$$

Then it follows that from Equation (3.17) that

$$\pi^\infty(v) = \arg\min_{\pi\in\Pi_d}\sum_{a\in\pi}\frac{c_eP\delta(a)}{v(a)} = \arg\min_{\pi\in\Pi_d}\sum_{a\in\pi}\frac{\delta(a)}{v(a)} = \arg\min_{\pi\in\Pi_d}\sum_{a\in\pi} t_a(v(a)),$$

since P and c_e are constant across all arcs $a\in A$.

Now, suppose that Π_d is an empty set. For this case, we use the total CO_2 emissions per unit load to compute $\pi^\infty(v)$. Thus, by Equation (3.20b),

$$\begin{aligned}\lim_{l\to\infty}\sum_{a\in\pi} e'_a(v(a)) =& \sum_{a\in\pi} c_eQ(g\sin\theta(a)+C_rg\cos\theta(a))^+\delta(a)\\ =& c_eQg\sqrt{1+{C_r}^2}\sum_{a\in\pi}\delta(a)\left(\sin\left(\theta(a)+\arctan C_r\right)\right)^+\\ =& c_eQg\sqrt{1+{C_r}^2}\sum_{a\in\pi}\delta(a)\sin\left(\theta(a)+\arctan C_r\right)^+,\end{aligned}$$

as $-90^\circ < \theta(a)+\arctan C_r < 90^\circ$ for all $a\in A$ by supposition. Again, since c_e, Q, and C_r are constant for all $a\in A$, by Equation (3.17),

$$\pi^\infty(v) = \arg\min_{\pi\in\Pi}\sum_{a\in\pi}\delta(a)\sin\left(\theta(a)+\arctan C_r\right)^+ = \arg\min_{\pi\in\Pi}\sum_{a\in\pi} h'(a).$$

Proposition 3.4 demonstrates the convergence of the $\pi^g(v,l)$ to the $\pi^\infty(v)$ for a very large payload. On the other hand, by Proposition 3.2 the shortest path, π^{sp}, is the greenest path under the dynamic speed policy, i.e. $\pi^g(v^d,l)$, if $w+l=0$ and $v^{\min}(a)\le c_v\le v^{\max}(a)$ for all $a\in A$. The reason is that if $w+l=0$, the dynamic speed policy equals the static speed policy ($v^d=v^s$) by Proposition 3.1. Therefore, one can argue that $\pi^g(v,l)$ diverges from π^{sp} and converges to the $\pi^\infty(v)$ as the load increases. We will explore this idea in Section 3.4.4 through numerical experiments.

Finally, if the payload l satisfies Inequality (3.19), by Proposition 3.1, $v^d(a,l)$, for arc $a\in A$ can be computed as follows.

$$v^d(a) = \begin{cases} v^{\min}(a), & \text{if } \tan\theta(a)\ge -C_r \wedge c_v\le v^{\min}(a)\\ c_v, & \text{if } \tan\theta(a)\ge -C_r\wedge v^{\min}(a)<c_v\le v^{\max}(a)\\ v^{\max}(a), & \text{if } \tan\theta(a)<-C_r\vee\tan\theta(a)\ge -C_r\wedge v^{\max}(a)<c_v.\end{cases}$$

Consequently, if $v^{\min}(a)$ and $v^{\max}(a)$ are constant for all arcs $a\in A$ and if Π_d is non-empty, then $\pi^\infty(v^d)=\pi^\infty(v^s)$, by Proposition 3.4. Evidently, if Π_d is empty then Proposition 3.4 requires $\pi^\infty(v,l)$ to be independent of the speed policy v. As a result, $\pi^\infty(v^d)=\pi^\infty(v^s)$ if $v^{\min}(a)$ and $v^{\max}(a)$ are constant for all arcs $a\in A$.

3.4 Numerical Experiments

In this section, we explore the value of using elevation data to inform routing and speed decisions to reduce emissions over a comprehensive data set. Additionally, we explore the major drivers of CO_2 emissions reduction. We benchmark the greenest path (π^g) and dynamic speed policy (v^d) against the shortest path (π^{sp}) and the static speed policy (v^s). Note that the shortest path is also the greenest path under a dynamic speed policy (i.e. $\pi^g(v^d, l) = \pi^{sp}$) if the effect of road gradients is ignored, as shown in Proposition 3.2. We also study how the greenest path changes, $\pi^g(v, l)$, as the payload l increases and how the asymptotic greenest path $\pi^\infty(v)$ performs in terms of CO_2 emissions reduction and similarity to $\pi^g(v, l)$. In our numerical experiments the asymptotic greenest path under the dynamic speed policy, i.e. $\pi^\infty(v^d)$, and the one under the static speed policy, i.e. $\pi^\infty(v^s)$, are identical since $v^{\min}(a)$ and $v^{\max}(a)$ are constant for all $a \in A$ (see Section 3.3.5), i.e. $\pi^\infty = \pi^\infty(v^d) = \pi^\infty(v^s)$.

Given a pair of source and target vertices and a constant payload l, we compute the relative CO_2 emissions reduction of one policy in comparison with another. In particular we study the CO_2 reduction of using path-speed policy 2, $d_2 = (\pi_2, v_2, l)$, relative to path-speed policy 1, $d_1 = (\pi_1, v_1, l)$, ($\%\mathcal{E}_{d_1}^{d_2}$) to quantify the benefit of using the elevation data in CO_2 reduction. That is to say,

$$\%\mathcal{E}_{d_1}^{d_2} = 100 \cdot \frac{\mathcal{E}(\pi_1, v_1, l) - \mathcal{E}(\pi_2, v_2, l)}{\mathcal{E}(\pi_1, v_1, l)},$$

where $\mathcal{E}(\pi_i, v_i, l), i = 1, 2$ is the total CO_2 emissions as per Equation (3.12). If π_i, $i = 1, 2$, is a greenest path then $\pi_i = \pi^g(v_i, l)$. Similarly, we compute the relative distinction between the paths of policies π_1 and π_2 ($\%\delta_{\pi_1}^{\pi_2}$) weighted by distance, as follows.

$$\%\delta_{\pi_1}^{\pi_2} = 100 \cdot \sum_{a \in \pi_1 \setminus \pi_2} \delta(a) \,/\, \sum_{a \in \pi_1} \delta(a).$$

Table 3.2 briefly summarizes the ratios that we use in our comparative studies.

In Section 3.4.1, we outline the test-bed that we consider. This test-bed comprises 25 cities and all the ratios in Table3.2 are computed for instances in this test-bed. We present the results of our computations in Sections 3.4.2 through 3.4.5. Section 3.4.2 focuses on the results for the CO_2 emissions reduced by π^g and v^d relative to π^{sp} and v^s. In Section 3.4.3, we address the distinctions between π^g and π^{sp} and the effect of the speed policies v^s and v^d on the greenest path. In Section 3.4.4, we study the asymptotic greenest path π^∞ and explore the performance of π^∞ relative to the shortest path π^{sp} and the greenest path π^g in terms of CO_2 emissions reduction. In Sections 3.4.2 through 3.4.4, we elaborate on how payload affects our results. Finally, Section 3.4.5 concentrates on the major determinants of CO_2 emissions reduction and path alteration.

Table 3.2: List of ratios used in the comparative studies.

Ratio	Description
$\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^s,l)}$	Relative CO_2 emissions reduction by selecting the greenest path with the dynamic speed policy relative to the shortest path with the static speed policy given the load l.
$\%\mathcal{E}^{(\pi^g,v^s,l)}_{(\pi^{sp},v^s,l)}$	Relative CO_2 emissions reduction by selecting the greenest path with the static speed policy relative to the shortest path with the static speed policy given the load l.
$\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^d,l)}$	Relative CO_2 emissions reduction by selecting the greenest path with the dynamic speed policy relative to the shortest path with the dynamic speed policy given the load l.
$\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{g},v^s,l)}$	Relative CO_2 emissions reduction by selecting the greenest path with the dynamic speed policy relative to the greenest path with the static speed policy given the load l.
$\%\mathcal{E}^{(\pi^\infty,v^d,l)}_{(\pi^{sp},v^d,l)}$	Relative CO_2 emissions reduction by selecting the asymptotic greenest path with the dynamic speed policy relative to the shortest path with the dynamic speed policy given the load l.
$\%\mathcal{E}^{(\pi^\infty,v^s,l)}_{(\pi^{sp},v^s,l)}$	Relative CO_2 emissions reduction by selecting the asymptotic greenest path with the static speed policy relative to the shortest path with the static speed policy given the load l.
$\%\mathcal{E}^{(\pi^\infty,v^d,l)}_{(\pi^{g},v^d,l)}$	Relative CO_2 emissions reduction by selecting the asymptotic greenest path with the dynamic speed policy relative to the greenest path with the dynamic speed policy given the load l.
$\%\mathcal{E}^{(\pi^\infty,v^s,l)}_{(\pi^{g},v^s,l)}$	Relative CO_2 emissions reduction by selecting the asymptotic greenest path with the static speed policy relative to the greenest path with the static speed policy given the load l.
$\%\delta^{\pi^g(v^d,l)}_{\pi^{sp}}$	Ratio of the length of the shortest path that is not shared with the greenest path under the dynamic speed policy given the load l.
$\%\delta^{\pi^g(v^s,l)}_{\pi^{sp}}$	Ratio of the length of the shortest path that is not shared with the greenest path under the static speed policy given the load l.
$\%\delta^{\pi^g(v^s,l)}_{\pi^g(v^d,l)}$	Ratio of the length of the greenest path under the dynamic speed policy that is not shared with the greenest path under the static speed policy given the load l.
$\%\delta^{\pi^\infty}_{\pi^g(v^d,l)}$	Ratio of the length of the greenest path under the dynamic speed policy that is not shared with the asymptotic greenest path given the load l.
$\%\delta^{\pi^\infty}_{\pi^g(v^s,l)}$	Ratio of the length of the greenest path under the static speed policy that is not shared with the asymptotic greenest path given the load l.

3.4.1 Data and test-bed

We consider the 25 cities shown in Table 3.3 and three truck types, namely heavy-duty diesel (HDD), medium-duty diesel (MDD), and light-duty diesel (LDD) for which we utilise the typical parameters as found in Table 3.4 of Koc et al. (2014).

Table 3.3: List of cities and sample sizes.

City	Country	Number of S-T pairs
Amsterdam	Netherlands	146842
Ankara	Turkey	114675
Athens	Greece	114687
Barcelona	Spain	100649
Canberra	Australia	131565
Geneva	Switzerland	125605
Guadalajara	Mexico	126600
Guangzhou	China	124940
Huston	US	148548
Istanbul	Turkey	94999
Johannesburg	South Africa	130054
Lima	Peru	119174
Los Angeles	US	142452
Luxembourg	Luxembourg	136201
Madrid	Spain	117686
Mexico City	Mexico	93076
Monterrey	Mexico	128619
Mumbai	India	143116
New York	US	146875
Rome	Italy	124452
San Francisco	US	93504
Santiago	Chile	142212
Shiraz	Iran	90935
Tehran	Iran	110104
Tel Aviv	Israel	131066

Table 3.4: Truck parameters.

Parameter	HDD	MDD	LDD
w	14000	5500	3500
L	26000	12500	4000
k	0.15	0.2	0.25
N	30	36.67	38.34
D	10.5	6.9	4.5
η	0.45	0.45	0.45
η_{tf}	0.45	0.45	0.45
ξ	1	1	1
κ	44	44	44
ϕ	737	737	737
C_d	0.9	0.7	0.6
ρ	1.2041	1.2041	1.2041
A	10	8	7
g	9.81	9.81	9.81
C_r	0.01	0.01	0.01

We use OpenStreetMap's database (OpenStreetMap contributors, 2017) to obtain the information of a 2D road network including all vertices within a 20 km radius around a manually selected point for each city. We only use roads that the database designates as public and driveable (OpenStreetMap contributors, 2022). We only consider arcs with a gradient ranging from -10% to 10% (i.e. $[-5.71°, 5.71°]$) so that gradients are in line with the implicit assumptions of the modified emissions model (3.5). We retrieve the elevation (height above sea level) of the vertices from the U.S. Geological Survey (2000)'s SRTM 1 Arc-Second Global data sets. We consider payloads of 30%, 40%, 50%, 60%, 70%, and 80% of the maximum capacity for each truck type. For all arcs the $v^{\max} = 90$ km/h and $v^{\min} = 20$ km/h.

We select several unique pairs of source and target vertices uniformly at random for each city. We make sure that the vertices in each pair are non-identical and connected. The number of selected pairs of vertices (sample size) for each city is presented in Table 3.3.

The Dijkstra algorithm (Dijkstra, 1959) is used to solve the shortest path and the greenest path problems. We use the arcs' distance $\delta(a)$, $a \in A$, to compute the shortest path π^{sp}. We consider two speed policies, namely dynamic speed policy, v^d, and static

speed policy, v^s to calculate the the CO_2 emissions, $e_a(v, l)$, for all arcs. Then we use the calculated $e_a(v, l)$ to compute the greenest paths ($\pi^g(v^d, l)$ and $\pi^g(v^s, l)$). We use the Dijkstra algorithm to compute π^∞ as per Proposition 3.4.

We consider two sets of ratios, as shown in Table 3.2, to compare the different path (π^{sp}, π^g, and π^∞) and speed (v^s and v^d) policies. The first group of ratios measure the relative CO_2 emissions reduction and the second group measures the geometrical distinctions between the paths. We compute the ratios for a full factorial combination of trucks and payloads traversing all samples, resulting in a total of more than 58.5 million path selection instances with a total shortest distance of more than 1.27 billion km. Evidently, it is hardly possible to determine CO_2 emissions experimentally by letting trucks drive 1.27 billion km as the approaches of Boriboonsomsin et al. (2012) and Scora et al. (2015). The confidence intervals of any estimate reported later are negligibly small due to the large sample size. Considering the large test-bed, we notice that the distribution and the sample mean of the ratios varies between different cities. For a given city, we use the overbar to denote the average of a ratio across all instances within a city. For instance, $\overline{\%\mathcal{E}}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^s,l)}$ for a city represents the sample mean of $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^s,l)}$ for that city.

3.4.2 Results: CO_2 Emissions Reduction by Greenest Path and Dynamic Speed Policy

In this section, we consider the payload as a percentage of the truck's maximum carrying capacity rather than the payload in kilograms, to make the notations simpler. For instance, $l = 60\%$ indicates that the payload equals 60% of the maximum capacity of the truck. Since the payload varies the results, we use $l = 60\%$ as our base case to maintain consistency.

Figures 3.1 through 3.4 visualize the empirical distribution of CO_2 emissions reduction ratios for the base case instances. We present the distributions separately for each truck type and each city. The sample size of each empirical distribution is listed in Table 3.3.

Figure 3.1 shows that $\overline{\%\mathcal{E}}^{(\pi^g,v^d,60\%)}_{(\pi^{sp},v^s,60\%)}$ lies between 4.11% and 10.15% for HDD trucks across all cities except Amsterdam. Figure 3.1 also shows that $\%\mathcal{E}^{(\pi^g,v^d,60\%)}_{(\pi^{sp},v^s,60\%)}$ decreases in truck class such that $\overline{\%\mathcal{E}}^{(\pi^g,v^d,60\%)}_{(\pi^{sp},v^s,60\%)}$ ranges from 3.27% to 8.65% for MDD, and from 2.41% to 7.00% for LDD trucks, in the same cities. Amsterdam, a known flat city, is the lone exception, but even here $\overline{\%\mathcal{E}}^{(\pi^g,v^d,60\%)}_{(\pi^{sp},v^s,60\%)}$ are 2.44%, 1.78%, and 1.19%, respectively, showing that it is possible to use significantly more fuel-efficient paths. The distribution of $\%\mathcal{E}^{(\pi^g,v^d,60\%)}_{(\pi^{sp},v^s,60\%)}$, on the other hand, sheds more light on the potential CO_2 emissions reduction by using the greenest path with a dynamic speed policy, $\pi^g(v^d, 60\%)$. In Los Angeles, for instance, 25% of cases have a $\%\mathcal{E}^{(\pi^g,v^d,60\%)}_{(\pi^{sp},v^s,60\%)}$ of at least 13.57% for HDD, 10.14% for MDD, and 7.00% for LDD trucks. It may appear that these effects are larger than the numerical results of earlier studies, for instance Scora et al. (2015); Schröder

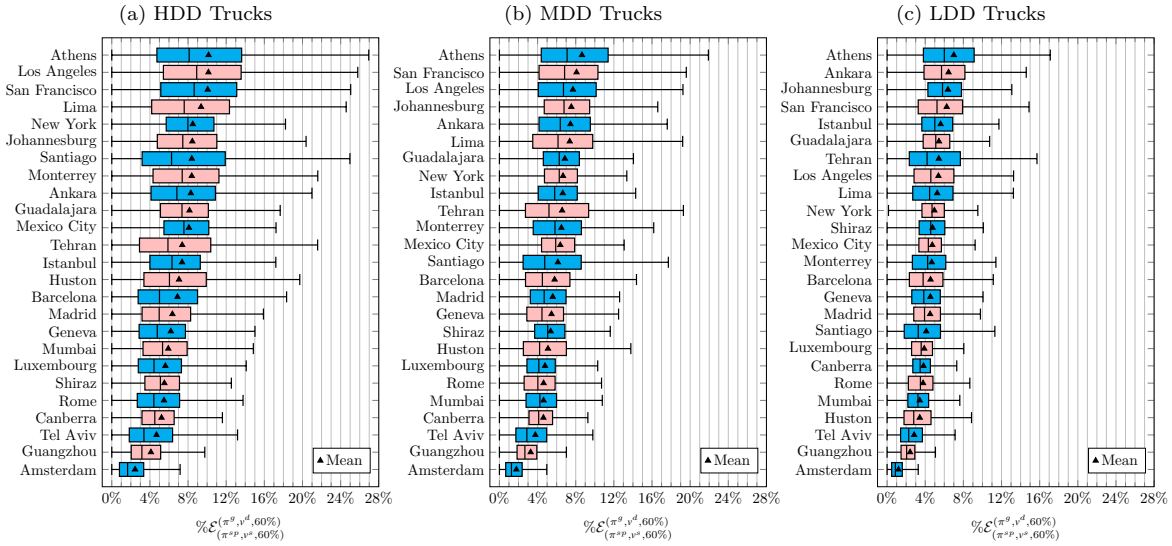

Figure 3.1: Relative CO_2 emissions reduction by selecting $(\pi^g, v^d, 60\%)$ rather than $(\pi^{sp}, v^s, 60\%)$.

and Cabral (2019) and Brunner et al. (2021). We submit that this is due to the long tail of the distribution of fuel savings which is found only with a sufficiently large sample.

To discern the individual effect of road gradient on CO_2 emissions reduction, we fix a speed policy $v \in \{v^s, v^d\}$ and then take into account the CO_2 emissions reduction by traveling along the greenest path $\pi^g(v, l)$ rather than the shortest path π^{sp}. We consider two ratios $\%\mathcal{E}^{(\pi^g, v^s, l)}_{(\pi^{sp}, v^s, l)}$ and $\%\mathcal{E}^{(\pi^g, v^d, l)}_{(\pi^{sp}, v^d, l)}$ to assess this effect. Figures 3.2 and 3.3 indicate the distribution and mean of these ratios for base case instances in different cities. Figure 3.2 demonstrates that the selection of $\pi^g(v^s, 60\%)$ rather than π^{sp} can reduce, on average, 1.76% to 8.15% of the CO_2 emissions, if $l = 60\%$ and v^s is decided. As explained before, this CO_2 emissions reduction capacity is lower for the MDD and LDD trucks, yet remains substantive. In the case of a dynamic speed policy v^d, the statistics, i.e. $\%\mathcal{E}^{(\pi^g, v^d, l)}_{(\pi^{sp}, v^d, l)}$, remain close to that of v^s, i.e. $\%\mathcal{E}^{(\pi^g, v^s, l)}_{(\pi^{sp}, v^s, l)}$, but they are slightly smaller. This implies that regardless of speed, taking into account the road gradient results in significant reductions in CO2 emissions.

Next, to investigate the effect of speed policies on fuel-efficient paths and CO2 emissions reduction, we appraise the carbon reduction by modifying the policy from (π^g, v^s, l) to (π^g, v^d, l) for the same truck, i.e. $\%\mathcal{E}^{(\pi^g, v^d, l)}_{(\pi^g, v^s, l)}$. Figure 3.4 shows that for most cities, $\overline{\%\mathcal{E}}^{(\pi^g, v^d, 60\%)}_{(\pi^g, v^s, 60\%)}$ is between 2% and 4%, and in all cases the estimates do not depend on the truck type.

We contrast $\%\mathcal{E}^{(\pi^g, v^d, 60\%)}_{(\pi^g, v^s, 60\%)}$ and $\%\mathcal{E}^{(\pi^g, v^d, 60\%)}_{(\pi^{sp}, v^d, 60\%)}$, as shown in Figure 3.5, in order to experimentally evaluate the relative efficacy of the greenest path and speed optimization in reducing CO2 emissions for each type of vehicle (i.e. HDD, MDD, and LDD). For HDD trucks, the road gradient is more crucial than the dynamic speed policy, whereas the dynamic speed policy has more impact for LDD trucks. The greenest path and dynamic

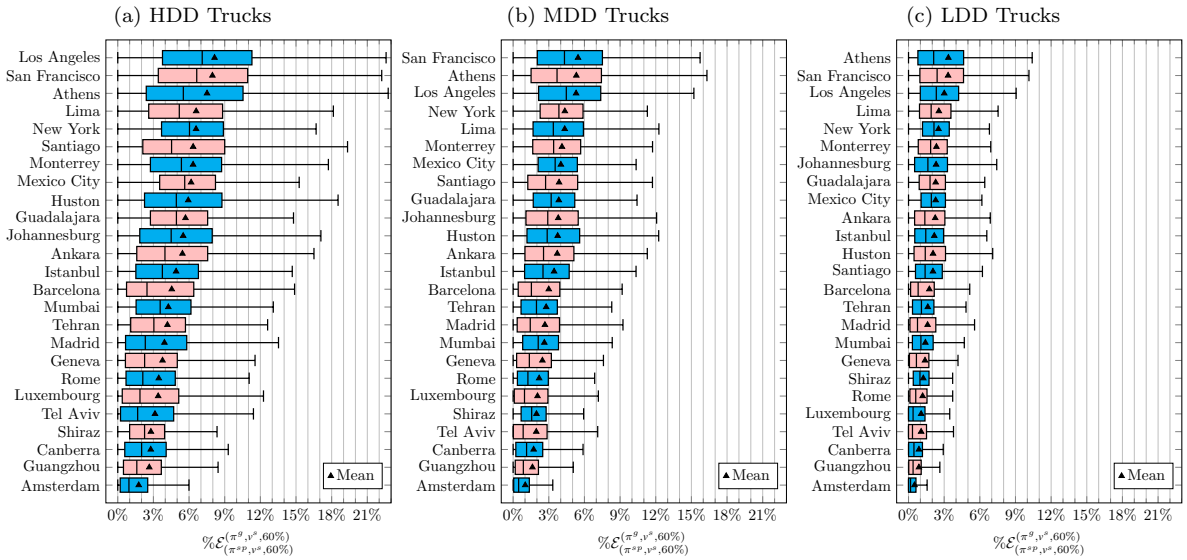

Figure 3.2: Relative CO_2 emissions reduction by selecting $(\pi^g, v^s, 60\%)$ rather than $(\pi^{sp}, v^s, 60\%)$.

speed policy can bring down the CO_2 emissions of MDD trucks to the same extent.

To analyze the effect of payload on CO_2 emissions reduction, we vary payload ratio for the base case instances (30%, 40%, 50%, 70%, and 80%) and repeat the same experiments. Figures 3.6 through 3.9 present the distributions of the sample mean of the relative CO_2 emissions reduction ratios over the 25 cities, where the sample size of each box plot is 25. The figures also present the alteration of the distributions as the payload increases. These results show that, on average, $\overline{\%\mathcal{E}}^{(\pi^g, v^d, l)}_{(\pi^{sp}, v^s, l)}$ (Figure 3.6), $\overline{\%\mathcal{E}}^{(\pi^g, v^s, l)}_{(\pi^{sp}, v^s, l)}$ (Figure 3.7), and $\overline{\%\mathcal{E}}^{(\pi^g, v^d, l)}_{(\pi^{sp}, v^d, l)}$ (Figure 3.8) are non-decreasing in payload. However, all graphs are concave, so that the growth rate of $\overline{\%\mathcal{E}}$ decreases in payload. In many cities, this phenomenon results

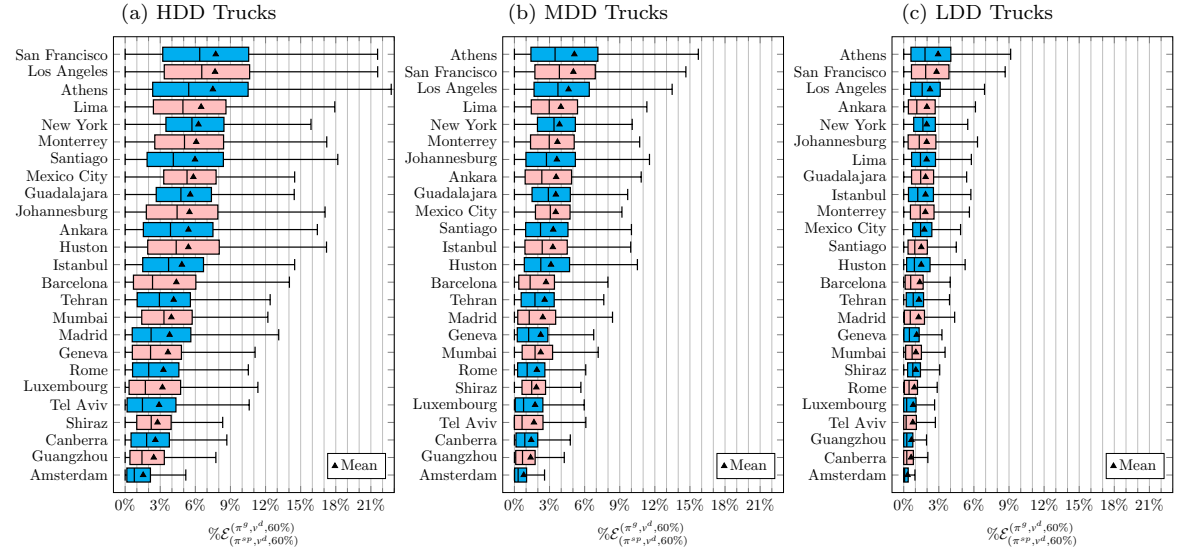

Figure 3.3: Relative CO_2 emissions reduction by selecting $(\pi^g, v^d, 60\%)$ rather than $(\pi^{sp}, v^d, 60\%)$.

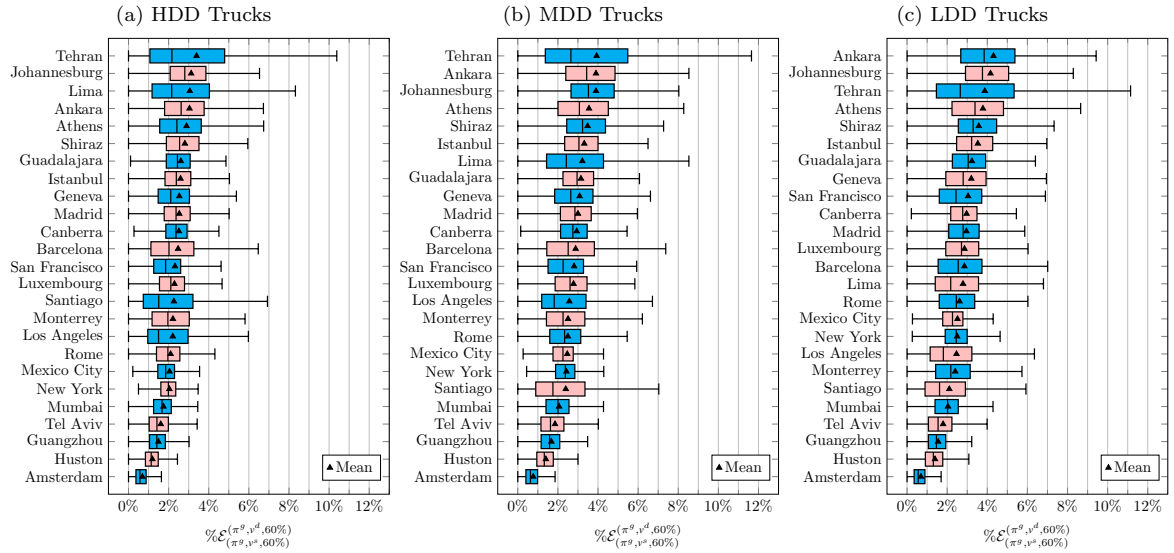

Figure 3.4: Relative CO_2 emissions reduction by selecting $(\pi^g, v^d, 60\%)$ rather than $(\pi^g, v^s, 60\%)$.

in a slow increase, and in one case (Shiraz) slight decrease of $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^s,l)}$ for HDD trucks. The same concave pattern takes place for $\overline{\%\mathcal{E}}^{(\pi^g,v^d,l)}_{(\pi^g,v^s,l)}$ (Figure 3.9) with the exception that the maxima of the concave functions are typically in the MDD or LDD regions. This result can explain the close range of $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^g,v^s,l)}$ across different truck types as shown in Figure 3.4.

3.4.3 Results: Paths of the $\pi^g(v^d, l)$, $\pi^g(v^s, l)$, and π^{sp}

The differences between the greenest path and the shortest path have been covered in earlier sections, along with an analysis of the impact of speed and road gradient. Although our findings indicate significant differences in fuel consumption and CO2 emissions, it is important to consider whether the shortest path's trajectory differs significantly from the

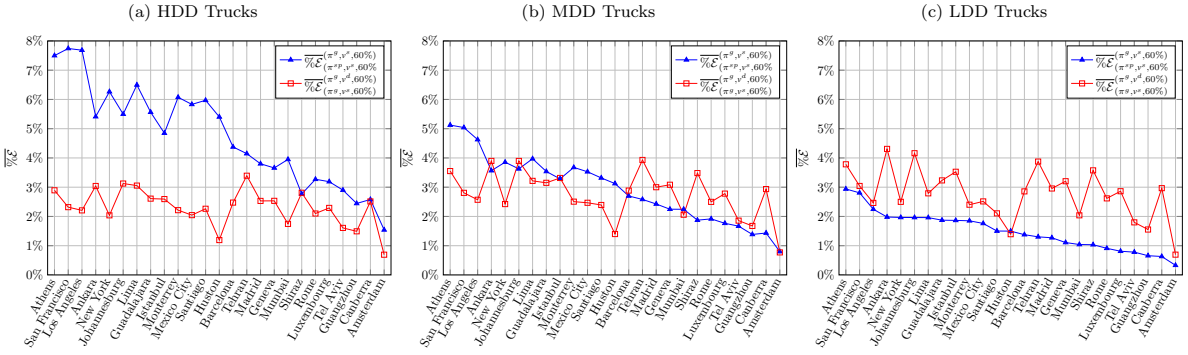

Figure 3.5: Comparison of the CO_2 emissions reduction potential: sole π^g $(\overline{\%\mathcal{E}}^{(\pi^g,v^s,60\%)}_{(\pi^{sp},v^s,60\%)})$ vs. sole v^d $(\overline{\%\mathcal{E}}^{(\pi^g,v^d,60\%)}_{(\pi^g,v^s,60\%)})$ across different cities.

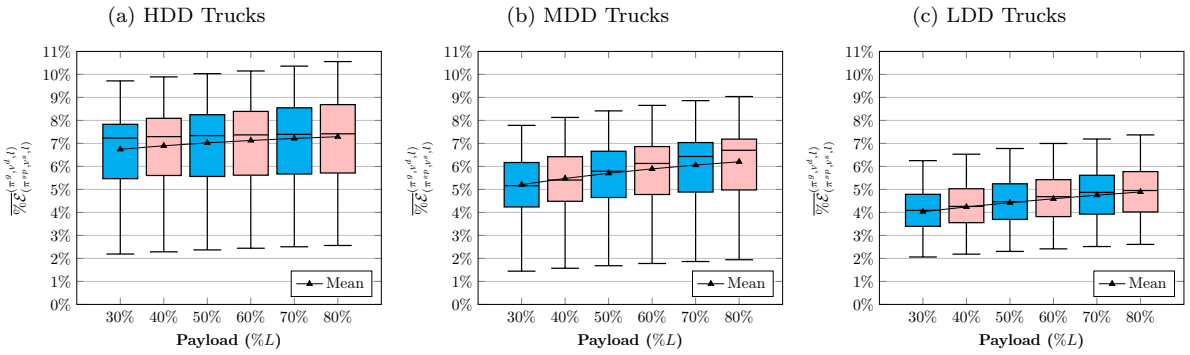

Figure 3.6: Effect of payload on $\overline{\%\mathcal{E}}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^s,l)}$ across 25 cities.

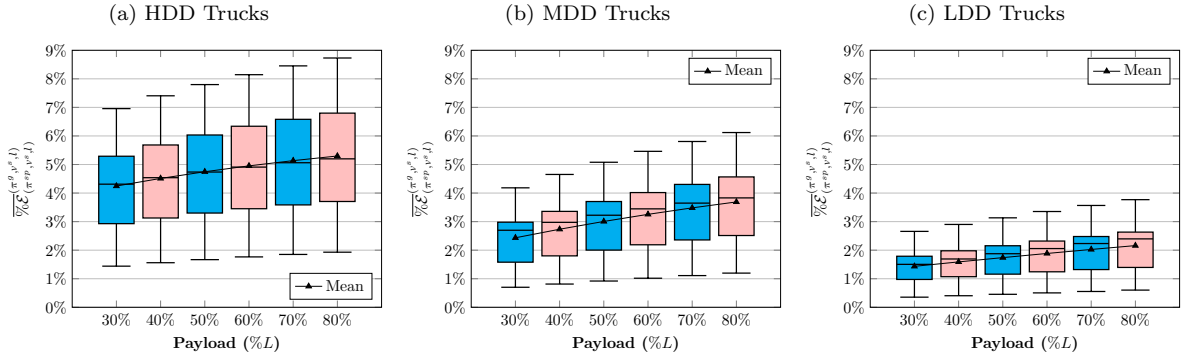

Figure 3.7: Effect of payload on $\overline{\%\mathcal{E}}^{(\pi^g,v^s,l)}_{(\pi^{sp},v^s,l)}$ across 25 cities.

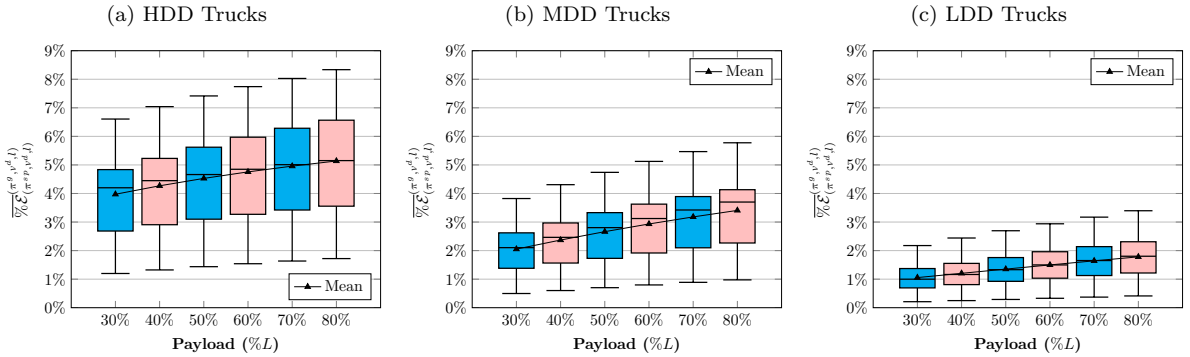

Figure 3.8: Effect of payload on $\overline{\%\mathcal{E}}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^d,l)}$ across 25 cities.

trajectory produced by the greenest path.

To illustrate this difference, we consider a LDD truck that delivers cargo weighing 60% of its maximum capacity from point A to B within a district of Los Angeles, see Figure 3.10. Figure 3.10 displays the greenest paths ($\pi^g(v^d,l)$ and $\pi^g(v^s,l)$) and the shortest path path (π^{sp}) on the map, and Figure 3.11 shows the elevation of the vertices and total CO_2 emissions for different path and speed choices. In this instance, π^{sp} differs significantly from $\pi^g(v^d,l)$ and $\pi^g(v^s,l)$, whereas the two greenest paths share a number of arcs. In this section, we examine whether such an observation is common throughout

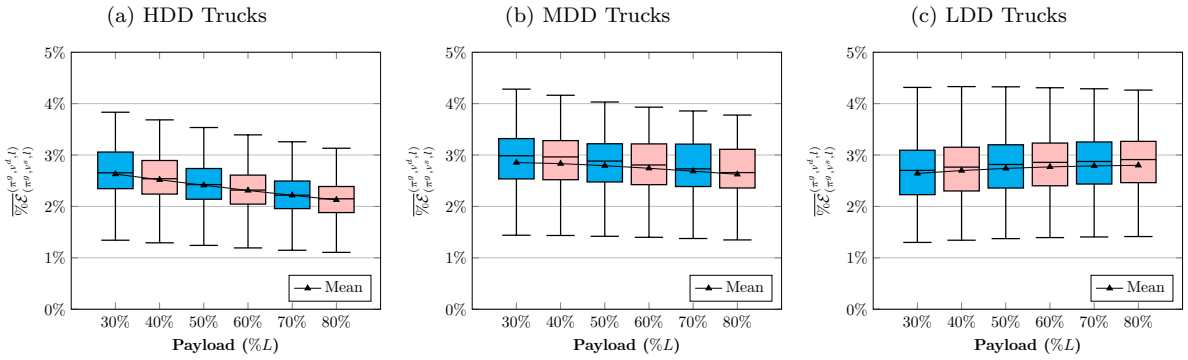

Figure 3.9: Effect of payload on $\overline{\%\mathcal{E}}^{(\pi^g,v^d,l)}_{(\pi^g,v^s,l)}$ across 25 cities.

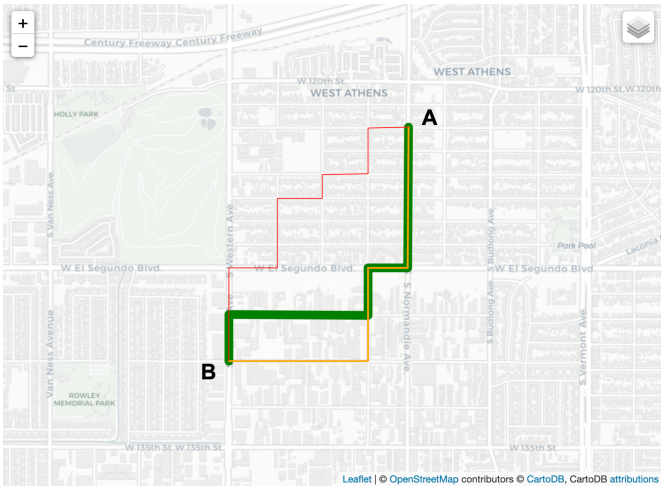

Figure 3.10: Example of $\pi^g(v^d, 60\%)$, $\pi^g(v^s, 60\%)$, π^{sp} (Green/Bold: $\pi^g(v^d, 60\%)$, Orange/Medium: $\pi^g(v^s, 60\%)$, and Red/Thin: π^{sp}).

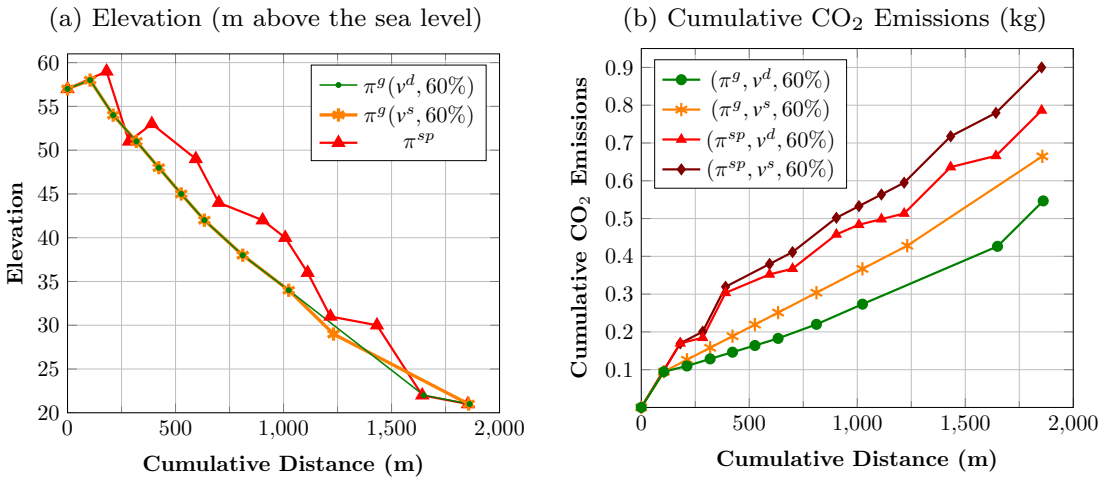

Figure 3.11: Elevation of the vertices and CO_2 emissions along the paths as per Figure 3.10.

our test-bed.

Figures 3.12 through 3.14 encapsulate the distribution and sample mean of $\%\delta^{(\pi^g,v^s,60\%)}_{(\pi^g,v^d,60\%)}$,

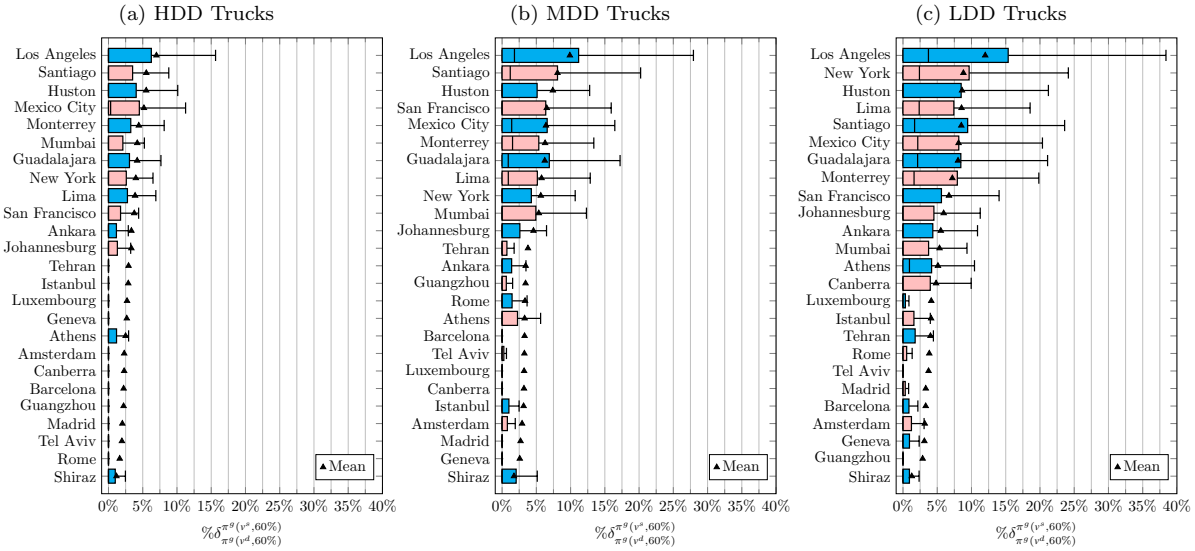

Figure 3.12: Ratio of the length of $\pi^g(\nu^d, 60\%)$ that is not shared with $\pi^g(\nu^s, 60\%)$.

$\%\delta_{\pi^{sp}}^{\pi^g(\nu^d,60\%)}$, and $\%\delta_{\pi^{sp}}^{\pi^g(\nu^s,60\%)}$ for the base cases. Figure 3.12 shows that the average difference of $\pi^g(\nu^s, 60\%)$ and $\pi^g(\nu^d, 60\%)$ is between 1.16% and 12.01% across the cities. In fact, the quartiles of $\%\delta_{(\pi^g,\nu^d,60\%)}^{(\pi^g,\nu^s,60\%)}$ show that for the most part $\pi^g(\nu^d, 60\%)$ are quite similar to $\pi^g(\nu^s, 60\%)$. In other words, in a majority of instances, the greenest path is independent of the speed policy. Additionally, for heavier trucks the greenest path is less likely to vary as a result of speed optimization. Figures 3.13 and 3.14 show that the distinction between the shortest and the greenest paths, i.e. $\%\delta_{\pi^{sp}}^{\pi^g(\nu^d,60\%)}$ and $\%\delta_{\pi^{sp}}^{\pi^g(\nu^s,60\%)}$, are conspicuously larger than the distinction between the greenest paths, i.e. $\%\delta_{\pi^g(\nu^d,60\%)}^{\pi^g(\nu^s,60\%)}$. This gap intensifies with heavier truck classes.

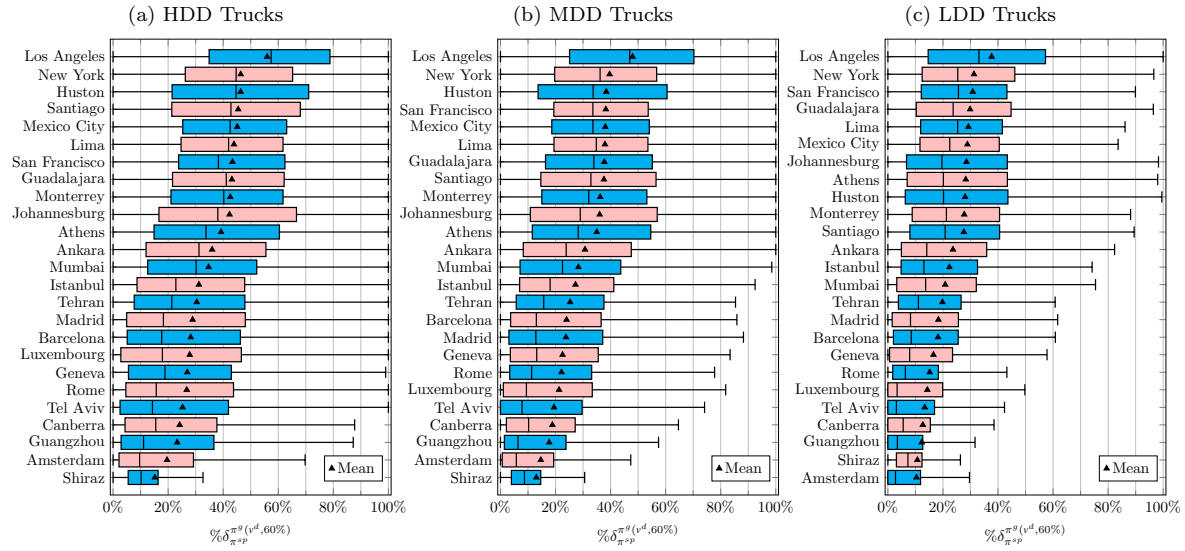

Figure 3.13: Ratio of the length of π^{sp} that is not shared with $\pi^g(\nu^d, 60\%)$.

To expand our understanding of the payload's influence on paths, we study whether

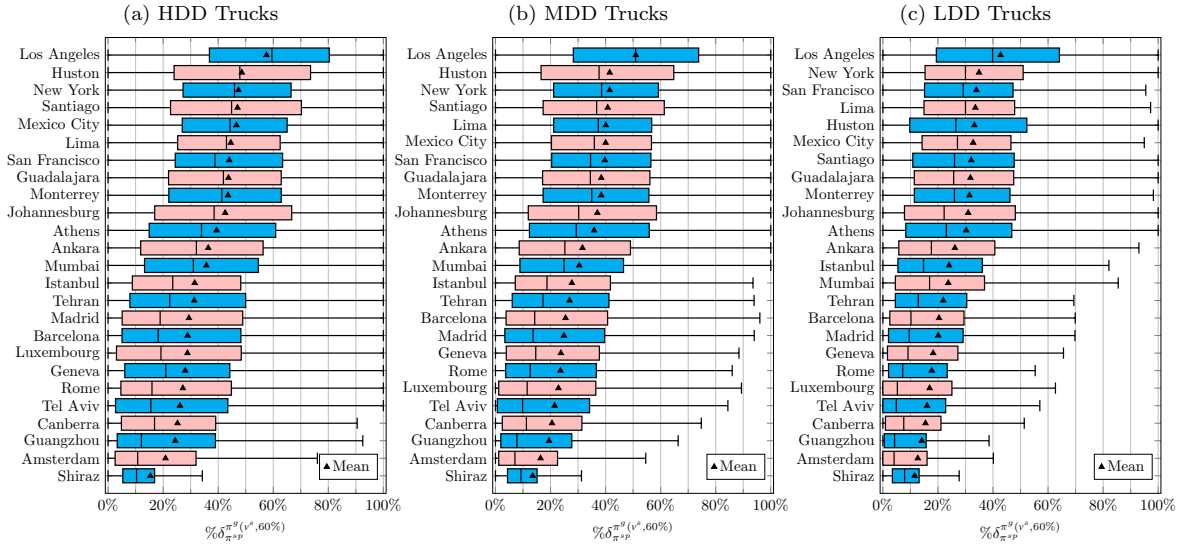

Figure 3.14: Ratio of the length of π^{sp} that is not shared with $\pi^g(v^s, 60\%)$.

$\pi^g(v^s, l)$ and $\pi^g(v^d, l)$ converge to each other and diverge from π^{sp} as the payload increases. Figures 3.15 and 3.16 demonstrate that both $\overline{\%\delta}_{\pi^{sp}}^{\pi^g(v^d,l)}$ and $\overline{\%\delta}_{\pi^{sp}}^{\pi^g(v^s,l)}$ are non-decreasing in payload in contrast to $\overline{\%\delta}_{(\pi^g,v^d,l)}^{(\pi^g,v^s,l)}$ which is mostly decreasing, as indicated by Figure 3.17. Note that, $\overline{\%\delta}_{\pi^{sp}}^{\pi^g(v^s,l)}$ is always higher than $\overline{\%\delta}_{\pi^{sp}}^{\pi^g(v^d,l)}$ since the more efficient dynamic speed policy of $\pi^g(v^d, l)$ usually allows for a shorter (and faster) path relative to $\pi^g(v^s, l)$. However, increase in payload erodes the impact of speed policy.

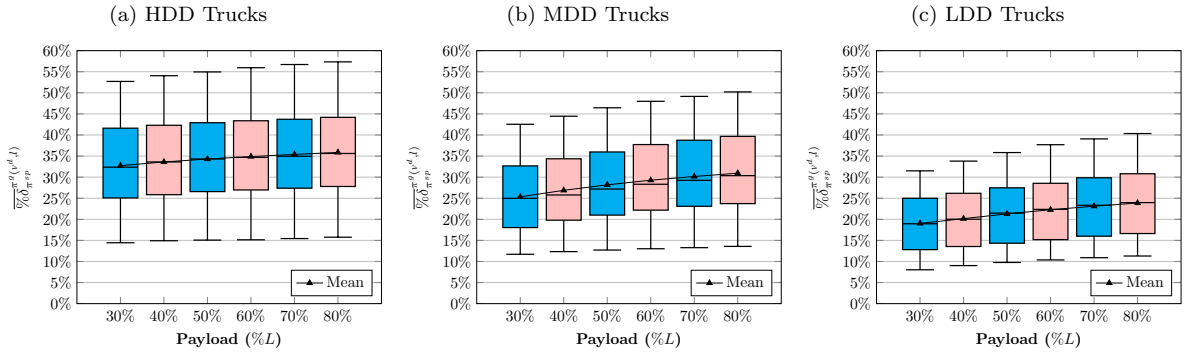

Figure 3.15: Effect of payload on $\overline{\%\delta}_{\pi^{sp}}^{\pi^g(v^d,l)}$ across 25 cities.

3.4.4 Results: Performance of the Asymptotic Greenest Paths

The greenest path converges to the asymptotic greenest path for arbitrarily large payloads as shown in Section 3.3.5. In this section, we study the performance of the asymptotic greenest path relative to the shortest path and the greenest path. Then we study the rate of convergence of the greenest path to the asymptotic greenest path for the dynamic

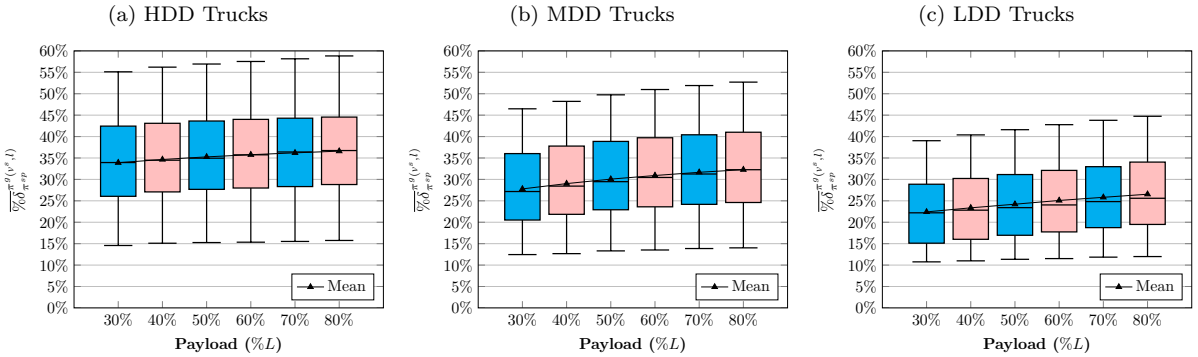

Figure 3.16: Effect of payload on $\overline{\%\delta}^{\pi^g(v^s,l)}_{\pi^{sp}}$ across 25 cities.

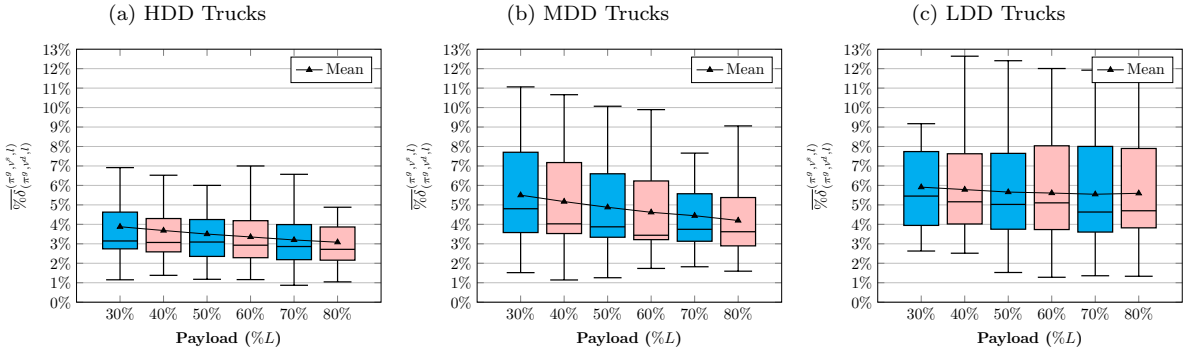

Figure 3.17: Effect of payload on $\overline{\%\delta}^{(\pi^g,v^s,l)}_{(\pi^g,v^d,l)}$ across 25 cities.

speed policy. In Appendix 3.B we study these things under the static speed policy. Figures 3.18 and 3.19 show that the distribution of the CO_2 emissions reduction of π^∞ relative to π^{sp} and $\pi^g(v^d, 60\%)$ for different cities. Similar to Section 3.4.2, we present the results for the base cases (60% payload ratio). Figure 3.18 shows that for the most

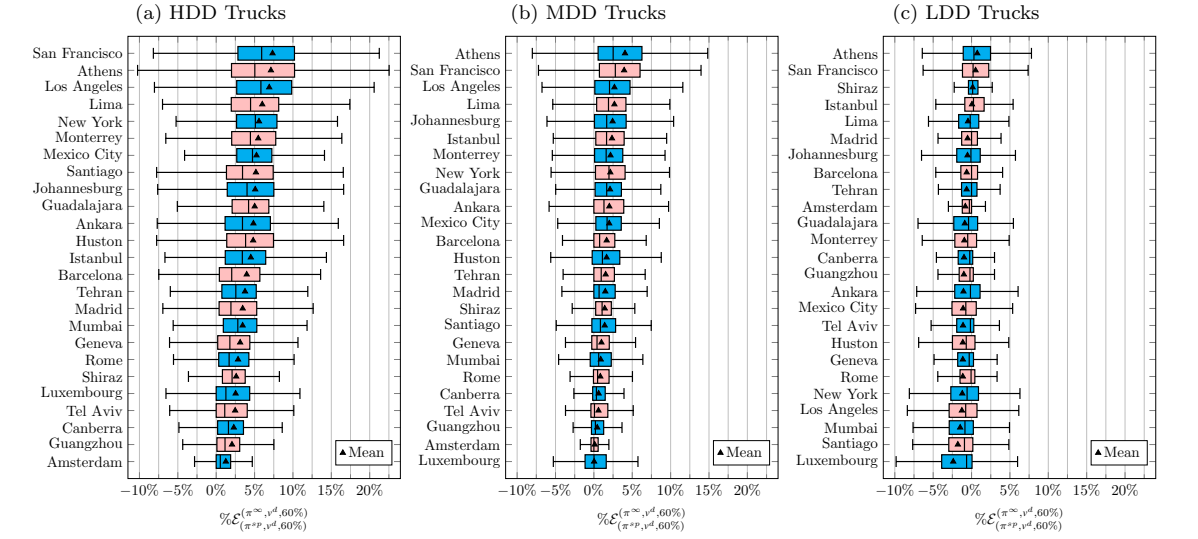

Figure 3.18: Relative CO_2 emissions reduction by selecting $(\pi^\infty, v^d, 60\%)$ rather than $(\pi^{sp}, v^d, 60\%)$.

part an LDD truck emits slightly more CO_2 if it traverses π^{∞} instead of π^{sp} in 18 cities. Whereas, the π^{∞} is greener than the π^{sp} for MDD and HDD trucks in more than 50% of the instances in all cities. The CO_2 emissions reduction of $(\pi^{\infty}, v^d, 60\%)$ relative to

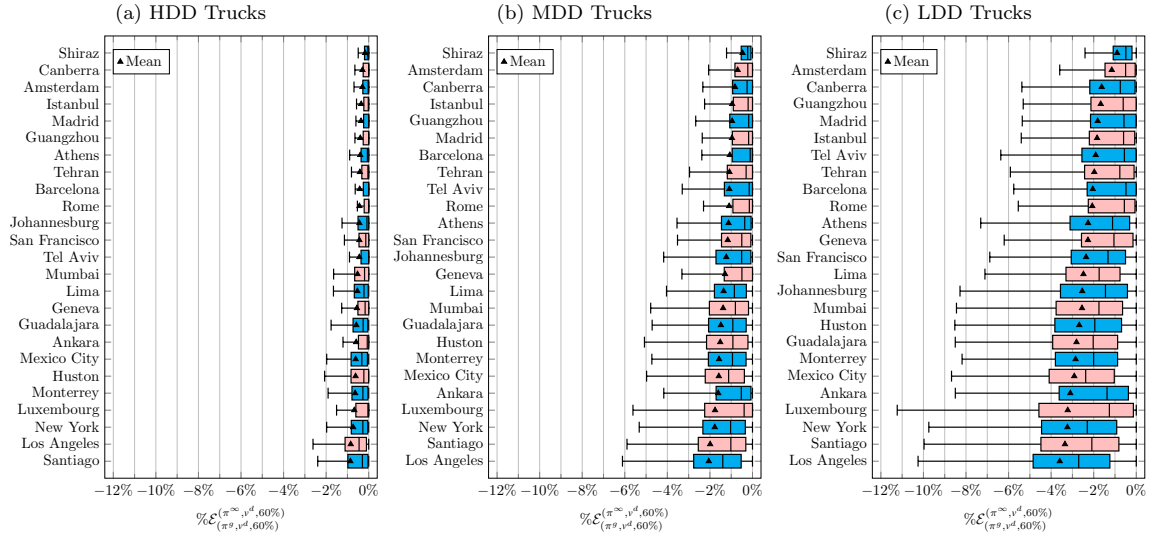

Figure 3.19: Relative CO_2 emissions reduction by selecting $(\pi^{\infty}, v^d, 60\%)$ rather than $(\pi^{g}, v^d, 60\%)$.

$(\pi^{g}, v^d, 60\%)$, i.e. $\%\mathcal{E}^{(\pi^{\infty}, v^d, 60\%)}_{(\pi^{g}, v^d, 60\%)}$, is consistent with this observation. Figure 3.19 shows that the median of extra CO_2 emissions along π^{∞} compared to the $\pi^{g}(v^d, 60\%)$ ranges from 0.48% to 2.70% for LDD trucks. This range decreases to between 0.11% and 1.40% for MDD trucks, and 0% and 0.45% for HDD trucks. Figure 3.20 shows the distribution of the sample mean of $\%\mathcal{E}^{(\pi^{\infty}, v^d, l)}_{(\pi^{sp}, v^d, l)}$ across the 25 cities for various payload ratios, i.e. $\overline{\%\mathcal{E}}^{(\pi^{\infty}, v^d, l)}_{(\pi^{sp}, v^d, l)}$. Correspondingly, Figure 3.21 presents $\overline{\%\mathcal{E}}^{(\pi^{\infty}, v^d, l)}_{(\pi^{g}, v^d, l)}$. The two figures show

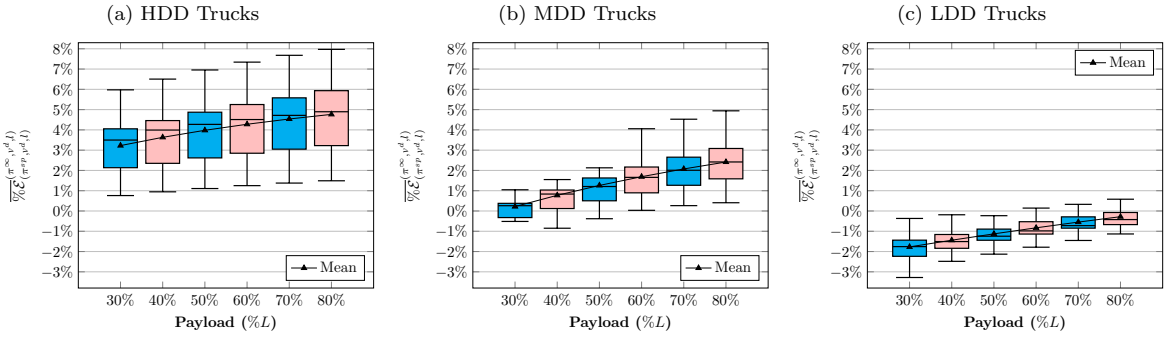

Figure 3.20: Effect of payload on $\overline{\%\mathcal{E}}^{(\pi^{\infty}, v^d, l)}_{(\pi^{sp}, v^d, l)}$ across 25 cities.

that the average CO_2 emissions along π^{∞} relative to π^{sp} and $\pi^{g}(v^d, l)$ is non-increasing in load, l. Evidently, π^{sp} outperforms π^{∞} in terms of average CO_2 emissions for LDD trucks with any payload ratio. Whereas, π^{∞} is on average greener than π^{sp} for MDD and HDD truck types for all payload ratios. The mean excess CO_2 emissions of π^{∞} relative to $\pi^{g}(v^d, l)$ is less than 1%, 2%, and 3% for HDD, MDD, and LDD trucks, respectively.

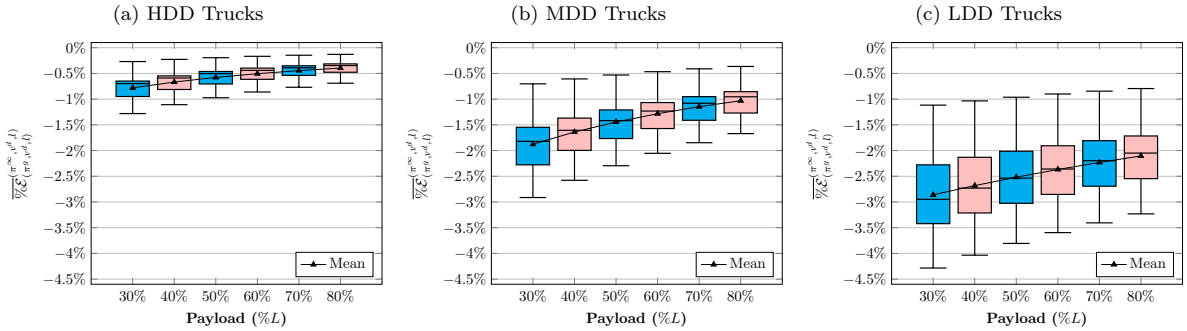

Figure 3.21: Effect of payload on $\overline{\%\mathcal{E}}^{(\pi^\infty,v^d,l)}_{(\pi^g,v^d,l)}$ across 25 cities.

The median of the difference between $\pi^g(v^d,l)$ and π^∞, i.e. $\%\delta^{\pi^\infty}_{\pi^g(v^d,60\%)}$, varies between 4.97% and 49.14% for the LDD trucks in base cases as Figure 3.22 shows. However, the similarity increases in MDD and HDD truck types as the median $\%\delta^{\pi^\infty}_{\pi^g(v^d,60\%)}$ ranges from 2.84% to 34.81% for MDD trucks and 0% to 18.14% for HDD trucks. Moreover,

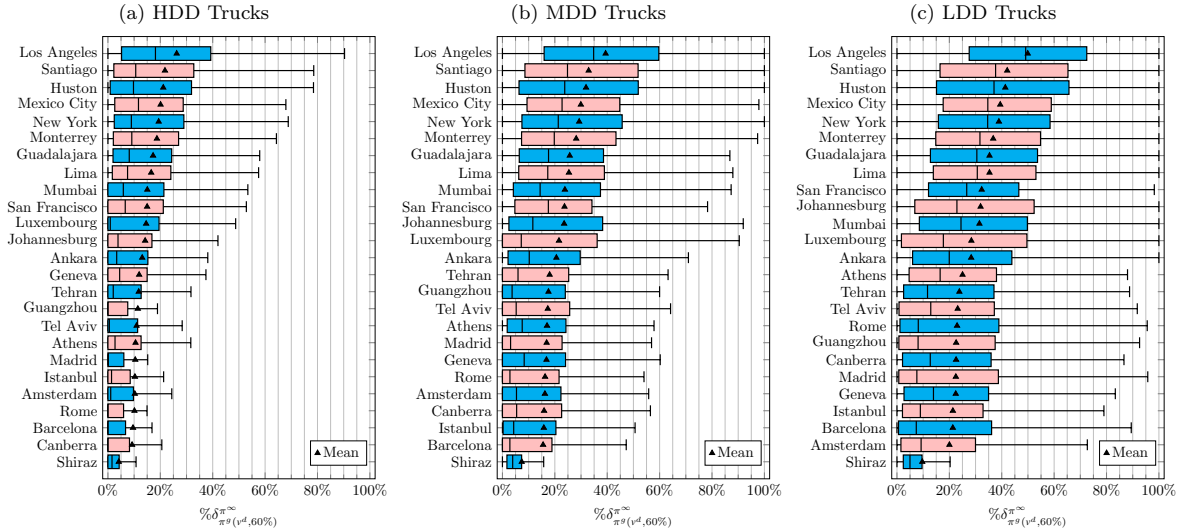

Figure 3.22: Ratio of the length of $\pi^g(v^d,60\%)$ that is not shared with π^∞.

the difference between the $\pi_{(\pi^g,v^d,l)}$ and π^∞ reduces in the payload ratio in all truck types. Figures 3.22 and 3.23 show that $\pi^g(v^d,l)$ converges to π^∞ in the payload ratio as established in Proposition 3.4. These results confirm that $\pi^g(v^d,l)$ diverges from π^{sp} (Figures 3.13 and 3.15) and converges to π^∞ (Figures 3.22 and 3.23) as the payload (and curb weight) increases.

3.4.5 Results: Main Determinants

In this section, we address the major determinants of the CO_2 emissions reduction and path alteration. We consider the following input features: city, truck type, payload, the

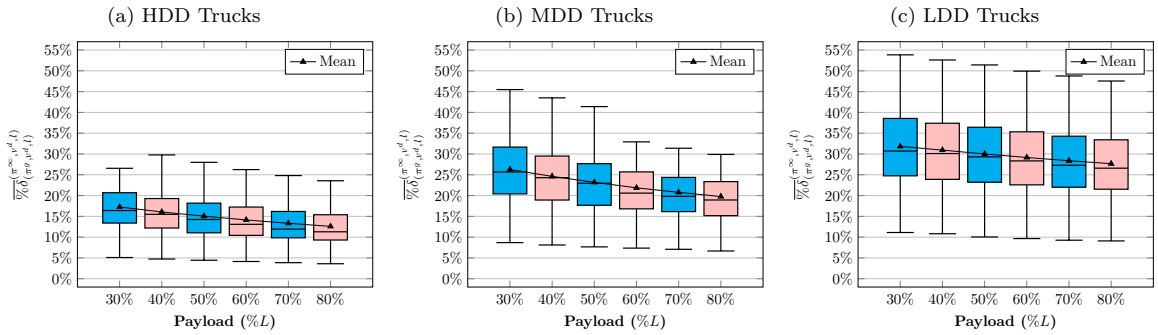

Figure 3.23: Effect of payload on $\overline{\%\delta}^{(\pi^\infty,v^d,l)}_{(\pi^g,v^d,l)}$ across 25 cities.

elevation difference of source and target (Δh), the distance of the shortest path (δ^{sp}) and the standard deviation of the gradients along the shortest path ($\sigma^{sp}(\theta)$). The latter characterizes the hilliness of the shortest path. All of these features can be efficiently computed. We use linear regression accompanied by the analysis of variance (ANOVA) to regress these features against seven responses, namely $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^s,l)}$, $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^d,l)}$, $\%\mathcal{E}^{(\pi^g,v^s,l)}_{(\pi^{sp},v^s,l)}$, $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^g,v^s,l)}$, $\%\delta^{\pi^g(v^d,l)}_{\pi^{sp}}$, $\%\delta^{\pi^g(v^s,l)}_{\pi^{sp}}$, and $\%\delta^{(\pi^g,v^s,l)}_{(\pi^g,v^d,l)}$. We apply min-max normalization for the continuous features and dummy encode the categorical features. The encoding removes the redundant dummy features including Canberra and HDD among cities and trucks, respectively. We use the type III sum of squares in the ANOVA. The full report is available in Appendix 3.A.

Table 3.5: Summary of the linear regression and ANOVA for seven different responses.

		$\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^s,l)}$		$\%\mathcal{E}^{(\pi^g,v^s,l)}_{(\pi^{sp},v^s,l)}$		$\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^d,l)}$		$\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^g,v^s,l)}$		$\%\delta^{\pi^g(v^s,l)}_{\pi^{sp}}$		$\%\delta^{\pi^g(v^d,l)}_{\pi^{sp}}$		$\%\delta^{(\pi^g,v^s,l)}_{(\pi^g,v^d,l)}$	
Features	*df*	R	S	R	S	R	S	R	S	R	S	R	S	R	S
$\sigma^{sp}(\theta)$	1	2	+	1	+	1	+	3	+	1	+	1	+	1	+
Δh	1	1	−	2	−	2	−	1	−	4	−	4	−	4	−
δ^{sp}	1	7	−	7	+	7	+	4	−	2	+	2	+	2	+
l	1	6	+	3	+	3	+	5	−	6	+	6	+	5	−
City	24	4	±	4	±	5	±	6	±	3	±	3	±	3	±
Truck	2	5	−	5	−	4	−	7	±	5	−	5	−	7	+
(Intercept)	1	3	−	6	−	6	−	2	−	7	−	7	−	6	+

R: Feature's ranking in ANOVA
S: Sign of the feature's weight in linear regression

Table 3.5 summarizes the ranking and sign of different features in the ANOVA as per Appendix 3.A. By the results, $\sigma^{sp}(\theta)$, i.e. the standard deviation of road gradient along the π^{sp} has the most explanatory power for CO_2 reduction capacity. In addition, $\sigma^{sp}(\theta)$ is has the strongest association with the dissimilarity of the greenest and shortest paths. That is to say, a higher $\sigma^{sp}(\theta)$ indicates a higher potential of CO_2 emissions reduction by selecting the greenest path instead of the shortest path. Next comes difference in elevation between the target and the source, Δh, which is negatively associated with the CO_2 emissions reduction capacity. This relation is strongest when comparing the dynamic speed policy with the static speed policy as in $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^g,v^s,l)}$. This implies that

using elevation data in routing policies is more pivotal for downward trips. Table 3.5 also reveals the positive association of $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^d,l)}$ and $\%\mathcal{E}^{(\pi^g,v^s,l)}_{(\pi^{sp},v^s,l)}$ with payload. Our analysis shows that relative dissimilarity of the shortest and greenest paths increases in distance of the shortest path, i.e. δ^{sp}. However, δ^{sp} is less important for CO_2 emissions reduction. A city's individual characteristics have a fair impact on the CO_2 emissions reduction capacity, albeit this effect is not comparable with that of $\sigma^{sp}(\theta)$ and Δh. Finally, the truck type has an effect that is similar to the payload. It follows that curb weight and payload of truck are more significant than other parameters for the CO_2 emissions reduction.

3.5 Numerical Experiments with Traffic Information

In this section, we examine how the simultaneous utilization of elevation and traffic data can guide routing decisions with lower CO_2 emissions over a large dataset. We compare the CO_2 emissions of different types of trucks traveling along three types of routes: the greenest path, $\pi^g(v,l)$, the asymptotic greenest path, $\pi^\infty(v)$, and the path with minimum possible travel duration, the fastest path π^{fp}. These comparisons are made under three different speed decisions: traffic speed v^f, dynamic speed v^d, and static speed v^s. For all arcs $a \in A$, the maximum speed $v^{\max}$ is set to the traffic speed v^f and the minimum speed $v^{\min}$ is set to zero.

We consider a strongly connected subgraph of New York city's road network for our experiments. The subgraph comprises 39,143 arcs and 23,091 vertices. Similar to Section 3.4, we obtain road network data from OpenStreetMap (OpenStreetMap contributors, 2017) and elevation data from the U.S. Geological Survey (2000)'s SRTM 1 Arc-Second Global datasets. The specifications of the trucks used in the study are provided in Table 3.4. Since traffic speed information is not publicly available for all arcs, we calculate traffic speeds using travel distance and duration inquiries from Google's Distance Matrix API. We select a time point with anticipated heavy traffic, particularly Wednesday, October 9, 2024, at 7:00 a.m., and set the traffic model to "best-guess". Given that $t_a(v^f(a))$ is the time to traverse arc $a \in A$ with traffic speed of arc a, i.e. $v^f(a)$, one can compute $v^f(a)$ by, $v^f(a) = \delta(a)/t_a(v^f(a))$.

We randomly select 20,098 unique pairs of non-identical source and target vertices. For each pair of source and target and each pair of path-speed policies $d_i = (\pi_i, v_i, l), i = 1, 2$, we compute three metrics including the relative CO_2 reduction $\%\mathcal{E}^{d_2}_{d_1}$, the relative path distinction $\%\delta^{\pi_2}_{\pi_1}$ and the relative time increase of selecting d_2 instead of d_1, $\%t^{d_2}_{d_1}$, defined by,

$$\%t^{d_2}_{d_1} = 100 \cdot \frac{\sum_{a\in\pi_2} t_a(v_2) - \sum_{a\in\pi_1} t_a(v_1)}{\sum_{a\in\pi_1} t_a(v_1)}.$$

Table 3.6 briefly summarizes the additional ratios that we use in our studies under

traffic conditions.

Table 3.6: List of ratios used in the comparative studies in addition to Table 3.2.

Ratio	Description
$\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{fp},v^f,l)}$	Relative CO_2 emissions reduction by selecting the greenest path with the dynamic speed policy relative to the fastest path with the traffic speed given the load l.
$\%\mathcal{E}^{(\pi^g,v^s,l)}_{(\pi^{fp},v^f,l)}$	Relative CO_2 emissions reduction by selecting the greenest path with the static speed policy relative to the fastest path with the traffic speed given the load l.
$\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{fp},v^d,l)}$	Relative CO_2 emissions reduction by selecting the greenest path with the dynamic speed policy relative to the fastest path with the dynamic speed policy given the load l.
$\%\mathcal{E}^{(\pi^g,v^s,l)}_{(\pi^{fp},v^s,l)}$	Relative CO_2 emissions reduction by selecting the greenest path with the static speed policy relative to the fastest path with the static speed policy given the load l.
$\%\mathcal{E}^{(\pi^g(v^f,l),v^d,l)}_{(\pi^{fp},v^d,l)}$	Relative CO_2 emissions reduction by selecting the greenest path under the assumption of driving at traffic speed but using the dynamic speed policy upon path traversal relative to the fastest path with the dynamic speed policy given the load l.
$\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^g(v^f,l),v^d,l)}$	Relative CO_2 emissions reduction by selecting the greenest path with the dynamic speed policy relative to the greenest path under the assumption of driving at traffic speed but using the dynamic speed policy upon path traversal given the load l.
$\%\mathcal{E}^{(\pi^\infty,v^d,l)}_{(\pi^{fp},v^d,l)}$	Relative CO_2 emissions reduction by selecting the asymptotic greenest path with the dynamic speed policy relative to the fastest path with the static speed policy given the load l.
$\%\delta^{\pi^g(v^d,l)}_{\pi^{fp}}$	Ratio of the length of the fastest path that is not shared with the greenest path under the dynamic speed policy given the load l.
$\%\delta^{\pi^g(v^s,l)}_{\pi^{fp}}$	Ratio of the length of the fastest path that is not shared with the greenest path under the static speed policy given the load l.
$\%t^{(\pi^g,v^d,l)}_{(\pi^{fp},v^f,l)}$	Relative time increase by selecting the greenest path with the dynamic speed policy relative to the fastest path with the traffic speed given the load l.
$\%t^{(\pi^g,v^d,l)}_{(\pi^{fp},v^d,l)}$	Relative time increase by selecting the greenest path with the dynamic speed policy relative to the fastest path with the dynamic speed policy given the load l.
$\%t^{(\pi^g,v^s,l)}_{(\pi^{fp},v^f,l)}$	Relative time increase by selecting the greenest path with the static speed policy relative to the fastest path with the traffic speed given the load l.
$\%t^{(\pi^g,v^s,l)}_{(\pi^{fp},v^s,l)}$	Relative time increase by selecting the greenest path with the static speed policy relative to the fastest path with the static speed policy given the load l.

3.5.1 Results: Impact of Path and Speed Decisions on CO_2 Emissions Reduction

Figure 3.24 demonstrates that $\overline{\%\mathcal{E}}^{(\pi^g,v^d,l)}_{(\pi^{fp},v^f,l)}$ ranges from 19.40% to 26.02% for different truck types and payloads, whereas Figure 3.25 shows that $\overline{\%\mathcal{E}}^{(\pi^g,v^d,l)}_{(\pi^{fp},v^d,l)}$ varies between 5.80% and 8.91%. Similar statistics are observed for $\overline{\%\mathcal{E}}^{(\pi^g,v^s,l)}_{(\pi^{fp},v^s,l)}$ as illustrated in Figure 3.26. These results provide significant evidence that both path selection and speed optimization can contribute to reducing CO_2 emissions in intra-city truck transportation. Additionally, the reduction potential in CO_2 emissions is greater on the greenest path during traffic conditions compared to free-flow situations (cf. Figures 3.2 and 3.3). However, the potential

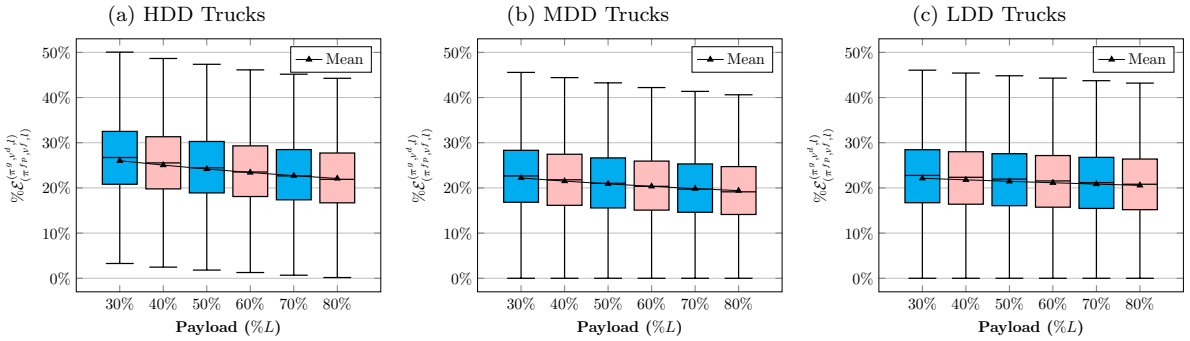

Figure 3.24: $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{fp},v^f,l)}$ **across truck types and payloads in traffic condition.**

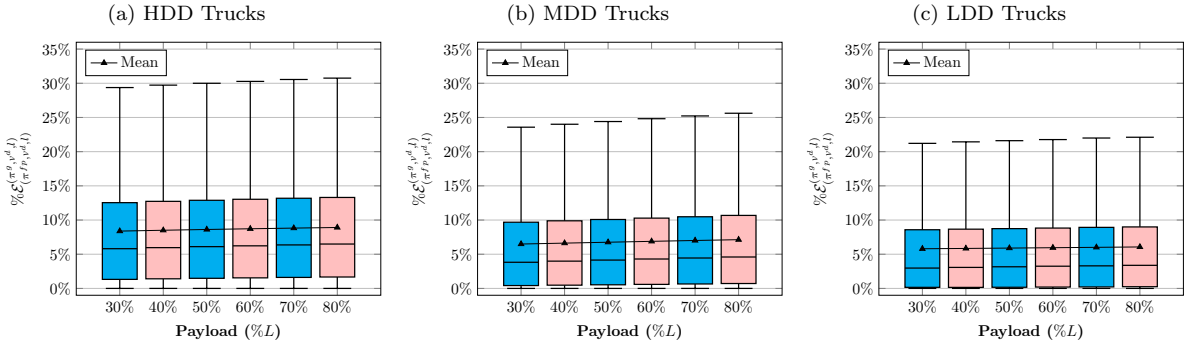

Figure 3.25: $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{fp},v^d,l)}$ **across truck types and payloads in traffic condition.**

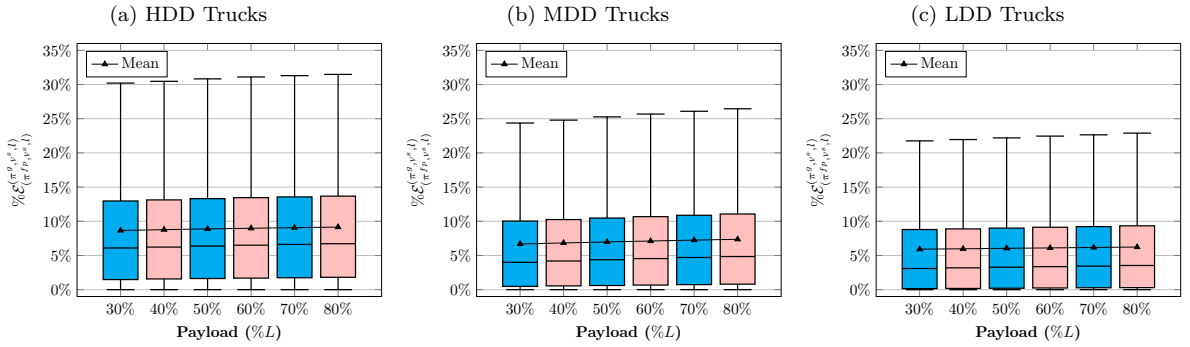

Figure 3.26: $\%\mathcal{E}^{(\pi^g,v^s,l)}_{(\pi^{fp},v^s,l)}$ **across truck types and payloads in traffic condition.**

reduction in CO_2 emissions through a dynamic speed policy versus a static speed policy is negligible in most instances, as illustrated in Figure 3.27. This figure demonstrates that the sample mean and third quartiles of $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^g,v^s,l)}$ are below 1% across all truck types and payloads. This result is primarily because the traffic conditions hinder trucks from utilizing gravity for acceleration on downhill segments, in most instances. Nevertheless, optimizing the speed on uphill segments can substantially reduce CO_2 emissions. If the route planner selects the greenest path for traffic speed, $\pi^g(v^f,l)$, rather than the fastest path, and the traveling speed is v^d, the average CO_2 emissions reduction, $\overline{\%\mathcal{E}}^{(\pi^g(v^f,l),v^d,l)}_{(\pi^{fp},v^d,l)}$,

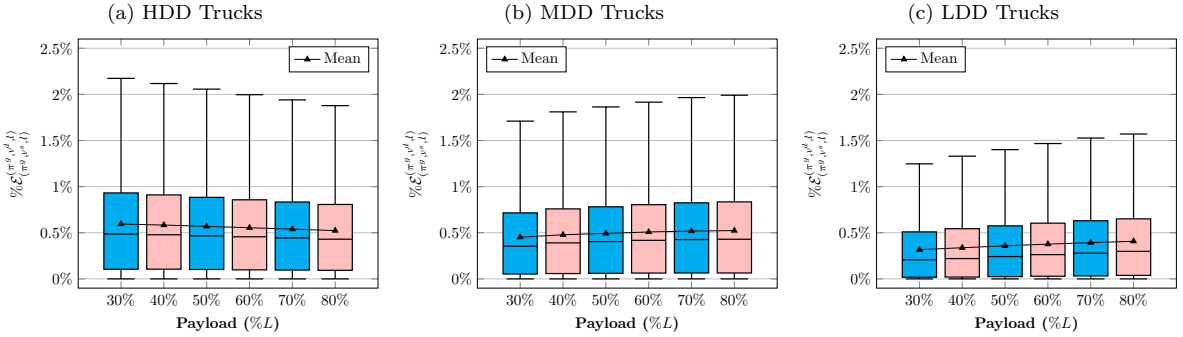

Figure 3.27: $\%\mathcal{E}^{(\pi^g,v^s,l)}_{(\pi^g,v^d,l)}$ **across truck types and payloads in traffic condition.**

ranges from 3.01% to 7.03% (see Figure 3.28). Although $\pi^g(v^f,l)$ is not the optimal

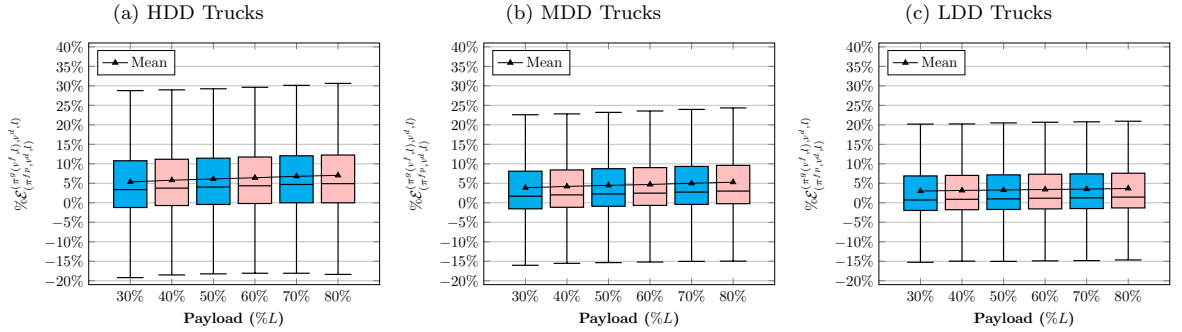

Figure 3.28: $\%\mathcal{E}^{(\pi^g(v^f,l),v^d,l)}_{(\pi^{fp},v^d,l)}$ **across truck types and payloads in traffic condition.**

path for minimizing CO_2 emissions when v^d is decided, a comparison of Figures 3.25 and 3.28 reveals that choosing $\pi^g(v^f,l)$ instead of π^{fp} can achieve more than half of the potential CO_2 emissions reduction in most instances (the same argument holds under v^s). It is worth noting that CO_2 reduction potential of $\pi^g(v^d,l)$ (or $\pi^g(v^s,l)$) over $\pi^g(v^f,l)$ is slightly higher for lower payloads (see Figure 3.29). This phenomenon is due to the convergence of the greenest paths to the asymptotic greenest paths.

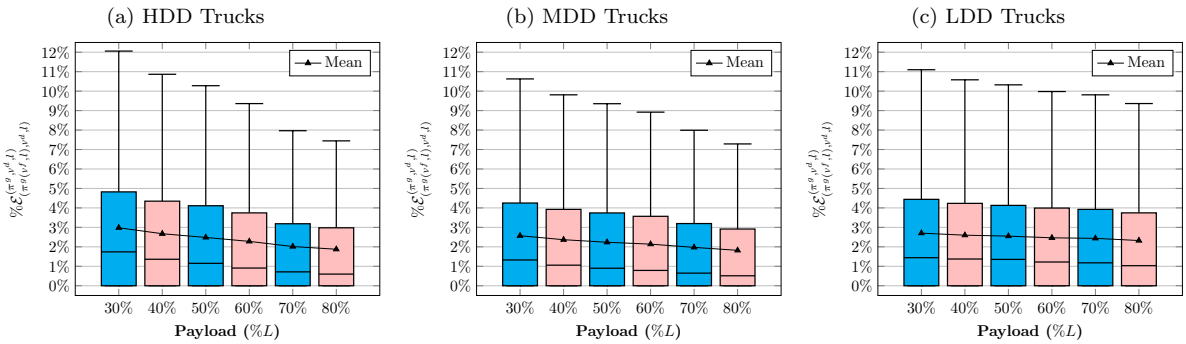

Figure 3.29: $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^g(v^f,l),v^d,l)}$ **across truck types and payloads in traffic condition.**

3.5.2 Results: Increased Travel Duration

Figures 3.30 presents statistics on the increased travel duration when trucks travel on the greenest paths with v^d instead of the fastest path with v^f. The figure indicates that, on

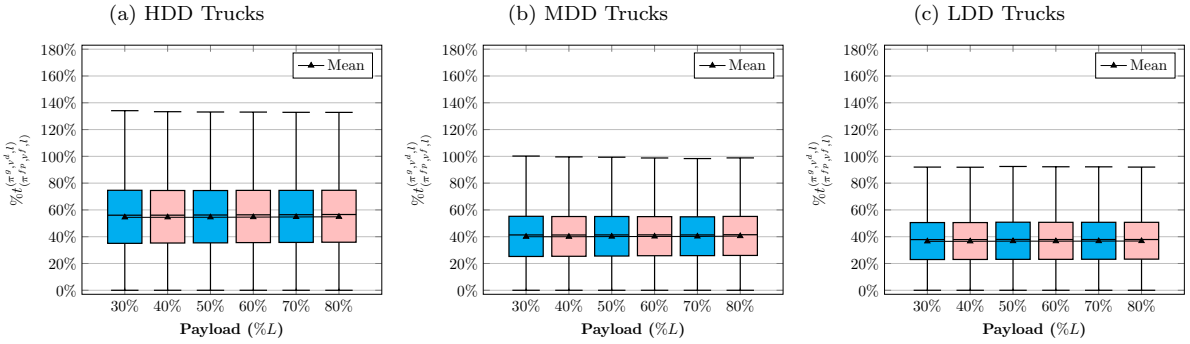

Figure 3.30: $\%t^{(\pi^g,v^d,l)}_{(\pi^{fp},v^f,l)}$ **across truck types and payloads in traffic condition.**

average, the travel duration increases relative to the fastest path (incorporating traffic speed) from 36.53% for LDD trucks with 30% payload to 54.48% for HDD trucks with 80% payload. However, if the speed policy for both greenest path and fastest path is v^d, the average advantage of selecting the fastest path in terms of travel duration is less than 2.28% (see Figures 3.31). In several instances, a truck traverses the greenest path

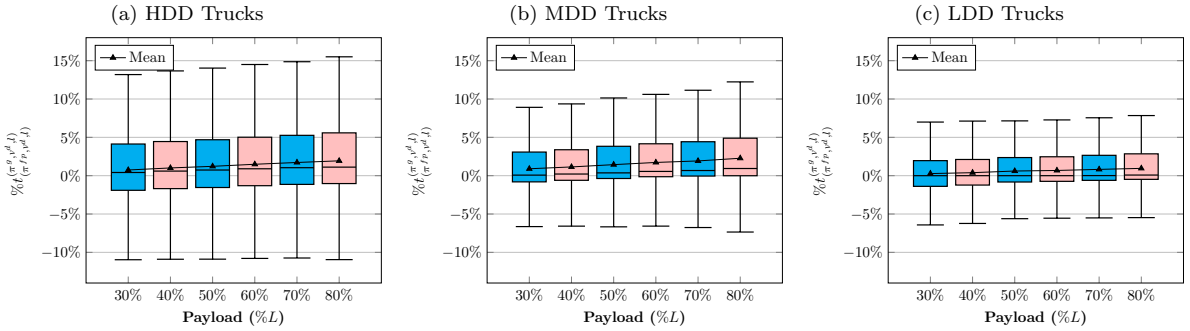

Figure 3.31: $\%t^{(\pi^g,v^d,l)}_{(\pi^{fp},v^d,l)}$ **across truck types and payloads in traffic condition.**

even faster than the fastest path when the selected speed policy is v^d because the fastest path is found under the assumption of the traffic speed policy. The statistics presented in Sections 3.5.1 and 3.5.2 clearly show that when the dynamic speed policy is selected, the CO_2 emissions reduction is larger than the increased travel duration. This argument holds for the static speed policy, v^s. The statistics of the static speed policy are presented in Appendix 3.5.

3.5.3 Results: Paths of the $\pi^g(v^d, l)$, $\pi^g(v^s, l)$, $\pi^g(v^f, l)$, and π^{fp}

The dissimilarities between the fastest path and the greenest paths is an important factor in the CO_2 reduction potential of the greenest paths, as highlighted in Section 3.5.1. Figures 3.32 to 3.33 present the statistics for these dissimilarities. Figure 3.32 shows

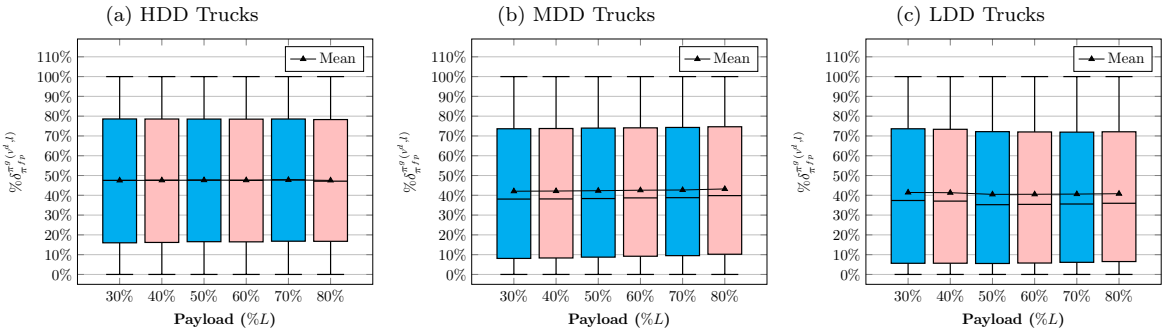

Figure 3.32: $\%\delta_{\pi^{fp}}^{\pi^g(v^d,l)}$ **across truck types and payloads in traffic condition.**

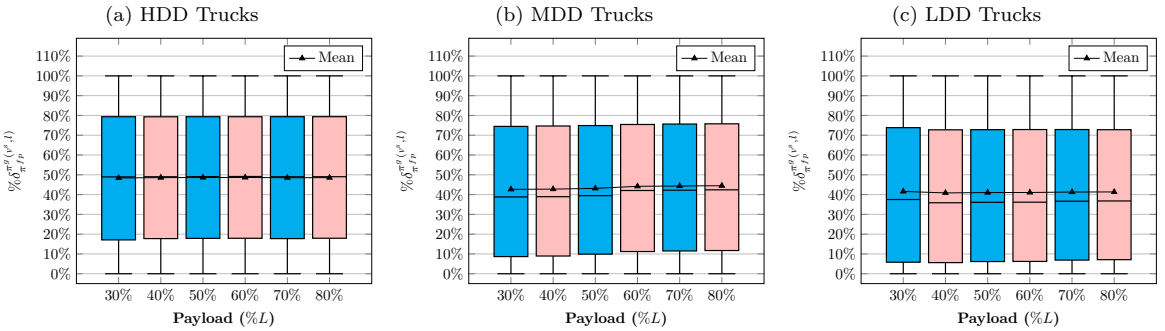

Figure 3.33: $\%\delta_{\pi^{fp}}^{\pi^g(v^s,l)}$ **across truck types and payloads in traffic condition.**

that, on average, π^{fp} does not share 40.47% to 47.81% of its paths with $\pi^g(v^d, l)$. A similar statistic for $\pi^g(v^s, l)$, i.e., $\overline{\%\delta}_{\pi^{fp}}^{\pi^g(v^s,l)}$, ranges from 41.01% to 48.62%, which is slightly higher than that of $\overline{\%\delta}_{\pi^{fp}}^{\pi^g(v^d,l)}$ (see Figure 3.33). We observed a similar pattern in Section 3.4.3 (see Figures 3.13 and 3.14), where we compared the greenest paths and the shortest path in the free flow conditions. Regarding the dissimilarity between $\pi^g(v^d, l)$ and $\pi^g(v^s, l)$, our experiments show that $\%\delta_{\pi^g(v^d,l)}^{\pi^g(v^s,l)}$ is zero or negligible for the majority of instances. Figure 3.34 indicates that the first, second, and third quartiles, as well as the upper whisker of $\%\delta_{\pi^g(v^d,l)}^{\pi^g(v^s,l)}$, are zero, and the maximum $\overline{\%\delta}_{\pi^g(v^d,l)}^{\pi^g(v^s,l)}$ is 2.62% (cf. Figure 3.12). As mentioned in Section 3.5.1 this result stems from the limitations on maximum speed on the downhill arcs due to traffic conditions.

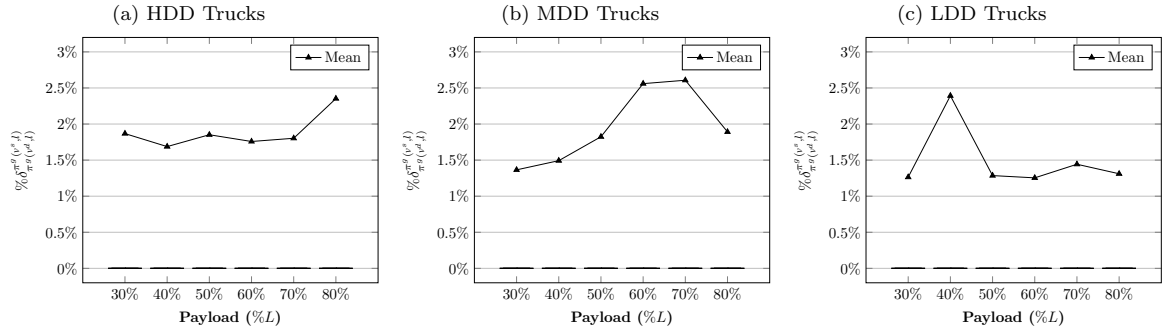

Figure 3.34: $\%\delta^{\pi^g(v^s,l)}_{\pi^g(v^d,l)}$ **across truck types and payloads in traffic condition.**

3.5.4 Results: Asymptotic Greenest Path under Traffic

In Section 3.4.4, we explained that the convergence of greenest paths to the asymptotic greenest paths is observable for all truck types as the payload increases. This section examines the asymptotic greenest paths under traffic conditions. We focus exclusively on $\pi^\infty(v^g)$, since similar trends can be expected for $\pi^\infty(v^s)$, as discussed in Section 3.5.3. Figure 3.35 illustrates the average CO_2 reduction achieved by the asymptotic greenest path compared to the fastest path with the same dynamic speed policy. Specifically,

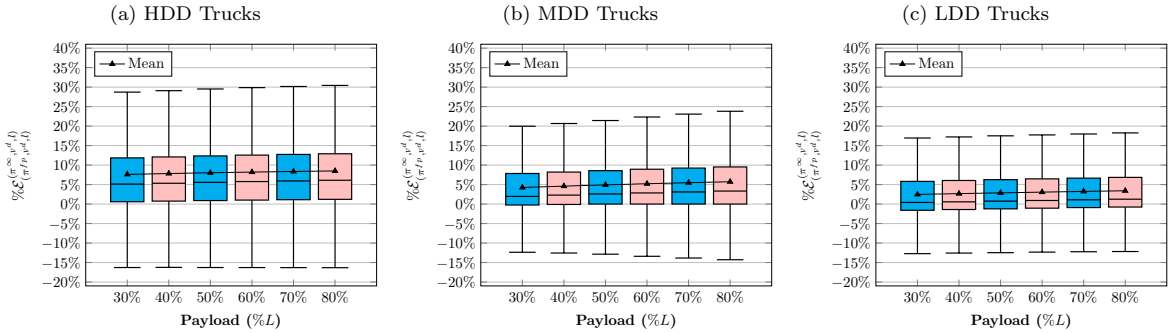

Figure 3.35: $\%\mathcal{E}^{(\pi^\infty,v^d,l)}_{(\pi^{fp},v^d,l)}$ **across truck types and payloads in traffic condition.**

$\overline{\%\mathcal{E}}^{(\pi^\infty,v^d,l)}_{(\pi^{fp},v^d,l)}$ ranges from 2.47% for LLD trucks with 30% payload to 8.50% for HDD trucks with 80% payload. This result indicates that the CO_2 reduction potential of $\pi^\infty(v^g)$ is similar to that of $\pi^g(v^f,l)$, even slightly higher for MDD and HDD trucks (cf. Figure 3.28). Figure 3.36 shows the CO_2 emissions reduction of the asymptotic greenest path relative to the greenest path, $\%\mathcal{E}^{\pi^\infty,v^d,l}_{\pi^g,v^d,l}$. It is straightforward to see that $\overline{\%\mathcal{E}}^{\pi^\infty,v^d,l}_{\pi^g,v^d,l}$ increases with truck weight, rising from -3.56% for LLD trucks with 30% payload to -0.45% for HDD trucks with 80% payload. Figure 3.37 highlights the convergence of the greenest paths to the asymptotic greenest paths, similar to the tendency observed in Section 3.4.4 for free flow conditions. Comparing the convergence results under free flow and traffic conditions in New York city, we can infer that, on average, convergence under traffic occurs more rapidly than free-flow condition. This phenomenon is due to the limitations on speed

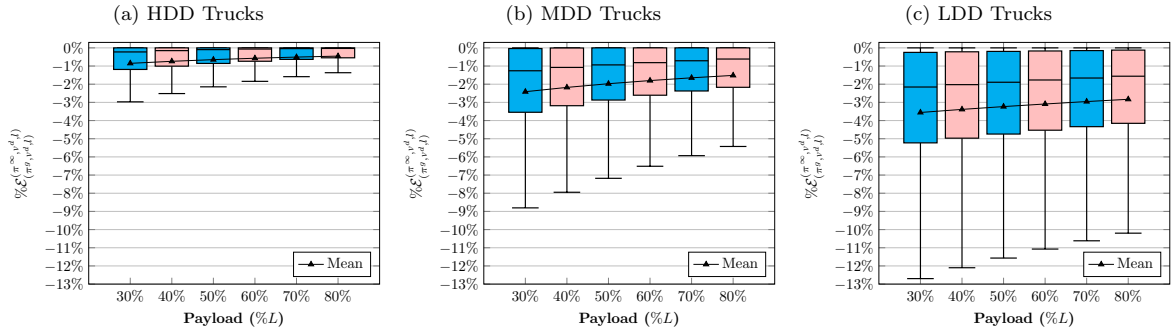

Figure 3.36: $\%\mathcal{E}^{(\pi^\infty, v^d, l)}_{(\pi^g, v^d, l)}$ **across truck types and payloads in traffic condition.**

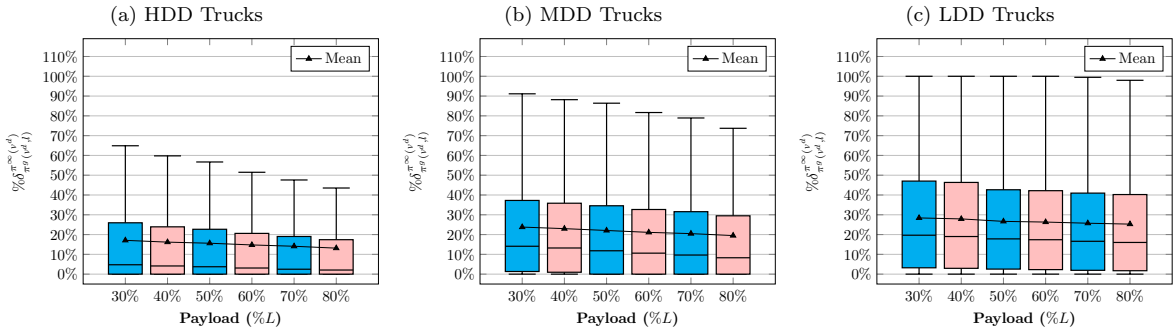

Figure 3.37: $\%\delta^{\pi^\infty(v^d)}_{\pi^g(v^d,l)}$ **across truck types and payloads in traffic condition.**

choices in traffic. In such a case, the selection of paths with less ascent, i.e. $h'(a)$, $a \in A$, plays a crucial role in the reduction of emissions CO_2 .

3.5.5 Conclusions on the role of traffic information

The incorparation of slope information to find the greenest path is even more important in heavy traffic than it is in free flow traffic. Comparing the results of dynamic and static speed optimizations in traffic, we find that optimizing speed on uphill arcs can significantly reduce CO_2 emissions. However, using gravity to accelerate on downhill arcs is limited due to traffic congestion. From a public policy perspective, this finding reinforces arguments for scheduling truck deliveries during "Off-Hours" when more environmentally-friendly options for path and speed selection are available (see e.g. New York City Department of Transportation, 2024). Truck delivery during "Off-Hours" can, additionally, reduce traffic congestion when trucks traverse uphill roads with optimized speeds that may be lower than the traffic speed. In Section 3.5.4, we demonstrate that the greenest path converges to the asymptotic greenest path even faster under traffic conditions than in free flow conditions.

3.6 Summary of Key Findings from Numerical Experiments

The first outcome of our experiments in Sections 3.4 and 3.5 is that high-resolution topographical data should be incorporated into urban truck transportation decisions when minimizing CO_2 emissions is the objective. Specifically, pre-computation of the greenest paths is not feasible due to the non-linear effects of speed decisions, road gradients, and payload. A similar argument has previously been made regarding the need to integrate high-resolution traffic speed data into emissions-minimizing transportation decisions (see e.g. Ehmke et al., 2016b).

Secondly, our results show that optimal speed decisions are dynamic, with dynamic speed choices reducing CO_2 emissions by 2% to 4% in free-flow conditions compared to static speed choices. While the difference between dynamic and static speed decisions is less significant in traffic, we found that optimized speeds still achieve significantly lower emissions than traffic speeds, even when acceleration is restricted by traffic congestion.

Thirdly, we observed that the greenest path is relatively insensitive to whether speed decisions are static or dynamic, even in free-flow conditions. Additionally, the greenest path begins to converge to the asymptotic greenest path at low payload ratios under both free-flow and traffic conditions. Therefore, a pre-computed greenest path for a given speed decision (e.g., static) and payload level (e.g., 50% or 100%) can be a good approximation for the greenest paths across different speed decisions and payloads. This approximation can help reduce the computational complexity of green transportation problems like PRP.

3.7 Conclusions

In this chapter, we studied the greenest path selection problem for a logistics service provider that operates a fleet of heavy-, medium-, and light-duty trucks in an urban environment. We established that the policies for the speed and path that minimize CO_2 emissions are slope-dependent (dynamic). We also showed that the greenest path converges to a fixed path as the payload increases and provided an efficient algorithm to compute the asymptotic greenest path. We conducted extensive numerical experiments using elevation data of 25 cities around the world to investigate the potential CO_2 reduction by such dynamic policies under free flow traffic conditions. The results in section 3.4.2 showed that, on average, the combined dynamic path and speed selection can reduce CO_2 emissions by 1.19% to 10.15% based on the truck type and city. Our analysis also showed that in most cities, the average emissions reduction potential of dynamic speed optimization lies between 2% to 4% regardless of the truck type. Nonetheless, the effect of slope-dependent path selection (the greenest path) depends on the payload and

truck type. In section 3.5, for the city of New York we also studied the effect of effective speed limits due to traffic congestion and found that choosing the greenest path can lead to even larger CO_2 reduction than in free flow traffic conditions. In section 3.4.3, we explained that the greenest path significantly differs from the shortest path. While the greenest path depends on the speed policy, the experiments show that this dependence is weak and that the greenest path under the static speed policy is usually optimal or near optimal especially when speed is determined by traffic. Moreover, we demonstrated, in Sections 3.3.5 and 3.4.4, that the greenest path diverges from the shortest path as the payload increases and converges to the asymptotic greenest path, i.e. the greenest path for the arbitrary large payloads. Convergence to the asymptotic greenest path is faster under heavy traffic conditions. These results could be used for the approximation of the greenest path to simplify complex transportation problems. The analysis of variance (ANOVA) indicated that the potential CO_2 emissions reduction by the greenest path and the dynamic speed policy is associated positively with the variability of arc gradients along the shortest path, and negatively with the relative elevation of the source and target.

3.A ANOVA Results

Table 3.7: Response: $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^s,l)}$

Feature	*df*	SS	*MSS*	$F-value$	$p-value$
$\sigma^{sp}(\theta)$	1	11771	11771	9061762	$< 10^{-15}$
Δh	1	22347	22347	17203995	$< 10^{-15}$
l	1	293	293	225597	$< 10^{-15}$
δ^{sp}	1	47	47	36175	$< 10^{-15}$
City	24	12477	520	400209	$< 10^{-15}$
Truck	2	926	463	356459	$< 10^{-15}$
(Intercept)	1	7265	7265	5592723	$< 10^{-15}$
Residuals	55415417	71983			

Table 3.8: Response: $\%\mathcal{E}^{(\pi^g,v^s,l)}_{(\pi^{sp},v^s,l)}$

Feature	*df*	SS	*MSS*	$F-value$	$p-value$
$\sigma^{sp}(\theta)$	1	10020	10020	5973305	$< 10^{-15}$
Δh	1	6261	6261	3732651	$< 10^{-15}$
l	1	908	908	541007	$< 10^{-15}$
δ^{sp}	1	225	225	134066	$< 10^{-15}$
City	24	16178	674	401852	$< 10^{-15}$
Truck	2	1156	578	344674	$< 10^{-15}$
(Intercept)	1	420	420	250138	$< 10^{-15}$
Residuals	55415417	92958			

Table 3.9: Response: $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^{sp},v^d,l)}$

Feature	*df*	SS	*MSS*	$F-value$	$p-value$
$\sigma^{sp}(\theta)$	1	7092	7092	6099148	$< 10^{-15}$
Δh	1	4822	4822	4147274	$< 10^{-15}$
l	1	816	816	701716	$< 10^{-15}$
δ^{sp}	1	82	82	70510	$< 10^{-15}$
City	24	10056	419	360351	$< 10^{-15}$
Truck	2	918	459	394610	$< 10^{-15}$
(Intercept)	1	347	347	298743	$< 10^{-15}$
Residuals	55415417	64434			

Table 3.10: Response: $\%\mathcal{E}^{(\pi^g,v^d,l)}_{(\pi^g,v^s,l)}$

Feature	*df*	SS	*MSS*	$F-value$	$p-value$
$\sigma^{sp}(\theta)$	1	1431	1431	3270532	$< 10^{-15}$
Δh	1	16335	16335	37344012	$< 10^{-15}$
l	1	147	147	335105	$< 10^{-15}$
δ^{sp}	1	870	870	1989427	$< 10^{-15}$
City	24	3102	129	295466	$< 10^{-15}$
Truck	2	81	40	91981	$< 10^{-15}$
(Intercept)	1	10731	10731	24531295	$< 10^{-15}$
Residuals	55415417	24240			

Table 3.11: Response: $\%\delta_{\pi^{sp}}^{\pi^g(v^s,l)}$

Feature	*df*	SS	*MSS*	$F-value$	$p-value$
$\sigma^{sp}(\theta)$	1	204303	204303	3344542	$< 10^{-15}$
Δh	1	15691	15691	256873	$< 10^{-15}$
l	1	5270	5270	86270	$< 10^{-15}$
δ^{sp}	1	44990	44990	736513	$< 10^{-15}$
City	24	640819	26701	437105	$< 10^{-15}$
Truck	2	13408	6704	109750	$< 10^{-15}$
(Intercept)	1	2732	2732	44716	$< 10^{-15}$
Residuals	55415417	3385080			

Table 3.12: Response: $\%\delta_{\pi^{sp}}^{\pi^g(v^d,l)}$

Feature	*df*	SS	*MSS*	$F-value$	$p-value$
$\sigma^{sp}(\theta)$	1	208923	208923	3601979	$< 10^{-15}$
Δh	1	17790	17790	306719	$< 10^{-15}$
l	1	7602	7602	131065	$< 10^{-15}$
δ^{sp}	1	37788	37788	651499	$< 10^{-15}$
City	24	577881	24078	415128	$< 10^{-15}$
Truck	2	18821	9411	162244	$< 10^{-15}$
(Intercept)	1	3005	3005	51805	$< 10^{-15}$
Residuals	55415417	3214225			

Table 3.13: Response: $\%\delta_{(\pi^g,v^d,l)}^{(\pi^g,v^s,l)}$

Feature	*df*	SS	*MSS*	$F-value$	$p-value$
$\sigma^{sp}(\theta)$	1	1190	1190	74122	$< 10^{-15}$
Δh	1	596	596	37147	$< 10^{-15}$
l	1	390	390	24271	$< 10^{-15}$
δ^{sp}	1	1096	1096	68270	$< 10^{-15}$
City	24	22738	947	59004	$< 10^{-15}$
Truck	2	310	155	9667	$< 10^{-15}$
(Intercept)	1	203	203	12647	$< 10^{-15}$
Residuals	55415417	889779			

3.B Results: Performance of the Asymptotic Paths with Static Speed policies

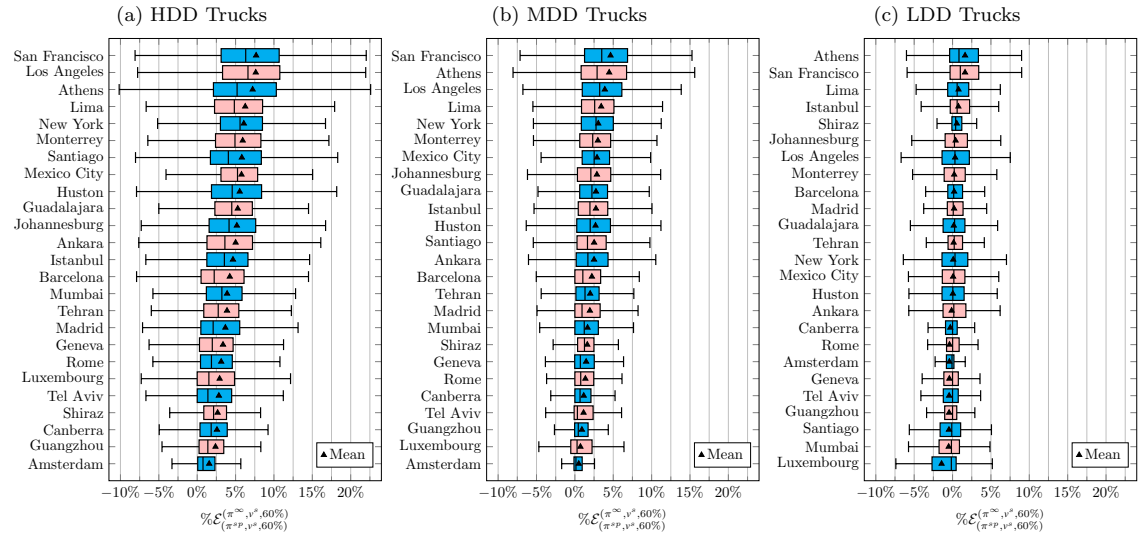

Figure 3.38: Relative CO_2 emissions reduction by selecting $(\pi^{\infty}, v^s, 60\%)$ rather than $(\pi^{sp}, v^s, 60\%)$.

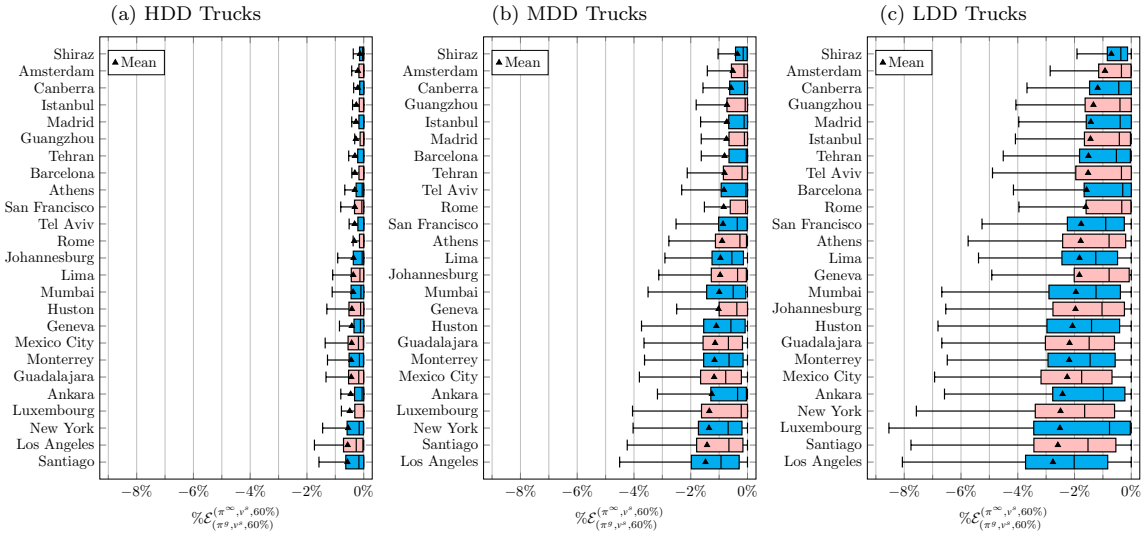

Figure 3.39: Relative CO_2 emissions reduction by selecting $(\pi^{\infty}, v^s, 60\%)$ rather than $(\pi^{g}, v^s, 60\%)$.

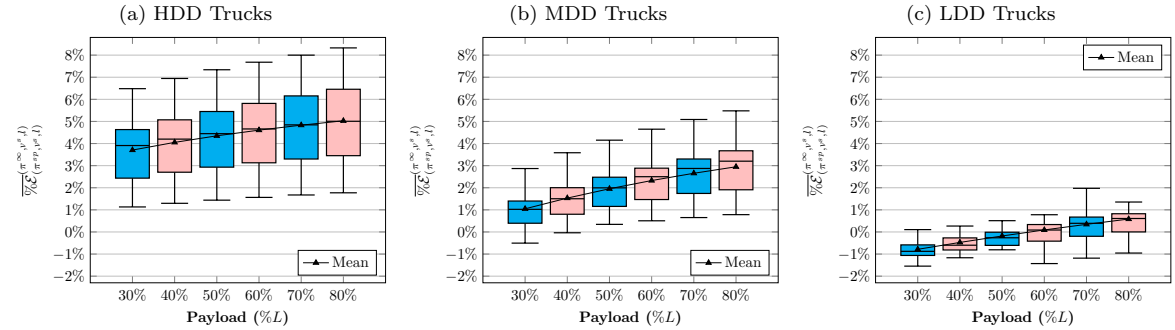

Figure 3.40: Effect of payload on $\overline{\%\mathcal{E}}^{(\pi^\infty, v^s, l)}_{(\pi^{sp}, v^s, l)}$ across 25 cities.

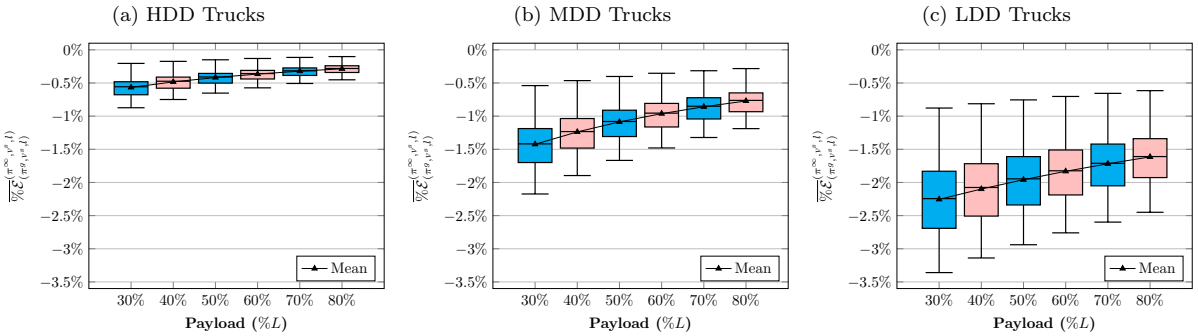

Figure 3.41: Effect of payload on $\overline{\%\mathcal{E}}^{(\pi^\infty, v^s, l)}_{(\pi^{g}, v^s, l)}$ across 25 cities.

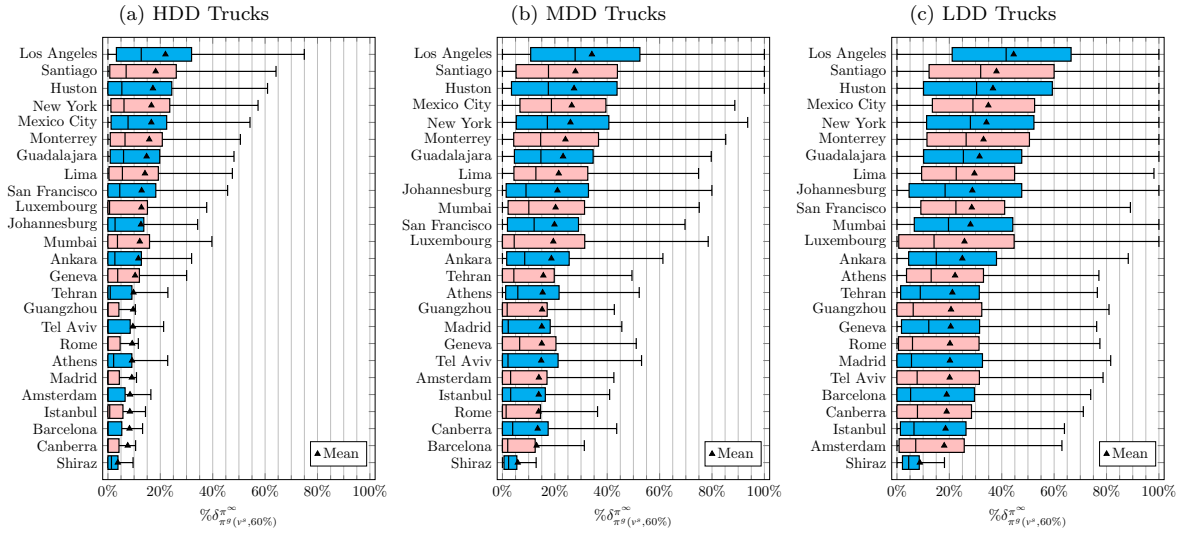

Figure 3.42: Ratio of the length of $\pi^g(v^s, 60\%)$ that is not shared with π^∞.

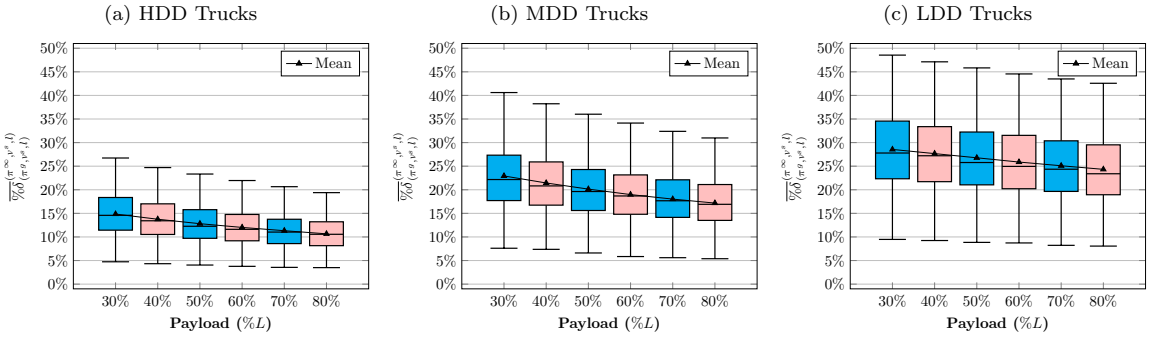

Figure 3.43: Effect of payload on $\overline{\%\delta}^{(\pi^\infty, v^s, l)}_{(\pi^g, v^s, l)}$ across 25 cities.

3.C Results: Increased Travel Duration under v^s

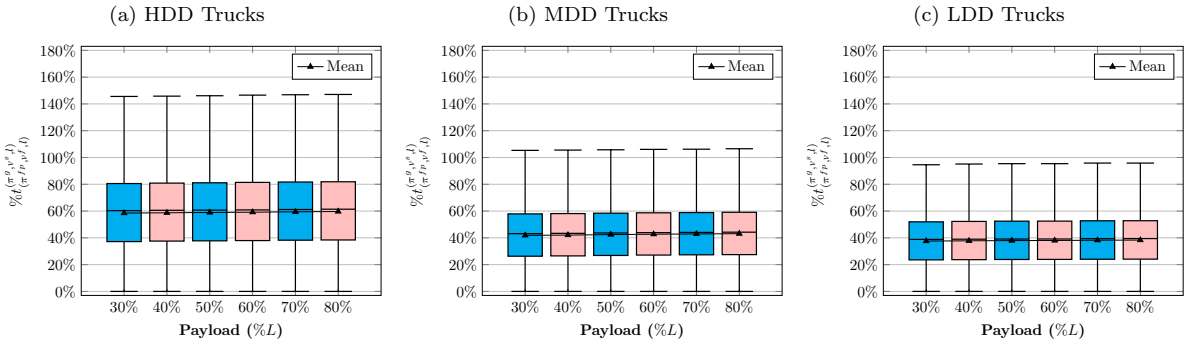

Figure 3.44: $\%t^{(\pi^g, v^s, l)}_{(\pi^{fp}, v^f, l)}$ across truck types and payload in traffic condition.

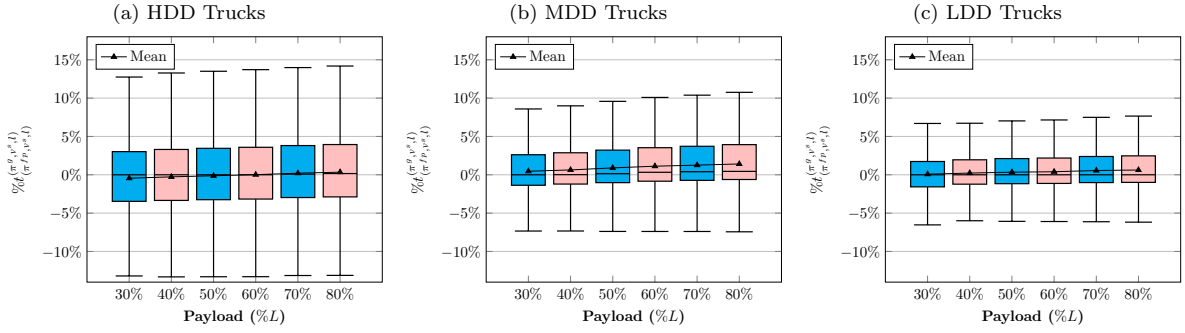

Figure 3.45: $\%t^{(\pi^g, v^s, l)}_{(\pi^{fp}, v^s, l)}$ across truck types and payloads in traffic condition.

Chapter 4

Asymptotic Optimality of Projected Inventory Level Policies for Lost Sales Inventory Systems with Large Leadtime and Penalty Cost

4.1 Introduction

The control of lost sales inventory systems remains a fundamental challenge in inventory theory. In such systems, unmet demand caused by stockouts is lost, often resulting in substantial penalty costs. We consider the canonical lost sales inventory system, which is a single-item, single-echelon, periodic-review inventory system with a positive leadtime and independent and identically distributed (i.i.d.) demand under the average cost criterion. This system serves as the foundation for more complex lost sales inventory models. Therefore, developing well-performing and computationally efficient control policies for the canonical system is crucial to derive effective policies for real-world lost sales inventory problems.

The optimal replenishment policy for the canonical system with negligible leadtime reduces to a newsvendor problem. When the leadtime is positive, the optimal policy can be found through dynamic programming but this is intractable due to the curse of dimensionality. Consequently, a key inventory research stream in stochastic lost sales inventory control focuses on developing simple heuristic policies that perform well under specific conditions, such as achieving asymptotic optimality in certain scaling regimes. We refer interested readers to Bijvank et al. (2023), Goldberg et al. (2021), and Bijvank and Vis (2011), for further discussions on lost sales inventory systems and related asymptotic optimality results.

Huh et al. (2009) and Bijvank et al. (2014) analyze base-stock policies that place orders to raise the inventory position to a fixed base-stock level. They establish that such

policies are asymptotically optimal as the cost of losing a sale grows for a fixed leadtime. Goldberg et al. (2016) and Xin and Goldberg (2016) demonstrate that a constant order policy, which places the same order quantity every period, is asymptotically optimal as the leadtime grows for a fixed cost of losing a sale. Both the base-stock policy and the constant order policy rely on a single parameter, making them easy to implement in practice. However, neither policy is optimal across both asymptotic regimes. To address this limitation, Xin (2021) proposes a two-parameter hybrid policy that integrates the base-stock and constant order policies, and proves its asymptotic optimality for large leadtimes. This policy, known as the capped base-stock policy, was initially studied by Johansen and Thorstenson (2008) and can be readily shown to be asymptotically optimal as the cost of losing a sale grows large. By adjusting its parameters, the capped base-stock policy can thus be tailored to achieve asymptotic optimality in either regime.

Recently, van Jaarsveld and Arts (2024) introduced the projected inventory level (PIL) policy, which places orders to ensure that the expected inventory level at the time of receipt reaches a fixed target. Unlike constant order and base-stock policies, the PIL policy dynamically adjusts order quantities by leveraging probabilistic information available at each decision epoch. van Jaarsveld and Arts (2024) demonstrate that the PIL policy consistently outperforms the base-stock policy for general demand distributions and prove that it also outperforms the constant order policy when demand is exponential. PIL policies are also asymptotically optimal for perishable inventory systems in several regimes (Bu et al., 2025a,b), and the projection idea is similarly employed by Drent and Arts (2022) for dual-sourcing inventory systems, where it yields both asymptotic optimality and strong empirical performance.

Policies developed for the canonical lost sales system can be extended to more complex settings, including systems with non-stationary demand, perishable items, continuous review, partially observable parameters, finite storage capacity, supply uncertainty, stochastic returns, joint inventory and pricing control, and finite horizon decision making (see, e.g., Bu et al., 2025a,b, 2024, 2020; Lyu et al., 2024; Bai et al., 2023; Xin, 2022; Chen et al., 2021).

The asymptotic regimes discussed in the literature only consider one parameter growing large while keeping all other parameters fixed. However, many practically important items exhibit both long lead times and high lost-sales costs, including critical components in aerospace and semiconductor manufacturing and pharmaceuticals used for rare or life-saving treatments. In this chapter we study the performance of the PIL policy under a general demand process as the leadtime grows large when the cost of losing a sale is sufficiently large. Under mild conditions on the demand distribution we show that:

1. The difference between the average cost-rate of the PIL policy and that of the constant order policy remains bounded by a finite constant.

2. The PIL policy is asymptotically optimal for sufficiently large lost sales penalty

costs as the leadtime approaches infinity.

Our analysis hinges on new bounds we derive for the solution of the Wiener-Hopf equation that characterizes the relative value function under the constant order policy. These bounds follow from studying the ladder processes of a random walk with increments equal to the per-period excess demand minus the constant order. We then apply a one-step policy improvement technique to analyze the cost-rate difference between the PIL and constant order policies.

The rest of the chapter is organized as follows. Section 4.2 introduces the model and optimization problem (Section 5.3), as well as the main result (Section 4.2.2). Section 4.3 provides the proof of the main result, including the introduction of ladder processes (Section 4.3.1), the solution to our Wiener-Hopf equation (Section 4.3.2), asymptotic inventory dynamics (Section 4.3.3), and policy improvement argument (Section 4.3.4) which completes the proof. A summary of results and final remarks are provided in Section 4.4. All proofs are included in the Appendix, unless otherwise specified.

4.2 Model and main result

4.2.1 Model

We consider an infinite-horizon periodic review lost sales inventory system. Demand in period t is denoted D_t and $\{D_t\}_{t\in\mathbb{N}_0}$ ($\mathbb{N}_0 := \mathbb{N}\cup\{0\}$) is a sequence of non-negative independent and identically distributed random variables with distribution function F_D supported on $[0,\infty)$, and finite mean $\mu_D := \mathbb{E}[D] < \infty$ and variance $\mathbf{Var}[D] := \sigma_D^2 \in (0,\infty)$. We assume $F_D(0) = 0$ for notational simplicity, though all results remain valid without this assumption. Each time period $t \in \mathbb{N}_0$ we receive an order, $q_t \in \mathbb{R}_+$, that is placed in period $t-L$, where $L \in \mathbb{N}_0$ is the deterministic leadtime. Let $\{J_t\}_{t\in\mathbb{N}_0}$ denote the sequence of inventory level random variables at the beginning of each period before receiving the order. The state of the system at the start of period $t \in \mathbb{N}_0$, denoted by $\mathbf{x}_t \in \mathbb{R}_+^{L+1}$, is a vector comprising the inventory level in period t as well as the outstanding orders in the pipeline. That is, $\mathbf{x}_t = (J_t, q_t, q_{t+1}, \ldots, q_{t+L-1})$. We assume that $\mathbf{x}_0$ is fixed and known, and $J_0 = 0$. Demand that exceeds the on-hand inventory $J_t + q_t$ is lost at the end of the period at the unit cost of $p \geq 0$. Any surplus inventory at the end of a period is held at a cost of $h \geq 0$ per item. The sequence of events in each period $t \in \mathbb{N}_0$ is as follows: (1) The state of the system $\mathbf{x}_t$ is observed and the order q_{t+L} is placed, (2) The order q_t is received, (3) The demand D_t is realized, and (4) the costs of period t are incurred as $p(D_t - q_t - J_t)^+ + h(J_t + q_t - D_t)^+$ where, $(x)^+ := \max(x, 0)$. The dynamics of the inventory level are

$$J_{t+1} = (J_t + q_t - D_t)^+. \tag{4.1}$$

A policy π is a set of mappings from the space of the states, $\mathbf{x}_t$, to the space of orders,

q_{t+L}, i.e., $\{\pi_t : \mathbb{R}_+^{L+1} \to \mathbb{R}_+\}_{t\in\mathbb{N}_0}$. We denote by Π the set of admissible policies. A policy π is stationary if $\pi_t(\mathbf{x}) = \pi_0(\mathbf{x})$ for all $t \in \mathbb{N}_0$ and $\mathbf{x} \in \mathbb{R}^{L+1}$. When a policy π is stationary, we omit the index t in π_t, for simplicity. We denote by $q_t(\pi)$ and $J_t(\pi)$ the random variables for the order quantity and inventory level respectively under policy π. We consider two stationary policies: the constant order policy C_r (cf. Xin and Goldberg, 2016) and the projected inventory level (PIL) policy P_ξ (cf. van Jaarsveld and Arts, 2024), where $r \in [0, \mu_D)$ is the constant order and $\xi \geq 0$ is the projected inventory level. For $t \in \mathbb{N}_0$, the constant order policy and PIL policy are expressed by:

$$C_r(\mathbf{x}_t) := r, \quad \text{and} \quad P_\xi(\mathbf{x}_t) := (\xi - \mathbb{E}[J_{t+L}|\mathbf{x}_t])^+.$$

Let $\{\{c_t(\pi)\}_{t\in\mathbb{N}_0}\}_{\pi\in\Pi}$ be the sequence of cost random variables given by:

$$c_t(\pi) := h\left(J_t(\pi) + q_t(\pi) - D_t\right)^+ + p\left(D_t - J_t(\pi) - q_t(\pi)\right)^+.$$

As a notational convenience, we define $D_{[a,b]} = \sum_{t=a}^{b} D_t$, and similarly define $J_{[a,b]}$, $q_{[a,b]}$, and $c_{[a,b]}(\pi)$. Accordingly, the cost-rate function $\mathcal{C} : \Pi \to \mathbb{R}_+$ is defined as:

$$\mathcal{C}(\pi) := \limsup_{T\to\infty} \mathbb{E}\left[\frac{c_{[L,T]}(\pi)}{T - L + 1}\right].$$

We will sometimes write the dependence of $\mathcal{C}(\pi)$ on p and L explicitly as $\mathcal{C}(\pi \mid p, L)$. Let $\mathcal{C}^*(p, L) := \inf_{\pi\in\Pi} \mathcal{C}(\pi \mid p, L)$ denote the optimal cost-rate. Huh et al. (2011) show that a stationary policy π^* exists such that $\mathcal{C}(\pi^*) = \mathcal{C}^*$. Throughout the chapter we say a function g is $o(f(x))$ and write $g(x) = o(f(x))$ if and only if $\lim_{x\to\infty} g(x)/f(x) = 0$.

4.2.2 Main result

In this section, we present the main result. For a fixed demand distribution F_D and h, we construct a sequence $\{\xi_p\}_{p\geq 0}$ such that $\xi_p \in \arg\min_{\xi\geq 0} \lim_{L\to\infty} \mathcal{C}(P_\xi \mid p, L)$, and a sequence $\{r_p\}_{p\geq 0}$ such that $r_p \in \arg\min_{r\in[0,\mu_D)} \lim_{L\to\infty} \mathcal{C}(C_r \mid p, L)$. We now state our main result.

Theorem 4.1. *1. There exists a constant $0 \leq M < \infty$ such that the optimality gap of the best PIL policy remains bounded by M for all $p \geq 0$ as L tends to infinity, that is*

$$\mathcal{C}\left(P_{\xi_p}\right) - \mathcal{C}(C_{r_p}) = \lim_{L\to\infty} \left(\mathcal{C}\left(P_{\xi_p}\right) - \mathcal{C}^*(p, L)\right) < M \quad \textit{for all } p \geq 0.$$

2. The PIL policy is asymptotically optimal for large p when L approaches infinity, that is,

$$\lim_{p\to\infty} \lim_{L\to\infty} \frac{\mathcal{C}\left(P_{\xi_p}\right)}{\mathcal{C}^*(p, L)} = 1.$$

Theorem 4.1 states that the cost-rate of the PIL policy exceeds that of the constant-order policy by at most a fixed constant M. The constant M depends on certain moments of a stochastic process induced by the inventory dynamics (see Section 4.G). Since the constant order policy is asymptotically optimal as the leadtime increases (Goldberg et al., 2016; Xin and Goldberg, 2016), the optimality gap of PIL policy remains bounded M under the same condition. This result extends our understanding of the asymptotic optimality of PIL policies beyond the special case of exponentially distributed demand addressed in Theorem 2 of van Jaarsveld and Arts (2024).

We note that Theorem 4.1 differs from the asymptotic optimality result for sufficiently large p presented in van Jaarsveld and Arts (2024) (Theorem 4) in terms of the asymptotic regime. The analysis in van Jaarsveld and Arts (2024) relies on the comparison of the PIL policy and the base-stock policy, which is not optimal as the leadtime approaches infinity. It is worth noting that, unlike the base-stock policy, whose optimality gap grows unboundedly with the lead time, the optimality gap of the PIL policy remains uniformly bounded.

4.3 Proof of Theorem 4.1

Assuming exponentially distributed demand, van Jaarsveld and Arts (2024) show that the relative value function of the constant-order policy is a parabola, and that a one-step policy improvement yields the PIL policy. This establishes the asymptotic optimality of the PIL policy as the leadtime grows, since it strictly improves upon the constant-order policy, which is itself asymptotically optimal (Goldberg et al., 2016). In this chapter, we extend this approach to general demand distributions with finite second moments by showing that the relative value function has a quadratic form and a term that grows sublinearly. Next we use that as the cost of losing a sale grows, the optimal constant order policy approaches a heavy traffic regime where this sublinear term turns out to be unimportant and a PIL policy will not be worse than a constant order policy within tight bounds.

Let Y be a random variable with distribution F_Y, defined as the difference between D and the constant order r, i.e. $Y := D - r$. It is straightforward to verify that F_Y is concentrated on $[-r, \infty)$, since $F_Y(x) = F_D(x + r)$, and $\mathbf{Var}[Y] = \mathbf{Var}[D] < \infty$. Let $\mu_Y := \mathbb{E}[Y]$ and $\sigma_Y^2 := \mathbf{Var}[Y]$.

Definition 4.2. For a constant order policy C_r, $r \in [0, \mu_D)$, the relative value function $v_r : \mathbb{R}_+ \to \mathbb{R}$ satisfies,

$$v_r(x) := \mathbb{E}_Y \left[h(x - Y)^+ + p(Y - x)^+ + v_r \left((x - Y)^+ \right) \right] - \mathcal{C}(C_r), \qquad v_r(0) = 0, \quad x \geq 0.$$

The difference $v_r(x_1) - v_r(x_2)$ represents the additional total long-run expected cost when the system starts from x_1 rather than x_2 under the constant order policy C_r (cf., Chapter 6

Tijms, 2003). Goldberg et al. (2016) show that $\mathcal{C}(C_r) = h\mathbb{E}[J_\infty] + p\mu_Y$, where J_∞ denotes the steady state inventory level under C_r, i.e., $\mathbb{P}\{J_\infty \leq x\} = \lim_{t\to\infty} \mathbb{P}\{J_t(C_r) \leq x\}$. By Definition 4.2, the relative value function $v_r(x)$ can be expressed as a convolution equation:

$$v_r(x) = a_r(x) + \int_{-r}^{x} v_r(x-y)F_Y(dy), \qquad v_r(0) = 0, \quad x \geq 0, \tag{4.2}$$

where,

$$a_r(x) := h\mathbb{E}_Y[(x-Y)^+] + p\mathbb{E}_Y[(Y-x)^+] - p\mu_Y - h\mathbb{E}[J_\infty]. \tag{4.3}$$

To simplify notation we introduce the convolution operator, $*$, as follows. Let $K : \mathbb{R} \to \mathbb{R}$ and $F : \mathbb{R} \to \mathbb{R}$ be two real functions. The convolution of K and F is defined as:

$$K * F(x) := \int_{-\infty}^{x} K(x-y)F(dy).$$

Therefore, Equation (4.2) can be rewritten as,

$$v_r(x) = a_r(x) + v_r * F_Y(x), \qquad v_r(0) = 0, \quad x \geq 0. \tag{4.4}$$

Deriving an explicit expression for $v_r(x)$ is non-trivial, as Equation (4.4) constitutes a Wiener-Hopf equation (cf. Asmussen, 1998). However, by analyzing a specific random walk with i.i.d. increments and its associated ladder processes in Section 4.3.1, we are able to derive an explicit solution in Section 4.3.2, which enables the remainder of our analysis.

4.3.1 Ladder processes

Consider a random walk $\{S_t := \sum_{i=1}^{t} Y_i\}_{t\in\mathbb{N}_0}$, with $S_0 = 0$, where $\{Y_t\}_{t\in\mathbb{N}_0}$ represents a sequence of random variables defined by $Y_t := D_t - r$. We introduce two stopping periods associated with the random walk S_t. The (weak) ascending ladder period, denoted by τ_+, is the first period (greater than zero) that the random walk attains a non-negative value, i.e., $\tau_+ := \inf\{t > 0 : S_t \geq 0\}$. The value of the stopped random walk at τ_+, i.e., S_{τ_+}, is a random variable known as the first (weak) ascending ladder height, with distribution function $G_+(x) = \mathbb{P}\{S_{\tau_+} \leq x\}$ supported on $[0, \infty)$. The mean and variance of S_{τ_+} are denoted by $\mu_+ := \mathbb{E}[S_{\tau_+}]$ and $\sigma_+^2 := \mathbf{Var}[S_{\tau_+}]$, respectively. Both μ_+ and σ_+ are finite for $r \in [0, \mu_D)$, and remain so as r approaches μ_D.

Lemma 4.3. *$\lim_{r\uparrow\mu_D} \mu_+$ and $\lim_{r\uparrow\mu_D} \sigma_+$ exist, and (a) $0 < \lim_{r\uparrow\mu_D} \mu_+ < \infty$, and (b) $\lim_{r\uparrow\mu_D} \sigma_+ < \infty$.*

Likewise, the (strict) descending ladder period is the first period (greater than zero) that the random walk takes a negative value, i.e., $\tau_- := \inf\{t > 0 : S_t < 0\}$. Accordingly, S_{τ_-}, is the first (strict) descending ladder height random variable with distribution

function $G_-(x) := \mathbb{P}\{S_{\tau_-} \leq x\}$ supported on $(-\infty, 0)$. We refer interested readers to Asmussen (2003) for a comprehensive overview of ladder processes. For any non-decreasing function $F : \mathbb{R} \to \mathbb{R}$, we let $\|F\| := \lim_{x\to\infty} F(x)$. A distribution, F, is called proper if $\|F\| = 1$ and defective if $\|F\| < 1$. There is a well-known result that G_+ is proper and G_- is defective since $\mu_Y > 0$ (cf. Theorem VIII 2.4. Asmussen, 2003). This implies that the probability that τ_- is finite cannot be one, i.e. $\lim_{x\to\infty} \mathbb{P}\{\tau_- < x\} < 1$, whereas $\tau_+ < \infty$ almost surely. Additionally, $\mathbb{E}[\tau_+] < \infty$, whereas $\mathbb{E}[\tau_-]$ is infinite. By Wald's identity (cf. Appendix A Tijms, 2003), μ_+ can be expressed as a function of $\mathbb{E}[\tau_+]$ and μ_Y:

$$\mu_+ = \mathbb{E}[S_{\tau_+}] = \mathbb{E}\left[\sum_{t=1}^{\tau_+} Y_t\right] = \mathbb{E}[\tau_+]\mathbb{E}[Y] = \mathbb{E}[\tau_+]\mu_Y.$$

Let m_n denote the partial minimum of the random walk within the first n periods, i.e. $m_n := \min_{0\leq t<n} S_t$. Then, the minimum of the entire random random walk, m, is defined as $m := \inf_{0\leq t<\infty} S_t$. Define the descending ladder height renewal measure $U_-(x) := \sum_{t=0}^{\infty} G_-^{*t}(x)$, where G_-^{*t} denotes the t-fold convolution of G_-, i.e., $G_-^{*t+1}(x) := G_-^{*t} * G_-(x)$, and $G_-^{*0}(x) = \delta_0(x)$, with δ_0 representing the probability measure degenerate at 0, i.e. $\delta_0(x) = 1$ if $x \geq 0$ and zero otherwise. We can express the distribution function of m as (cf. Theorem VIII, 2.2. Asmussen, 2003):

$$\mathbb{P}\{m \leq x\} = \frac{U_-(x)}{\|U_-\|}. \tag{4.5}$$

Next, J_t is distributed as the waiting time of the t-th customer of a GI/G/1 queue with inter-arrival distribution F_D and service time r. Thus, similar to Proposition, X.1.1. of Asmussen (2003) $J_\infty \stackrel{d}{=} -m$ ($\stackrel{d}{=}$ denotes equality in distribution) which implies by Equation (4.5) that:

$$\mathbb{E}[J_\infty] = -\mathbb{E}[m] = \frac{1}{\|U_-\|}\int_{-\infty}^{0} U_-(x)dx. \tag{4.6}$$

Similar to U_-, we define the ascending ladder height renewal measure, U_+, by $U_+ := \sum_{t=0}^{\infty} G_+^{*t}$.

4.3.2 Solution to the Wiener-Hopf equation

We build on the methodology developed by Asmussen (1998) to derive a solution to Equation (4.4). Asmussen (1998) shows that a solution to the Wiener-Hopf equation satisfies $v_r(x) = a_r * U_- * U_+(x)$. Using this fact leads, after multiple intricate steps, to the characterization of $v_r(x)$ in Theorem 4.4:

Theorem 4.4. *The relative value function characterized by Equation* (4.4) *is given by*

$$v_r(x) = \frac{h\mu_+}{\mu_D - r}\int_0^x U_+(y)dy - (h+p)x, \qquad x \geq 0.$$

Observe that the ascending ladder process is in fact a renewal process. As such, it possesses all the general properties of the renewal processes including the following lemma. Let $\kappa \in \mathbb{R}_+$ be expressed by

$$\kappa := \begin{cases} \dfrac{\sigma_+^2 + \mu_+^2}{2\mu_+^2} & \text{if } D \text{ is non-lattice,} \\ \dfrac{\sigma_+^2 + \mu_+^2 + \mu_+}{2\mu_+^2} & \text{if } D \text{ is lattice.} \end{cases}$$

Lemma 4.5. *The ascending ladder height renewal measure U_+ can be expressed as*

$$U_+(x) = \frac{1}{\mu_+}x + \kappa + g_r(x).$$

where $g_r : \mathbb{R}_+ \to \mathbb{R}$ satisfies $|g_r(x)| \leq \kappa$ for all $x \geq 0$, and $g_r(x) = o(1)$.

van Jaarsveld and Arts (2024) show that v_r is a quadratic function in the case of exponential demand. Next, we demonstrate that for a general demand distribution, v_r can be expressed as the sum of a quadratic function and an $o(x)$ term. This result holds under the sole mild assumption that the demand distribution has a finite second moment.

Theorem 4.6. *If $r \in [0, \mu_D)$ then for all $x \geq 0$,*

$$v_r(x) = b(r)\left(\left(x - \tilde{\xi}(r)\right)^2 - \tilde{\xi}^2(r) + 2\mu_+ \int_0^x g_r(y)dy\right),$$

with,

$$b(r) := \frac{h}{2(\mu_D - r)}, \quad \tilde{\xi}(r) := (\mu_D - r)\left(\frac{p}{h} + 1\right) - \kappa$$

and g_r as specified in Lemma 4.5.

Proof. This follows after some computation from Theorem 4.4 and Lemma 4.5. □

Indeed, g_r vanishes faster than $o(1)$ for most practical demand processes. For instance, it decays exponentially fast, i.e., $g_r(x) = o(e^{-\alpha x})$ with $\alpha > 0$, if D is non-lattice and sub-exponential, i.e., $\int_0^\infty e^{\delta x} G_+(dx) < \infty$, for some $\delta > 0$ (cf. VII Section 2. Asmussen, 2003). In this case, $v_r(x)$ is asymptotically quadratic as x grows large.

Next, suppose that Z is a non-negative random variable. We introduce a sufficient condition for the existence and finiteness of $\mathbb{E}[v_r(Z)]$.

Lemma 4.7. *Let Z have a finite second moment and $r \in [0, \mu_D)$, then $|\mathbb{E}[v_r(Z)]| < \infty$.*

4.3.3 Inventory dynamics

We next establish useful properties of constant order policies and PIL policies. Recall that $r_p \in [0, \mu_D)$ represents the best constant order quantity under a lost sales unit

penalty cost of $p \in \mathbb{R}_+$, given a fixed holding cost h, i.e., $r_p \in \arg\min_{r\in[0,\mu_D)} \mathcal{C}(C_r \mid p, L)$. As p increases, it is intuitive that r_p converges to μ_D to minimize the expected lost sales cost. In this case, the steady-state inventory level $J_\infty(C_{r_p})$ grows large as $\mu_Y \to 0$. Next, we provide a more detailed elaboration on this intuition. Consider the sequences of the steady state inventory levels $\{J_\infty(C_r)\}_{r\in[0,\mu_D)}$.

Lemma 4.8. *(a)* $\mathbb{E}[J_\infty(C_r)]$ *is non-decreasing and convex in* r*,*

(b) $2\mu_Y \mathbb{E}[J_\infty(C_r)]/\sigma_D^2 \to 1$ *as* $r \to \mu_D$*.*

(c) r_p *is non-decreasing in* $p \geq 0$*,*

(d) $r_p \to \mu_D$*, as* $p \to \infty$*,*

(e) $\sqrt{\frac{2p}{\sigma_D^2 h}}(\mu_D - r_p) \to 1$*, as* $p \to \infty$*.*

Note that part (e) of Lemma 4.8 implies that $\lim_{p\to\infty}(r_p - \mu_D)/\sqrt{p} \in (0, \infty)$, that is, r_p approaches μ_D at the same rate as $1/\sqrt{p}$ approaches 0. Bu et al. (2020) have also established the same result in a more general setting with stochastic supply yield. Our setting allows for a shorter proof that we provide to make our argument self-contained.

We now shift our attention to the dynamics of the inventory level under PIL policies. Let $\{\{q_t(P_\xi)\}_{t\in\mathbb{N}_0}\}_{\xi\geq 0}$ be a sequence of random variables representing orders under PIL policies $\{P_\xi\}_{\xi\geq 0}$, where $\{q_t(P_\xi)\}_{t=0,\ldots,L-1}$ are fixed for all $\xi \geq 0$ and known almost surely. Let $\{\{J_t(P_\xi)\}_{t\in\mathbb{N}_0}\}_{\xi\geq 0}$ be the corresponding sequence of inventory level random variables.

Lemma 4.9. *For all* $t \geq L+1$ *and* $\xi \geq 0$*, the order size* $q_t(P_\xi)$ *satisfy:* $\mathbb{E}[q_t(P_\xi)] \leq \min\{\xi, \mu_D\}$.

The result from Lemma 4.9 leads to Lemma 4.10 which will later be used to show that the impact of g_r on the optimality gap of the PIL policy remains bounded for any p and large L. This is the sense in which the term of $v_r(x)$ that grows sublinearly becomes unimportant as p grows.

Lemma 4.10. *There exists* $0 \leq M < \infty$ *such that for all* $t \geq L+1$*,* $r \in [0, \mu_D)$*, and* $\xi \geq 0$*,*

$$b(r)\mu_+ \mathbb{E}\left[\int_{J_t(P_\xi)}^{J_t(P_\xi)+q_t(P_\xi)-r} g_r(y)dy\right] < M.$$

Next, we define the projected inventory level $\xi(r)$, $r \in [0, \mu_D)$ by

$$\xi(r) := \tilde{\xi}(r) + r = \frac{\mu_Y p}{h} + \mu_D - \kappa. \tag{4.7}$$

Notice that, by Lemma 4.8(e), in combination with Lemma 4.3 and Equation (4.7), we have

$$0 < \xi(r_p)/\sqrt{p} < \infty, \qquad \text{as } p \to \infty. \tag{4.8}$$

That is, $\xi(r_p)$ goes to infinity in the order of $\sqrt{p}$, as $p \to \infty$. The next section provides a more detailed analysis of the cost-rate difference between C_{r_p} and $P_{\xi(r_p)}$.

4.3.4 Cost-rate difference between PIL and constant order policy

Next, we derive an upper bound on the cost rate of a family of PIL policies by comparing it to that of corresponding constant order policies. One classical way of comparing the performance of two policies is by using the improvement theorem (cf. Theorem 6.2.1. Tijms, 2003). In general, applying the improvement theorem to our problem requires the consideration of $L+1$-dimensional state space. However, the state space collapses to one dimensional for a system under a constant order policy since all order quantities as identical. Lemma 4.11 adapts the improvement theorem for a constant order policy.

Lemma 4.11. *(Similar to Lemma 4 of van Jaarsveld and Arts, 2024) Let $t_1 \leq t_2$, $t_1, t_2 \in \mathbb{N}$ and suppose $q_t = r$ for all $t \in \{t_1, \ldots, t_2\}$. Then,*

$$\mathbb{E}[c_{[t_1,t_2]}(C_r)|J_{t_1}] = v_r(J_{t_1}) - \mathbb{E}[v_r(J_{t_2+1})|J_{t_1}] + (t_2 + 1 - t_1)\mathcal{C}(C_r).$$

We are now in the position to prove the main results.

Theorem 4.12. *There exists $0 \leq M < \infty$ such that $\mathcal{C}(P_{\xi_p}) \leq \mathcal{C}(P_{\xi(r_p)}) \leq \mathcal{C}(C_{r_p}) + M$ for every $p \geq 0$.*

Proof of Theorem 4.12. In this proof we bound $\mathbb{E}[c[L,T]\left(P_{\xi(r_p)}\right) - c[L,T]\left(C_{r_p}\right)]$ for $p \geq 0$. Similar to van Jaarsveld and Arts (2024), the proof relies on a policy $\mathcal{P}^{\tilde{t}}$, $\tilde{t} \in \mathbb{N}_0$, which uses the PIL policy $P_{\tilde{\xi}(r_p)}$ to order for $t = 1, \ldots, \tilde{t} + L$, and then orders r_p when $t \geq \tilde{t} + L + 1$, that is,

$$\mathcal{P}_t^{\tilde{t}}(\mathbf{x}) = \begin{cases} P_{\xi(r_p)}(\mathbf{x}), & t = 1, \ldots, \tilde{t} + L, \\ r_p, & t = \tilde{t} + L + 1, \ldots \end{cases}$$

Then,

$$\mathbb{E}[c_{[L,T]}(\mathcal{P}^{\tilde{t}}) - c_{[L,T]}(\mathcal{P}^{\tilde{t}-1})] = \mathbb{E}[c_{[\tilde{t}+L,T]}(\mathcal{P}^{\tilde{t}}) - c_{[\tilde{t}+L,T]}(\mathcal{P}^{\tilde{t}-1})] =$$
$$\mathbb{E}\left[\mathbb{E}\left[c_{[\tilde{t}+L,T]}(C(r_p))|J_{\tilde{t}+L} = J_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}}) + q_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}}) - r_p\right] - \mathbb{E}\left[c_{[\tilde{t}+L,T]}(C(r_p))|J_{\tilde{t}+L} = J_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}})\right]\right]. \tag{4.9}$$

The first equality in (4.9) holds because c_t remains the same under $\mathcal{P}^{\tilde{t}-1}$ and $\mathcal{P}^{\tilde{t}}$ for $t \leq \tilde{t} + L - 1$. To justify the second equality, first observe that $J_{\tilde{t}+L-1}$ remains unchanged under $\mathcal{P}^{\tilde{t}-1}$ and $\mathcal{P}^{\tilde{t}}$ due to the dynamics of the inventory levels. Second, observe that under both policies the system receives r_p in periods $t > \tilde{t} + L$. Third, observe that a system initiated at $J_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}})$ and receiving the order quantity $q_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}})$ is equivalent to one starting at $J_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}}) + q_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}}) - r_p$ and receiving an order quantity r_p. Thus, the second equality compares the total cost of two systems under C_{r_p} with different initial

inventory levels. Next, we use Lemma 4.11 to expand Equation (4.9) as follows

$$\mathbb{E}\left[c_{[L,T]}(\mathcal{P}^{\tilde{t}}) - c_{[L,T]}(\mathcal{P}^{\tilde{t}-1})\right] = \mathbb{E}\Big[v_{r_p}\left(J_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}}) + q_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}}) - r_p\right) - v_{r_p}\left(J_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}+1})\right) - v_{r_p}\left(J_{T+1}(\mathcal{P}^{\tilde{t}})\right) + v_{r_p}\left(J_{T+1}(\mathcal{P}^{\tilde{t}})\right)\Big].$$

Let $\mathcal{P}^{-1} := C_{r_p}$. We use a telescopic sum,

$$\begin{aligned}\mathbb{E}\left[c_{[L,T]}\left(P_{\xi(r_p)}\right) - c_{[L,T]}\left(C_{r_p}\right)\right] =& \mathbb{E}\left[c_{[L,T]}\left(\mathcal{P}^{T}\right) - c_{[L,T]}\left(\mathcal{P}^{-1}\right)\right] \\ =& \sum_{\tilde{t}=0}^{T-L} \mathbb{E}[c_{[L,T]}(\mathcal{P}^{\tilde{t}}) - c_{[L,T]}(\mathcal{P}^{\tilde{t}-1})] \\ =& \sum_{\tilde{t}=0}^{T-L} \mathbb{E}[v_{r_p}(J_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}}) + q_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}}) - r_p) - v_{r_p}(J_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}}))] - \\ & \mathbb{E}[v_{r_p}(J_{T+1}(\mathcal{P}^{\tilde{t}}))] + \mathbb{E}[v_{r_p}(J_{T+1}(\mathcal{P}^{\tilde{t}-1}))]. \end{aligned} \tag{4.10}$$

By Theorem 4.6 we notice that for any $r \in [0, \mu_D)$, $\mathbf{x}_t \in \mathbb{R}_+^{L+1}$, $t \geq 0$, $\tilde{q} \in \mathbb{R}$

$$\begin{aligned}&\mathbb{E}\left[v_r(J_{t+L} + \tilde{q})|\mathbf{x}_t\right] = b(r)\mathbb{E}\left[\left(J_{t+L} + \tilde{q} - \tilde{\xi}(r)\right)^2 - \tilde{\xi}^2(r) + 2\mu_+ \int_0^{J_{t+L}+\tilde{q}} g_r(y)dy|\mathbf{x}_t\right] = \\ &b(r)\left(\mathbf{Var}[J_{t+L}|\mathbf{x}_t] + \left(\mathbb{E}[J_{t+L}|\mathbf{x}_t] + \tilde{q} - \tilde{\xi}(r)\right)^2 - \tilde{\xi}^2(r) + 2\mu_+\mathbb{E}\left[\int_0^{J_{t+L}+\tilde{q}} g_r(y)dy|\mathbf{x}_t\right]\right).\end{aligned} \tag{4.11}$$

Using Equation (4.10) combined with (4.11) and some algebra we have

$$\begin{aligned}&\mathbb{E}[c_{[L,T]}(P_{\xi(r_p)}) - c_{[L,T]}(C_{r_p})] = \mathbb{E}[v_{r_p}(J_{T+1}(P_{\xi(r_p)}))] - \mathbb{E}[v_{r_p}(J_{T+1}(C_{r_p}))] + \\ &\sum_{\tilde{t}=0}^{T-L} b(r_p)\mathbb{E}\left[-(P_{\xi(r_p)}(\mathbf{x}_{\tilde{t}}) - r_p)^2 + 2\mu_+\mathbb{E}\left[\int_{J_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}})}^{J_{\tilde{t}+L}(\mathcal{P}^{\tilde{t}})+P_{\xi(r_p)}(\mathbf{x}_{\tilde{t}})-r_p} g_{r_p}(y)dy\Big|\mathbf{x}_{\tilde{t}}\right]\right].\end{aligned}$$

Then, it follows from Lemma 4.10 that there exists $0 \leq M < \infty$ such that for all $p \geq 0$

$$\begin{aligned}&\mathbb{E}[c_{[L,T]}(P_{\xi(r_p)}) - c_{[L,T]}(C_{r_p})] < \\ &\mathbb{E}\left[v_{r_p}(J_{T+1}(P_{\xi(r_p)}))\right] - \mathbb{E}\left[v_{r_p}(J_{T+1}(C_{r_p}))\right] + (T-L+1)M - b(r_p)\sum_{\tilde{t}=0}^{T-L}\mathbb{E}\left[(P_{\xi(r_p)}(\mathbf{x}_{\tilde{t}}) - r_p)^2\right] \\ &\leq \mathbb{E}\left[v_{r_p}(J_{T+1}(P_{\xi(r_p)}))\right] - \mathbb{E}[v_{r_p}(J_{T+1}(C_{r_p}))\big] + (T-L+1)M.\end{aligned} \tag{4.12}$$

Notice that the last inequality of (4.12) holds since $b(r_p)$ and $(P_{\xi(r_p)}(\mathbf{x}_{\tilde{t}}) - r_p)^2$ are non-negative. We use (4.12) to find an upper bound on the cost-rate of the PIL policy, i.e.,

$\mathcal{C}(P_{\xi(r_p)})$ with respect to the cost-rate of constant order policy $\mathcal{C}(C_{r_p})$ when $p \geq 0$:

$$\begin{aligned}
C(P_{\xi(r_p)}) =& \limsup_{T\to\infty} \frac{1}{T-L+1}\mathbb{E}[c_{[L,T]}(P_{\xi(r_p)})] \\
<& \limsup_{T\to\infty} \frac{1}{T-L+1}\left(\mathbb{E}[c_{[L,T]}(C_{r_p}) + v_{r_p}(J_{T+1}(P_{\xi(r_p)})) - v_{r_p}(J_{T+1}(C_{r_p}))] + (T-L+1)M\right) \\
=& C(C_{r_p}) + M + \limsup_{T\to\infty} \frac{1}{T-L+1}\left(\mathbb{E}[v_{r_p}(J_{T+1}(P_{\xi(r_p)}))] - \mathbb{E}\left[v_{r_p}(J_{T+1}(C_{r_p}))\right]\right) \\
=& C(C_{r_p}) + M \qquad (4.13)
\end{aligned}$$

The last equality holds since both $\mathbb{E}[v_{r_p}(J_{T+1}(P_{\xi(r_p)}))$ and $\mathbb{E}[v_{r_p}(J_{T+1}(C_{r_p}))]$ remain finite as $T \to \infty$. First, observe, as van Jaarsveld and Arts (2024) do, that $0 \leq J_{T+1}(P_{\xi(r_p)}) \leq \xi(r_p) + L\mu_D$ for all $T \geq L$ which ensures that $J_{T+1}(P_{\xi(r_p)})$ has finite first and second moments. Then it follows from Lemma 4.7 that $|\mathbb{E}[v_{r_p}(J_{T+1}(P_{\xi(r_p)}))| < \infty$. Second, $J_{T+1}(C_{r_p})$ converges to the steady state distribution of the inventory level under the constant order policy C_{r_p}, i.e., J_∞, as $T \to \infty$. We note that J_∞ has a finite first moment because D has a finite second moment. Additionally, J_∞ has a finite second moment since $0 \leq ((r_p - D)^+)^3 \leq r_p^3$, implying that $\mathbb{E}\left[((r_p - D)^+)^3\right] < \infty$ (cf. Theorem X. 2.1. Asmussen, 2003). Hence, $|\mathbb{E}[v_{r_p}(J_{T+1}(C_{r_p}))]| < \infty$ due to Lemma 4.7. The optimality of ξ_p, i.e., $C(P_{\xi_p}) \leq C(P_{\xi(r_p)})$ together with Inequality (4.13) complete the proof. □

Proof of Theorem 4.1. Combining Theorem 4.12 with asymptotic optimality of the constant order policy as L approaches infinity (Goldberg et al., 2016; Xin and Goldberg, 2016) provides the result: There exists $0 \leq M < \infty$ such that for all $p \geq 0$:

$$\lim_{L\to\infty}\left(\mathcal{C}(P_{\xi_p}) - \mathcal{C}^*(p,L)\right) \leq \lim_{L\to\infty}\left(\mathcal{C}(P_{\xi(r_p)}) - \mathcal{C}^*(p,L)\right) \leq M.$$

Observe that $\lim_{p\to\infty}\lim_{L\to\infty}\mathcal{C}^*(p,L) = \infty$. Thus,

$$\lim_{p\to\infty}\lim_{L\to\infty}\frac{\mathcal{C}(P_{\xi_p})}{\mathcal{C}^*(p,L)} = 1. \quad \square$$

4.4 Concluding remarks

In this chapter, we proved that the PIL policy is asymptotically optimal for sufficiently large lost sales unit costs as the leadtime approaches infinity, under mild assumptions on the i.i.d. demand process. This result, combined with van Jaarsveld and Arts (2024), demonstrates that the PIL policy is asymptotically optimal when the lost sales penalty cost is large, both in the case of a small leadtime and when the leadtime grows at a rate faster than the unit cost of lost sales. This makes the PIL policy the only single-parameter policy that guarantees optimality in both regimes under a general i.i.d. demand. It remains an open question whether the PIL policy is asymptotically optimal when both the leadtime and the lost sales unit penalty cost grow at the same rate. To the best of our knowledge, no simple policies are known to achieve optimality in this asymptotic regime.

4.A Proof of Lemma 4.3

Part (a) follows from Theorem XVIII.5.1. Feller (1991). The rest is the proof of Part (b). Observe that $\mathbb{E}[Y^2] < \infty$ only if for some $\alpha > 2$,

$$1 - F_Y(x) = O(x^{-\alpha}) \quad \text{as } x \to \infty.$$

This condition is equivalent to

$$\begin{aligned}\lim_{x\to\infty} \mathbb{E}\left[Y^2 | Y \geq x\right] &= \lim_{x\to\infty} \frac{\int_x^\infty y^2 F_Y(dy)}{1 - F_Y(x)} \\ &= \lim_{x\to\infty} \frac{x^2(1 - F_Y(x)) + 2\int_x^\infty y(1 - F_Y(y))dy}{1 - F_Y(x)} < \infty. \end{aligned} \tag{4.14}$$

Since for any $x \in (0, \infty)$, $\mathbb{E}\left[Y^2 | Y \geq x\right] < \infty$, and it is finite at the limit $x \to \infty$ by (4.14), we conclude that,

$$\sup_{x\geq 0} \mathbb{E}\left[Y^2 | Y \geq x\right] < \infty. \tag{4.15}$$

Next, we notice that $S_{\tau_+} = S_{\tau_+ - 1} + Y_{\tau_+} \stackrel{d}{=} S_{\tau_+ - 1} + Y | Y \geq -S_{\tau_+ - 1}$, and $S_{\tau_+ - 1} < 0$ by the definition of τ_+. This implies in particular that $S_{\tau_+} \leq Y_{\tau_+}$ almost surely and

$$\mathbb{E}\left[S_{\tau_+}^2\right] \leq \mathbb{E}[Y_{\tau_+}^2] = \mathbb{E}\left[\mathbb{E}\left[Y^2 | Y \geq -S_{\tau_+ - 1}\right]\right] \leq \sup_{x>0} \mathbb{E}\left[Y^2 | Y \geq x\right] < \infty.$$

Finally $\mathbb{E}[D^2] < \infty$ is equivalent to $\mathbb{E}[Y^2] < \infty$ which completes the proof. □

4.B Proof of Theorem 4.4

We use the methodology of solving Wiener-Hopf equations introduced by Asmussen (1998). A key distinction between our approach and that of Asmussen (1998) lies in the class of admissible solutions: While Asmussen (1998) restricts attention to non-negative solutions, we allow for all possible solutions, including non-positive ones. We use the following lemma to solve Equation (4.4) under this general class of admissible solutions.

Lemma 4.13. *(Corollary 3.1 and Proposition 3.3 of Asmussen, 1998)*

$$v_r(x) = a_r * U_- * U_+(x), \qquad v_r(0) = 0, \forall x \geq 0.$$

Lemma 4.13 provides a powerful approach for solving the Wiener-Hopf equation (4.4). Applying Lemma 4.13 to derive $v_r(x)$ is intricate and involves multiple steps. The proof of Theorem 4.4 is provided at the end of this section. The first step in deriving $v_r(x)$, following Lemma 4.13, involves expressing $a_r(x)$ in terms of $F_Y(x)$. This step is necessary due to the lack of a standard result in the literature that allows direct convolution of a_r in Equation (4.3) with U_-. However, as we will later show, existing results from random walk theory enable the convolution of F_Y with both U_- and U_+.

Lemma 4.14.

$$a_r(x) = (h+p)\int_{-r}^{x} F_Y(y)dy - px - h\mathbb{E}[J_\infty].$$

Proof of Lemma 4.14. By Equation (4.3):

$$\begin{aligned} a_r(x) &= h\mathbb{E}_Y[(x-Y)^+] + p\mathbb{E}_Y[(Y-x)^+] - p\mu_Y + h\mathbb{E}[J_\infty] \\ &= (h+p)\mathbb{E}_Y[(x-Y)^+] - px + p\mu_Y - p\mu_Y + h\mathbb{E}[J_\infty]. \end{aligned}$$

Now, we express $\mathbb{E}_Y[(x-Y)^+]$ in terms of $F_Y(x)$ as follows:

$$\mathbb{E}_Y[(x-Y)^+] = \int_{-r}^{x}(x-y)F_Y(dy) = xF_Y(y)\Big|_{-r}^{x} - \int_{-r}^{x} yF_Y(dy).$$

By assumptions, $F_Y(-r) = 0$. Additionally, we use integration by parts to compute $\int_{-r}^{x} yF_Y(dy)$:

$$\int_{-r}^{x} yF_Y(dy) = yF_Y(y)\Big|_{-r}^{x} - \int_{-r}^{x} F_Y(y)dy.$$

Thus

$$\mathbb{E}_Y[(x-Y)^+] = \int_{-r}^{x} F_Y(y)dy,$$

and

$$a_r(x) = (h+p)\int_{-r}^{x} F_Y(y)dy - px - h\mathbb{E}[J_\infty]. \quad \square$$

By Lemma 4.14, a_r is expressed as a linear combination of $\int_{-r}^{x} F_Y(y)dy$, x, and the constant 1. Importantly, the convolution operator possesses both distributive and homogeneous properties. These properties enable the separate convolution of $\int_{-r}^{x} F_Y(y)dy$, x, and 1 with U_- and U_+, providing the basis for the proof of Theorem 4.4.

Derivation of $\int_{-\infty}^{x} F_Y(y)dy * U_- * U_+(x)$: The convolution operator satisfies the associativity property. Furthermore, the following well-known lemma indicates the relation between the integration and convolution operators.

Associativity and commutativity of convolution imply that:

$$\int_{-\infty}^{x} F_Y(y)dy * U_- * U_+(x) = \int_{-\infty}^{x} F_Y * U_-(y)dy * U_+(x). \tag{4.16}$$

By Equation (4.16), the next steps involve first calculating $F_Y * U_-$, then convolving the result with U_+, and finally integrating the outcome. The following lemma is crucial to our computations.

Lemma 4.15. *(Theorem VIII 3.1. and Corollary 3.2 Asmussen, 2003)*

$$U_- * F_Y = U_- + G_+ - \delta_0,$$

It follows from commutativity of convolution and Lemma 4.15 that:

$$\int_{-\infty}^{x} U_- * F_Y(y)dy * U_+(x) = \int_{-\infty}^{x} (U_- + G_+ - \delta_0)(y)dy * U_+(x) = \int_{-\infty}^{x} (G_+ - \delta_0) * U_+(y)dy + \int_{-\infty}^{x} U_-(y)dy * U_+(x). \tag{4.17}$$

By definition of U_+:

$$G_+ * U_+ = G_+ * \sum_{t=0}^{\infty} G_+^{*t} = \sum_{t=1}^{\infty} G_+^{*t} = U_+ - \delta_0. \tag{4.18}$$

Furthermore, it is a well-known result that the convolution of any function with δ_0 returns the same function. Consequently, the first term of Equation (4.17) can be calculated as follows:

$$\int_{-\infty}^{x} (G_+ - \delta_0) * U_+(y)dy = \int_{-\infty}^{x} (U_+ - \delta_0 - U_+)dy = -\int_{0}^{x} \delta_0 dy = -x. \tag{4.19}$$

Now we address the second term of Equation (4.17),

$$\int_{-\infty}^{x} U_-(y)dy * U_+(x) = \int_{-\infty}^{0} U_-(y)dy * U_+(x) + \int_{0}^{x} U_- * U_+(y)dy.$$

We notice that for all $x \geq 0$, $U_-(x) = \|U_-\|$. Additionally, $1 * U_+ = U_+$, since for all $x \leq 0$, $U_+(x) = 0$. Thus, by Equation (4.6):

$$\int_{-\infty}^{x} U_-(y)dy * U_+(x) = \|U_-\|\mathbb{E}[J_\infty]U_+ + \|U_-\| \int_{0}^{x} U_+(y)dy. \tag{4.20}$$

The following lemma allows us to relate Equation (4.20) to $\mathbb{E}[\tau_+]$.

Lemma 4.16. *(Theorem VIII 2.3. (c) Asmussen, 2003)*

$$\|U_-\| = \mathbb{E}[\tau_+] = (1 - \|G_-\|)^{-1}.$$

By Lemma 4.16 and Equation (4.20), we can compute the second term of Equation (4.17):

$$\int_{-\infty}^{x} U_-(y)dy * U_+(x) = \mathbb{E}[\tau_+]\mathbb{E}[J_\infty]U_+ + \mathbb{E}[\tau_+] \int_{0}^{x} U_+(y)dy. \tag{4.21}$$

We combine Equations (4.17), (4.19), and (4.21) to compute $\int_{-\infty}^{x} F_Y(y)dy * U_- * U_+(x)$:

$$\int_{-\infty}^{x} F_Y(y)dy * U_- * U_+(x) = \mathbb{E}[\tau_+] \int_{0}^{x} U_+(y)dy + \mathbb{E}[\tau_+]\mathbb{E}[J_\infty]U_+ - x. \tag{4.22}$$

Derivation of $x*U_-*U_+(x)$: It is straightforward to verify that $x*U_-(x) = \int_{-\infty}^{x} U_-(y)dy$, given the definition and commutativity of the convolution operator. Thus, by Equation (4.21):

$$x * U_- * U_+(x) = \int_{-\infty}^{x} U_-(y)dy * U_+(x) = \mathbb{E}[\tau_+] \int_{0}^{x} U_+(y)dy + \mathbb{E}[\tau_+]\mathbb{E}[J_\infty]U_+. \tag{4.23}$$

Derivation of $1 * U_- * U_+(x)$: By definition of the convolution operator:

$$1 * U_-(x) = \int_{-\infty}^{x} U_-(dy) = U_-(x) = \|U_-\| = \mathbb{E}[\tau_+],$$

which implies that:

$$1 * U_- * U_+(x) = \mathbb{E}[\tau_+]U_+. \tag{4.24}$$

At this point we have all the tools available to prove Theorem 4.4.

Proof of Theorem 4.4. By Lemma 4.13 and Lemma 4.14, for $x \geq 0$, $v_r(x)$ can be calculated by:

$$\begin{aligned} v_r(x) &= a_r * U_- * U_+(x) = \left((h+p)\int_{-r}^{x} F_Y(y)dy - px - h\mathbb{E}[J_\infty]\right) * U_- * U_+(x), \\ &= (h+p)\int_{-r}^{x} F_Y(y)dy * U_- * U_+(x) - px * U_- * U_+(x) - h\mathbb{E}[J_\infty]1 * U_- * U_+(x). \end{aligned}$$

By Equations (4.22), (4.23), and (4.24):

$$\begin{aligned} v_r(x) = (h+p)\left(\mathbb{E}[\tau_+]\int_0^x U_+(y)dy + \mathbb{E}[\tau_+]\mathbb{E}[J_\infty]U_+ - x\right) + \\ - p\left(\mathbb{E}[\tau_+]\int_0^x U_+(y)dy + \mathbb{E}[\tau_+]\mathbb{E}[J_\infty]U_+\right) - h\mathbb{E}[J_\infty]\mathbb{E}[\tau_+]U_+. \end{aligned}$$

Simplifying the last expression, we can calculate $v_r(x)$ as follows:

$$v_r(x) = h\mathbb{E}[\tau_+]\int_0^x U_+(y)dy - (h+p)x.$$

By Wald's equality $\mathbb{E}[\tau_+]\mu_Y = \mathbb{E}[S_{\tau_+}] = \mu_+$, since τ_+ is a stopping time for the $\{S_t\}_{t\in\mathbb{N}}$ process. Hence:

$$v_r(x) = \frac{h\mu_+}{\mu_Y}\int_0^x U_+(y)dy - (h+p)x. \quad \square$$

4.C Proof of Lemma 4.5

For this proof we need two observations. First, for all $x \geq 0$:

$$\frac{1}{\mu_+}x \leq U_+(x) \leq \frac{1}{\mu_+}x + \kappa. \tag{4.25}$$

The left inequality of (4.25) deals with the fact that the expected time until the next renewal after x (residual life) is non-negative (cf. V. 6. Asmussen, 2003). The right inequality of (4.25) is Lorden's Inequality (Lorden, 1970). Next, as $x \to \infty$,

$$U_+(x) = \frac{1}{\mu_+}x + \kappa + o(1). \tag{4.26}$$

Equation (4.26) is due to the asymptotic expansion of the expected residual life function (cf. Proposition V 6.1. Asmussen, 2003, for non-lattice D). (4.25) together with (4.26) provide the result. $\square$

4.D Proof of Lemma 4.7

By Theorem 4.6,

$$\left(Z-\tilde{\xi}(r)\right)^2-\kappa\leq\frac{1}{b(r)}\left(v_r(Z)+\tilde{\xi}^2(r)\right)\leq\left(Z-\tilde{\xi}(r)\right)^2+\kappa.$$

We take the expectation with respect to Z on all sides,

$$\mathbb{E}\left[\left(Z-\tilde{\xi}(r)\right)^2\right]-\kappa\leq\frac{1}{b(r)}\left(\mathbb{E}[v_r(Z)]+\tilde{\xi}^2(r)\right)\leq\mathbb{E}\left[\left(Z-\tilde{\xi}(r)\right)^2\right]+\kappa.$$

Observe that by definition, $\mathbf{Var}\left[Z-\tilde{\xi}(r)\right]=\mathbb{E}\left[\left(Z-\tilde{\xi}(r)\right)^2\right]-\left(\mathbb{E}[Z]-\tilde{\xi}(r)\right)^2$. It follows that,

$$\mathbf{Var}[Z]+\left(\mathbb{E}[Z]-\tilde{\xi}(r)\right)^2-\kappa\leq\frac{1}{b(r)}\left(\mathbb{E}[v_r(Z)]+\tilde{\xi}^2(r)\right)\leq\mathbf{Var}[Z]+\left(\mathbb{E}[Z]-\tilde{\xi}(r)\right)^2+\kappa.$$

Notice that Z has finite first and second moments and $0<\mu_+,\sigma_+<\infty$ for $r\in[0,\mu_D)$, which implies the result. □

4.E Proof of Lemma 4.8

Consider the sequences of random variables $\{\{Y_t(r)=D_t-r\}_{t\in\mathbb{N}}\}_{r\in[0,\mu_D)}$, sequences of random walks $\{\{S_t(r)=\sum_{i=1}^t Y_t(r)\}_{t\in\mathbb{N}}\}_{r\in[0,\mu_D)}$.

(a) We recall that $J_\infty\overset{d}{=}-m$. It is a known result (cf. Proposition VIII 4.5 Asmussen, 2003) that,

$$\mathbb{E}[J_\infty(C_r)]=\sum_{t=1}^\infty\frac{1}{t}\mathbb{E}[S_t^-]=\sum_{t=1}^\infty\frac{1}{t}\mathbb{E}[(-S_t)^+]=\sum_{t=1}^\infty\frac{1}{t}\mathbb{E}\left[\left(tr-\sum_{i=1}^t D_i\right)^+\right].$$

Let $r_1,r_2\in[0,\mu_D)$ and $r_1\leq r_2$. First we prove monotonicity. Observe that

$$tr_1-\sum_{i=1}^t D_i\leq tr_2-\sum_{i=1}^t D_i,$$

almost surely and so

$$\mathbb{E}\left[(tr_1-\sum_{i=1}^t D_i)^+\right]\leq\mathbb{E}\left[(tr_2-\sum_{i=1}^t D_i)^+\right].$$

Hence,

$$\mathbb{E}[J_\infty(C_{r_1})]=\sum_{t=1}^\infty\frac{1}{t}\mathbb{E}\left[\left(tr_1-\sum_{i=1}^t D_i\right)^+\right]\leq\sum_{t=1}^\infty\frac{1}{t}\mathbb{E}\left[\left(tr_2-\sum_{i=1}^t D_i\right)^+\right]=\mathbb{E}[J_\infty(C_{r_2})].$$

Next, we prove convexity. For all $0 \leq \alpha \leq 1$,

$$\left(t(\alpha r_1 + (1-\alpha)r_2) - \sum_{i=1}^{t} D_i\right)^+ = \left(\alpha\left(tr_1 - \sum_{i=1}^{t} D_i\right) + (1-\alpha)\left(tr_2 - \sum_{i=1}^{t} D_i\right)\right)^+$$
$$\leq \alpha\left(tr_1 - \sum_{i=1}^{t} D_i\right)^+ + (1-\alpha)\left(tr_2 - \sum_{i=1}^{t} D_i\right)^+, \qquad \text{almost surely.}$$

Hence,

$$2\mathbb{E}\left[\left(t(\alpha r_1 + (1-\alpha)r_2) - \sum_{i=1}^{t} D_i\right)^+\right] \leq$$
$$\alpha\mathbb{E}\left[\left(tr_1 - \sum_{i=1}^{t} D_i\right)^+\right] + (1-\alpha)\mathbb{E}\left[\left(tr_2 - \sum_{i=1}^{t} D_i\right)^+\right],$$

which gives,

$$\mathbb{E}[J_\infty(C_{\alpha r_1+(1-\alpha)r_2})] = \sum_{t=1}^{\infty}\mathbb{E}\left[\left(t(\alpha r_1 + (1-\alpha)r_2) - \sum_{i=1}^{t} D_i\right)^+\right]$$
$$\leq \alpha\sum_{t=1}^{\infty}\mathbb{E}\left[\left(tr_1 - \sum_{i=1}^{t} D_i\right)^+\right] + (1-\alpha)\sum_{t=1}^{\infty}\mathbb{E}\left[\left(tr_2 - \sum_{i=1}^{t} D_i\right)^+\right]$$
$$= \alpha\mathbb{E}[J_\infty(r_1)] + (1-\alpha)\mathbb{E}[J_\infty(r_2)].$$

(b) Part (b) presents the expected waiting time of a GI/G/1 queue in a heavy traffic condition. Interested readers may refer to Kingman (1961).

(c) Recall that $\mathcal{C}(C_r) = h\mathbb{E}[J_\infty(r)] + p(\mu_D - r)$. By part (a), $\mathcal{C}(C_r)$ is convex in r. Let $\partial\mathcal{C}(C_r)$ denote the sub-differential of the cost-rate function at r, that is:

$$\partial\mathcal{C}(C_r) := \{x \in \mathbb{R} : \mathcal{C}(C_{\bar{r}}) - \mathcal{C}(C_r) \geq x(\bar{r} - r), \forall \bar{r} \geq 0\}.$$

By the optimality condition $0 \in \partial\mathcal{C}(C_{r_p})$ which is equivalent to $\frac{p}{h} \in \partial\mathbb{E}[J_\infty(C_{r_p})]$, $p \geq 0$. It is straightforward to verify that $\partial\mathbb{E}[J_\infty(C_r)]$, $r \in [0, \mu_D)$ is an interval $[a_r, b_r]$ where a_r, b_r are some non-negative real numbers due to part (a). Additionally, for any $0 \leq r_1 \leq r_2 < \infty$, $b_{r_1} \leq a_{r_2}$ due to the convexity of $\mathbb{E}[J_\infty(C_r)]$. This implies that for $p_1 \leq p_2$, $r_{p_1} \leq r_{p_2}$, since $p_i/h \in \partial\mathbb{E}\left[J_\infty\left(C_{r_{p_i}}\right)\right]$, for $i \in \{1, 2\}$, and either $b_{r_{p_1}} \leq a_{r_{p_2}}$ or $b_{r_{p_2}} \leq a_{r_{p_1}}$.

(d) Next we prove that r_p approaches μ_D as $p \to \infty$. This statement is equivalent to showing that there exists no $0 \leq \tilde{r} < \mu_D$ such that for some $\tilde{p} \geq 0$, $r_p \leq \tilde{r}$ for all $p \geq \tilde{p}$, considering part (c). Assume the contrary that there exist such $\tilde{r}$ and $\tilde{p}$. Consider the sequence $\{p_r = \max(\frac{h\sigma_D^2}{2(\mu_D - r)^2}, \tilde{p})\}_{r \in (\tilde{r}, \mu_D)}$. By assumption and convexity of $\mathcal{C}(C_{r_p})$, for any $r \in (\tilde{r}, \mu_D)$:

$$h\mathbb{E}[J_\infty(C_{\tilde{r}})] + p_r(\mu_D - \tilde{r}) \leq h\mathbb{E}[J_\infty(C_r)] + p_r(\mu_D - r).$$

It follows that:

$$\frac{p_r(r-\tilde{r})}{h} \leq \mathbb{E}[J_\infty(C_r)] - \mathbb{E}[J_\infty(C_{\tilde{r}})] \leq \mathbb{E}[J_\infty(C_r)].$$

Therefore by definition of p_r

$$\frac{\sigma_D^2(r-\tilde{r})}{2(\mu_D - r)^2} \leq \mathbb{E}[J_\infty(C_r)]. \tag{4.27}$$

Now by part (b), for every $\epsilon > 0$ there exists $\tilde{r}_\epsilon \in [0, \mu_D)$ such that for all $r \geq \tilde{r}_\epsilon$

$$\frac{\mathbb{E}[J_\infty(C_r)]}{\sigma_D^2/(\mu_D - r)} < 1 + \epsilon.$$

Let $\max\{\tilde{r}, \tilde{r}_\epsilon\} < r < \mu_D$ for some $\epsilon > 0$. We divide both sides of Inequality (4.27) by $\sigma_D^2/(\mu_D - r)$ which implies that for all $\epsilon > 0$

$$\frac{r-\tilde{r}}{\mu_D - r} < 1 + \epsilon \quad \text{or} \quad \frac{r-\tilde{r}}{\mu_D - r} \leq 1 \tag{4.28}$$

for all $r \in (\max\{\tilde{r}, \tilde{r}_\epsilon\}, \mu_D)$. Inequality (4.28) cannot hold for all $r \in (\max\{\tilde{r}, \tilde{r}_\epsilon\}, \mu_D)$ unless $\tilde{r} = \mu_D$ which contradicts the assumption.

(e) Let $\hat{\mathcal{C}} : [0, \mu_D) \to \mathbb{R}_+$ be defined by $\hat{\mathcal{C}}(r) := h\sigma_D^2/\left(2(\mu_D - r)\right) + p(\mu_D - r)$. Let $\hat{r}$ denote the minimizer of $\hat{\mathcal{C}}$, i.e., $\hat{r}_p := \mu_D - \frac{\sqrt{h\sigma^2/2}}{\sqrt{p}}$. Now we have

$$\begin{aligned}\lim_{p\to\infty} \frac{\mathcal{C}(C_{\hat{r}_p})}{\hat{\mathcal{C}}(\hat{r}_p)} &= \lim_{p\to\infty} \frac{h\mathbb{E}[J_\infty(C_{\hat{r}_p})] + p(\mu_D - \hat{r}_p)}{h\sigma_D^2/\left(2(\mu_D - \hat{r}_p)\right) + p(\mu_D - \hat{r}_p)} \\ &= \lim_{p\to\infty} \frac{h\mathbb{E}[J_\infty(C_{\hat{r}_p})] + \sqrt{ph\sigma_D^2/2}}{\sqrt{2ph\sigma_D^2}} = 1, \end{aligned} \tag{4.29}$$

where the first two equalities use definitions and algebra and the final equality follows from part (b). Similarly, by part (d) we obtain

$$\lim_{p\to\infty} \frac{\mathcal{C}(C_{r_p})}{\hat{\mathcal{C}}(r_p)} = \lim_{p\to\infty} \frac{h\mathbb{E}[J_\infty(C_{r_p})] + p(\mu_D - r_p)}{h\sigma_D^2/\left(2(\mu_D - r_p)\right) + p(\mu_D - r_p)} = 1. \tag{4.30}$$

Optimality of r_p together with (4.30) implies $\lim_{p\to\infty} \frac{\mathcal{C}(C_{\hat{r}_p})}{\hat{\mathcal{C}}(r_p)} \geq 1$. Combining this with (4.29), we obtain

$$\lim_{p\to\infty} \frac{\mathcal{C}(C_{\hat{r}_p})}{\hat{\mathcal{C}}(r_p)} = \lim_{p\to\infty} \frac{\sqrt{2h\sigma_D^2}}{h\sigma_D^2/\left(2\sqrt{p}(\mu_D - r_p)\right) + \sqrt{p}(\mu_D - r_p)} \geq 1. \tag{4.31}$$

Observe that $h\sigma_D^2/\left(2\sqrt{p}(\mu_D - r_p)\right) + \sqrt{p}(\mu_D - r_p) \geq \sqrt{2h\sigma_D^2}$ and $\sqrt{2h\sigma_D^2} < \infty$. This result together with (4.31) yields the result: $\lim_{p\to\infty} \sqrt{p}(\mu_D - r_p) = \sqrt{h\sigma_D^2/2}$.

□

4.F Proof of Lemma 4.9

We drop P_ξ in $q_t(P_\xi)$, for simplicity of notation. Then we use iteration (4.1) L times to find

$$\begin{aligned} q_{t+1} =& \xi - \mathbb{E}[J_{t+1} \mid \mathbf{x}_{t-L+1}] \\ =& \xi - \mathbb{E}_{D_{t-L+1},\dots,D_t}[(((J_{t-L+1} + q_{t-L+1} - D_{t-L+1})^+ + \dots)^+ + q_t - D_t)^+] \\ \leq& \xi - \mathbb{E}_{D_{t-L+1},\dots,D_{t-1}}[J_t] - q_t + \mu_D = \mu_D + \mathbb{E}_{D_{t-L},\dots,D_{t-1}}[J_t] - \mathbb{E}_{D_{t-L+1},\dots,D_{t-1}}[J_t]. \end{aligned} \tag{4.32}$$

The first equality holds since for any $a \in \mathbb{R}$, $a^+ \geq a$. The final equality follows from $q_t = \xi - \mathbb{E}[J_t \mid \mathbf{x}_{t-L}]$. Next observe that $\mathbb{E}[\mathbb{E}_{D_{t-L},\dots,D_{t-1}}[J_t]] = \mathbb{E}[\mathbb{E}_{D_{t-L+1},\dots,D_{t-1}}[J_t]]$. This observation combined with (4.32), and Lemma 1 of van Jaarsveld and Arts (2024) imply the results. □

4.G Proof of Lemma 4.10

Define $\overline{\kappa} := \sup_{r\in[0,\mu_D)} \kappa$. Note that κ is strictly positive and finite for all $r \in [0, \mu_D)$. This fact, combined with Lemma 4.3 ensures that $0 < \overline{\kappa} < \infty$. Furthermore, $\sup_{r\in[0,\mu_D)} \mu_+ \leq 2\overline{\kappa}$, and $g_r(x) \leq \overline{\kappa}$ for all $x \geq 0$ by Lemma 4.5. For $\xi \geq 0$, let the random variable I_ξ be given by:

$$I_\xi := \frac{h\mu_+}{2(\mu_D - r)} \int_{J_t(P_\xi)}^{J_t(P_\xi)+q_t(P_\xi)-r} g_r(y)dy.$$

Then

$$\mathbb{E}[I_\xi] \leq \frac{h\mu_+}{2(\mu_D - r)} \mathbb{E}[q_t(P_\xi) - r] \sup_x g_r(x) \leq h\overline{\kappa}^2$$

The first inequality follows from applying the mean value theorem to I_ξ, and the second inequality from Lemma 4.9and from $\mu_+ \leq 2\overline{\kappa}$. □

4.H Proof of Lemma 4.11

We prove the result by induction. The case $t_2 = t_1$ holds by the definition of v_r (cf. Definition 4.2). Next assume the result holds for $t_2 \geq t_1$. Then, by Definition 4.2,

$$\begin{aligned} &\mathbb{E}[v_r(J_{t_2+1})|J_{t_1}] = \mathbb{E}[\mathbb{E}[c_{t_2+1}(C_r) + v_r(J_{t_2+2}) - \mathcal{C}(C_r)|J_{t_2+1}]|J_{t_1}] \\ &= \mathbb{E}[c_{t_2+1}(C_r) + v_r(J_{t_2+2})|J_{t_1}] - \mathcal{C}(C_r). \end{aligned}$$

Plugging this relation into the induction hypothesis we have,

$$\mathbb{E}[c_{[t_1,t_2]}(C_r)|J_{t_1}] = v_r(J_{t_1}) - \mathbb{E}[c_{t_2+1}(C_r) + v_r(J_{t_2+2})|J_{t_1}] + (t_2 + 2 - t_1)\mathcal{C}(C_r),$$

which gives the result by algebraic rearrangement. □

Chapter 5

Risk or Replace: Efficient Asymptotics for Data-Driven Maintenance

5.1 Introduction

Maintenance strategies are essential to sustain the reliable and efficient operation of critical systems that degrade over time. Failure of critical systems can cause safety hazards and substantial economic losses. The Guardian (Goodier and Campbell, 2025) reports that equipment malfunctions in the UK's National Health Service (NHS) have led to almost 100 deaths and harm to nearly 4,000 individuals, since 2022. Some of these incidents are linked to critical device failure in high-risk areas such as neonatal wards and emergency care. They also estimate that the costs of addressing NHS maintenance problems have nearly tripled over the past decade, rising from "£4.5 billion in 2012–2013 to £13.8 billion in 2023–2024". In the industrial sector, unplanned downtime also imposes enormous costs. Siemens (2024) find that the 500 globally largest companies lose 11% of their revenue on unplanned downtime, which amounts to $1.4 trillion per year, more than the 2024 GDP of the Netherlands (World Bank Group, 2025). The financial impact varies sharply across industries, from roughly $36,000 per hour in Fast-Moving Consumer Goods to $2.3 million per hour in the automotive sector. Maintenance decisions affect not only the reliability and economic performance of industrial systems, but also their environmental footprint. Components that can still function are often replaced unnecessarily. This replacement creates extra demand for raw materials, energy, and manufacturing, and leads to avoidable waste and emissions. The transportation of new equipment and technicians further contributes to the environmental impact. With lifetime extension and resource efficiency recognized among the core sustainability strategies by international and national organizations (see e.g., UN Environment Programme International Resource Panel, 2018; International Energy Agency, 2019; Thomas, 2023), maintenance

decisions informed by system condition monitoring offer a practical approach to reduce the frequency of failure while reducing the footprint of early replacements.

Condition-based maintenance (CBM) is a prominent maintenance strategy that enables organizations to track the condition of their systems and perform maintenance interventions precisely when necessary. The advancement of sensor technologies has resulted in an increasing adoption of CBM policies to maintain critical equipment. These technologies also provide more condition monitoring data, often in real-time, which enables and motivates the use of more complex yet better-performing data-driven maintenance decision models. Siemens (2024) estimate that full adoption of CBM by Fortune 500 companies could increase productivity by 5% and reduce maintenance costs by 40%, saving more than $600 billion per year.

The growing availability of condition monitoring data also transforms CBM strategies and allows organizations to make better use of the information. A traditional, yet simplistic, approach relies on the data to model the degradation process of an operating system and then makes CBM decisions under the assumption that all systems degrade identically (Arts et al., 2024; de Jonge and Scarf, 2020). This assumption has often been necessary because degradation data were traditionally available only during downtime. However, recent developments in automatic real-time monitoring (see e.g. GE HealthCare, 2010; Radiology Business, 2010) allow decision makers to account for differences in degradation behaviors across individual systems. With such data, decision makers can integrate learning directly into the CBM problem and make data-driven maintenance decisions that are optimally tailored to each system's condition and degradation trajectory. Although this approach yields better decisions, it also introduces substantial computational complexity, making the CBM problem tractable only under restrictive modeling assumptions.(Drent et al., 2023a; Chen et al., 2015; Elwany et al., 2011). Even in such simplified settings, implementation and computation of optimal policies remains a challenge for practitioners. We therefore develop efficient, data-driven CBM policies that are computationally tractable, achieve high performance in realistic operating environments, and offer provable performance guarantees.

Brief Model Explanation. In this chapter, we study a critical component subject to stochastic degradation over time that must be replaced at the appropriate time to minimize long-run maintenance and operating costs. We allow both continuous-time and discrete-time degradation models, while interventions (decisions and replacements) are restricted to evenly spaced time epochs, typically coinciding with planned site visits or shutdowns. We consider a general class of degradation models where the total degradation sustained between any two consecutive epochs are non-negative, independent, and identically distributed, and satisfy mild technical conditions. These classes contain many stochastic shock models with cumulative damage, such as compound Poisson processes, as well as stochastic wear processes, such as the gamma processes. All operating components follow the same type of stochastic degradation process, but the parameters of

these processes differ across components, modeling heterogeneity in the population. For instance, when degradation is modeled as a compound Poisson shock process, both the shock arrival rate and the parameters of the compounding (damage) distribution vary from one component to another. Furthermore, the degradation model parameters for an operating component are not directly observable and can only be inferred from real-time sensor measurements. For example, in turbomachinery, measurements of vibration, pressure, and temperature can be used to infer the underlying degradation process of a blade. Similarly, in medical imaging, changes in electrical resistance provide information to estimate the degradation process of X-ray tube filaments. We refer to the collection of such data as the degradation data, which contains all available relevant information about the degradation behavior of a component. This degradation data is critical to learn about the (degradation) parameters of any individual component. The variation in parameters from one component to the next is modeled by a prior distribution with hyperparameters. This assumption is common in practice as the prior distribution models variation in component characteristics that the manufacturer may provide. A component is considered to have failed once its degradation level exceeds a failure threshold. In our model, planned preventive interventions incur a fixed cost, whereas a component failure incurs both this cost and an additional fixed failure cost. The model also accounts for a cumulative operating cost that reflects the system's performance deterioration as the critical component degrades. Specifically, the operating cost of a new component is zero, and as the component's condition worsens, higher costs accumulate in each period. An example of operating cost is the revenue loss resulting from reduced production due to component aging. Another example is the increased energy expenditure resulting from the deterioration of system performance. Such costs vary with the system state and must be accounted for when determining optimal maintenance decisions. At the beginning of each period, the decision maker uses the available degradation data to decide whether to continue operating the component or replace it with a new one. Our objective is to determine maintenance policies that minimize the long-run average cost-rate, conditional on the degradation data observed by the decision maker.

A similar problem, under the discounted total cost criterion and for certain degradation models, has been formulated as a multi-dimensional POMDP and solved to optimality (see e.g. Drent et al., 2023a). However, POMDP optimization models become intractable for many degradation models, for large finite state spaces, and for all denumerable and continuous state spaces, due to the curse of dimensionality. In this chapter, we propose a tractable approach that can handle both discrete and continuous state spaces, accommodate a broad range of degradation models, and support high-dimensional data. This approach uses the available data to estimate the actions of a hypothetical "*Oracle*" who has full knowledge of the true parameters of the operating component's degradation process. We first characterize the Oracle's optimal maintenance policies, which can then be combined with parameter estimates to produce an estimate of the Oracle's optimal policy. Since computing the Oracle's optimal policies poses significant challenges, we em-

ploy a renewal-theoretic asymptotic approximation instead, which allows us to introduce efficient data-driven replacement policies. Then we introduce a scaling regime where both the failure threshold and the cost parameters increase proportionally, so that the Oracle's optimal cost-rate converges to a strictly positive and finite limit. In many real systems, components operate over long periods relative to the time between two consecutive CBM decision moments, and the cost of maintenance represents a notable portion of the system's running costs, conditions that motivate our proposed scaling regime. This regime also characterizes settings where POMDP optimization faces the greatest computational challenges, as the state space becomes large.

The contributions of this chapter are summarized as follows:

1. We propose simple and efficient data-driven CBM policies capable of handling a wide range of degradation processes and parameter learning mechanisms. From a computational perspective, our approach separates parameter learning from the determination of the replacement policy.

2. We show that the regret of these efficient data-driven replacement policies, that is the difference between the cost-rates of our data-driven policies and the Oracle's optimal cost-rate, converges to zero in our scaling regime, when parameters are estimated using a consistent estimator. To the best of our knowledge, this is the first strong convergence result for a data-driven CBM policy under general degradation and learning processes.

3. We evaluate our policy against the state-of-the-art POMDP-based Integrated Bayes policy (Drent et al., 2023a) using real degradation data and find that our policy consistently achieves superior cost-rate performance.

4. In an extensive simulation study, we test our data-driven policies in both discrete and continuous state spaces and observe that, in both cases, the regret of using our approach relative to the Oracle's optimal cost-rate is small, particularly when failure thresholds are high. We could not statistically distinguish the cost rates of our data-driven CBM policy from those of the Integrated Bayes policies when the latter are optimal.

5. Our policy is easily interpretable by practitioners as it relies on three intuitive steps: (i) The replacement threshold is expressed by specifying the right distance from the failure threshold, i.e., we specify a safety margin around the failure threshold such that replacement occurs when the degradation crosses this safety margin; (ii) we determine what the optimal safety margins are for an Oracle that can observe the degradation parameters of each component; (iii) the decision maker who cannot observe the degradation parameters uses a consistent estimator of the parameter based on the available information and decides to replace or not based on the safety margin that corresponds to the current estimate of the parameter.

The remainder of the chapter is organized as follows. Section 5.2 presents a brief review of the related literature. We introduce our model, outline our arguments, and state the main theoretical result in Section 5.3. Sections 5.4 and 5.5 develop our proposed adaptive policies and rigorously present the main results, which are proved through an asymptotic analysis in Section 5.6. In Section 5.8, we compare the performance of our data-driven policy with the Integrated Bayes policy using real degradation datasets, and Section 5.7 benchmarks our approach against the Oracle's optimal policy through extensive simulations with discrete and continuous state spaces. Finally, we provide our concluding remarks in Section 5.9.

We finish this section with a few technical definitions. A random variable X is said to be *lattice* if and only if there exists a $d > 0$ and such that $\sum_{n\in\mathbb{Z}} \mathbb{P}(X = nd) = 1$ and *non-lattice* otherwise. For a lattice random variable, the largest $d > 0$ that such that $\sum_{n\in\mathbb{Z}} \mathbb{P}(X = nd) = 1$ is called the period. A lattice random variable with period d and its distribution function are called $d-$*lattice*.

Let $\{X_n\}$ be a sequence of random vectors in $\mathbb{R}^d$ and let X be another random vector in $\mathbb{R}^d$. We say X_n converges to X in probability or weakly, written $X_n \xrightarrow{p} X$, if for every $\epsilon > 0$

$$\lim_{n\to\infty} \mathbb{P}\{\|X_n - X\| < \epsilon\} = 1,$$

where $\|\cdot\|$ is any norm on $\mathbb{R}^d$. We say X_n converges to X almost surely or strongly, written $X_n \xrightarrow{a.s.} X$, if

$$\mathbb{P}\left\{\lim_{n\to\infty} \|X_n - X\| = 0\right\} = 1.$$

5.2 Literature review

Maintenance optimization models for stochastically deteriorating systems have been extensively explored in the literature. Comprehensive reviews of the field can be found in de Jonge and Scarf (2020) and Arts et al. (2024). When the stochastic process for the deteriorating system is known and the state transitions are observable, a majority of works on discrete-time CBM use either Markov decision processes (MDP) (see e.g. Derman, 1963a,b; Kolesar, 1966; Ross, 1969; Andersen et al., 2022) or renewal theory (see e.g. Poppe et al., 2018; Zhang et al., 2020) to compute optimal maintenance decisions. These studies typically assume that all components share the same degradation process (i.e. a homogeneous population of components), this degradation process is known by decision makers, and they can perfectly observe the state of the system at any time. They show that, under these assumptions the optimal policy has a simple structure: a fixed threshold exists such that it is optimal to replace a component once its degradation level exceeds this threshold, and to continue operation otherwise. Even when such simple threshold policies are not optimal, some studies focus on finding the best threshold policy, as non-threshold policies are difficult to implement in practice (see e.g. Feldman, 1976).

However, the assumption of homogeneous component behavior seldom holds in practice, as components typically exhibit distinct degradation dynamics due to differences in their physical properties and operating conditions. Moreover, the underlying degradation process of a component is rarely known in advance and must be inferred from condition monitoring data. In addition, condition monitoring data can contain noise, and provide imperfect information about the component's state. Handling these uncertainties requires more complex decision making frameworks.

One major research stream to address these assumptions focuses on systems where low-cost, real-time condition monitoring data is available. This stream has gained attention over the past two decades thanks to technological advancements that make such data accessible. Studies in this area generally consider a specific stochastic process whose parameters, drawn from a prior distribution, are learned from observed condition monitoring data. Elwany et al. (2011) consider a Brownian motion wear model with a drift parameter which is estimated from the latest degradation observation and formulates a POMDP to compute the optimal condition-based maintenance decisions. In a similar vein, Chen et al. (2015) focus on the inverse Gaussian wear process, Zhang et al. (2016) on the gamma process, and van Oosterom et al. (2017) on a Markovian stochastic process characterized by a finite set of transition matrices. Recently, Drent et al. (2023a) study compound Poisson shock models, where both the shock arrival rate and the damage process parameters are unknown and drawn from some priors. They establish that if the damage process belongs to a one-parameter exponential family with conjugate priors, the optimal policy can be computed using a tractable POMDP framework. Drent et al. (2023a) demonstrate that the optimal replacement policy relies not only on the observed degradation level but also on other available information, such as the number of shocks occurring between decision epochs and the component ages.

Our work advances this research stream by broadening both its scope and methodology to address a wider and more realistic class of CBM problems. Specifically, our model accommodates a general non-decreasing independent and identically distributed degradation process with unknown parameters, which are estimated using a general consistent estimator. To the best of our knowledge, no prior work addresses such a broad range of degradation and learning processes. Although this problem can in principle be formulated as a POMDP, such an approach becomes computationally intractable for most realistic degradation and learning processes, particularly when component lifetimes are long. To overcome this limitation, we propose an alternative renewal-theoretic framework and establish its asymptotically optimality when both the average component lifetime and the associated maintenance costs are large. A further notable study in this stream, Kim (2016), formulates an optimal CBM policy designed to ensure robustness against posterior mis-specification. It is worth noting that another research stream addresses uncertainty in observations when real-time condition monitoring data are unavailable or costly to obtain (Girshick and Rubin, 1952; Ross, 1971; Maillart, 2006; Maillart and Zheltova, 2007; Kim and Makis, 2013; van Staden and Boute, 2021; Khaleghei and Kim, 2021; Zhang and

Zhang, 2023).

From an optimization perspective, this chapter adopts a renewal-theoretic framework, in which the replacement of a component constitutes a renewal event, the component's lifetime defines the renewal cycle length, and the associated maintenance cost represents the renewal cycle cost. The use of this paradigm dates back to the classical work of Barlow and Hunter (1960) on age based maintenance models. Since then, this approach has been used in some discrete-time CBM papers (Kim and Makis, 2013; Poppe et al., 2018; Zhang et al., 2020; van Staden and Boute, 2021). Nevertheless, this paradigm has been less common than MDP models in discrete-time CBM research, as MDP formulations facilitate both the derivation of structural results and the computation of optimal policies. It is noteworthy that a few other studies, such as Bérenguer et al. (2003), also use asymptotic approximations based on the renewal theory for maintenance problems.

5.3 Model

5.3.1 Degradation Process

We consider components that degrade with stationary non-negative independent increments. The components are indexed in the natural numbers $\mathbb{N} = \{1, 2, \ldots\}$. A decision maker can decide to replace a component at the beginning of evenly spaced time periods. Without loss of generality, we rescale time such that a period is one time unit. The periods are numbered on non-negative integers $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ and forward in time. We assume that replacement occurs instantaneously, and we define the beginning of a period $t \in \mathbb{N}_0$, when a decision is made, as epoch t. The age of an operating component is the number of epochs elapsed since its installation, with the installation epoch numbered as 0. The degradation increment between age $\tau - 1 \in \mathbb{N}_0$ and $\tau \in \mathbb{N}_0$ of component $i \in \mathbb{N}$ is denoted by $X_{i,\tau}$ so that the total degradation at age τ of component i is given by

$$S_{i,\tau} = \sum_{j=1}^{\tau} X_{i,j}.$$

The empty sum is zero such that $S_{i,0} = 0$, implying that a new component has no cumulative degradation; that is, the component is in an as-good-as-new state. The component i fails when the degradation level, $S_{i,\tau}$, exceeds the failure level $L < \infty$. For each item $i \in \mathbb{N}$, the sequence $\{X_{i,\tau}\}_{\tau \in \mathbb{N}}$ are non-negative independently and identically distributed random variables with a common probability distribution function $F_X(\cdot \mid \theta_i)$, where $\theta_i \in \Theta$ is the parameter of $F_X(\cdot \mid \theta_i)$, and $\Theta \subseteq \mathbb{R}^\rho, \rho \in \mathbb{N}$ is the parameter space. The parameters, θ_i, vary across components with a probability distribution function $F_\theta : \Theta \to [0, 1]$ which is well-defined. Furthermore, the degradation process of components are mutually independent, with each θ_i representing an independent realization from F_θ.

We impose the following conditions on the distribution functions F_θ and $F_X(\cdot \mid \theta)$, $\theta \in \Theta$, and the random variable $X_{i,1}$, $i \in \mathbb{N}$.

A.1. F_θ is continuous with a density f_θ.

A.2. Fix $x \geq 0$; $F_X(x \mid \theta)$ is a continuous function of θ.

A.3. Fix $\theta \in \Theta$; $F_X(\cdot \mid \theta)$ is continuous with a density function $f_X(\cdot \mid \theta) := dF_X(x \mid \theta)/dx$ or is $d-$lattice, for some $d > 0$, with a mass function $f_X(\cdot \mid \theta)$.

A.4. There exists a finite length $x_{\text{flat}} \geq 0$ such that for all $x_1, x_2 \in \mathbb{R}_+, \theta \in \Theta$, if $x_2 \geq x_1 + x_{\text{flat}}$ then $F_X(x_1 \mid \theta) < F_X(x_2 \mid \theta)$.

A.5. The increments have strictly positive mean $\mu(\theta) := \mathbb{E}[X_{i,1} \mid \theta_i = \theta] < \infty$ and finite variance $\sigma^2(\theta) := \mathbf{Var}[X_{i,1} \mid \theta_i = \theta] < \infty$.

A.6. $\mu(\theta_i)$ and $\sigma(\theta_i)$ are continuous functions of θ_i.

A.7. The following integrals are finite,

(a) $\int_\Theta \frac{1}{\mu(\mathbf{y})} dF_\theta(\mathbf{y})$,

(b) $\int_\Theta \frac{\sigma^2(\mathbf{y})}{\mu^2(\mathbf{y})} dF_\theta(\mathbf{y})$.

5.3.2 Observation Process

Let T_i denote the lifetime of the component $i \in \mathbb{N}$, which is the age of the component when it is replaced. T_i is a random variable that depends on the replacement decisions made for all previous components and on the replacement decision for component i itself. Let $N(t)$ denote the index of the component that operates at epoch t (before potential replacement):

$$N(t) := \min \left\{ i \in \mathbb{N}_0 : \sum_{j=1}^{i} T_j \geq t \right\}.$$

Moreover, let $A(t)$ denote the age of the component that operates at epoch t (before potential replacement):

$$A(t) := t - \sum_{i=1}^{N(t)-1} T_i.$$

Note that $N(t)$ and $A(t)$ are random variables. At each epoch, every decision maker has access to the complete past of components numbers, ages and degradation increments: $\left\{ N(j), A(j), S_{N(j),A(j)} \right\}_{j=0}^{t}$. Some decision makers may have access to even more information and we let Q_t denote additional information that becomes available at epoch t.

For example, Q_t may contain (i) the degradation level(s) measured in real-time between epochs $t-1$ and t (not only the accumulated level at epoch t) or it may contain (ii) (real-time) sensor readings that relate to the condition of the component or its degradation parameter $\theta_{N(t)}$. Thus, in general, the decision maker has access to the degradation data

$$\mathbf{I}_t := \left\{N(j), A(j), S_{N(j),A(j)}, Q_j\right\}_{j=0}^{t}$$

at epoch t, taking values in $\mathcal{I}_t$. Note that $\mathbf{I}_0$ contains all historical data before the start of the planning horizon. We let $\mathcal{F}$ denote the sequence of sigma-algebras generated by $\{\mathbf{I}_t\}_{t\in\mathbb{N}_0}$, i.e., $\mathcal{F} := \{\mathcal{F}_t\}_{t\in\mathbb{N}_0}$, where $\mathcal{F}_t = \sigma(\mathbf{I}_t)$. (Note that σ is used here to denote an induced sigma-field, not a standard deviation.) That is, $\mathcal{F}$ is the natural filtration of the underlying probability space.

Note that the true parameters of the degradation processes for the operating components $N(t)$ are not included in $\mathbf{I}_t$. However, we construct a hypothetical decision maker that does have access to the true degradation parameters $\theta_{N(t)}$. We call this decision maker the "Oracle" and note that at epoch t it has access to $\mathbf{O}_t := \mathbf{I}_t \cup \left\{\theta_{N(j)}\right\}_{j=0}^{t}$. As before, we let $\mathcal{O}$ denote the filtration of the Oracle.

5.3.3 Decision Problem

At each epoch $t \in \mathbb{N}_0$, the decision maker chooses either to continue operating the component or to replace it instantaneously, based on the available degradation data up to t, i.e., $\mathbf{I}_t$. Furthermore, the costs incurred between epochs $t-1$ and t are realized at epoch t. The chosen action determines the evolution of the system's state. We next introduce the cost structure of the model and then formulate the cost optimization problem faced by the decision maker.

Let $c_p \in \mathbb{R}_+$ denote the cost to replace a component preventively. Correctively replaced components incur a cost of $c_p + c_f$ such that $c_f \in \mathbb{R}_+$ is the additional cost of corrective replacement. Finally, the system incurs an operating cost between epochs $t-1$ and t whose expectation depends on the degradation level observed at epoch t and is given by $\ell : \mathbb{R}_+ \to \mathbb{R}_+$ and we assume that ℓ is non-decreasing, bounded, and continuous and satisfies the following conditions:

A.8. $\ell(0) = 0$ and $\lim_{x\to 0} \ell(x)/x = 0$,

A.9. ℓ attains its maximum at L; that is, $\ell(x) = \ell(L)$ for all $x \geq L$.

This last condition serves mainly to simplify notation, and the main results hold (with minor modifications) without it. The maintenance cost incurred by component i, denoted by C_i, is then given by,

$$C_i := c_p \mathbf{1}_{S_{i,T_i}\leq L} + (c_p + c_f)\mathbf{1}_{S_{i,T_i}>L} + \sum_{\tau=1}^{T_i} \ell(S_{i,\tau}), \tag{5.1}$$

where, **1** is the indicator function.

An admissible policy (or simply a policy) $\pi := \{\pi_t\}_{t\in\mathbb{N}_0}$ is a sequence of functions $\pi_t : \mathcal{I}_t \to \{\text{continue}, \text{replace}\}$ that takes in the data available at time t and maps it to a decision to continue to operate the current component, or replace it. Specifically, for the decision maker, π_t is $\mathcal{F}_t-$measurable, which is equivalent to saying that π is adapted to $\mathcal{F}$. For the Oracle, π is adapted to $\mathcal{O}$. Under a policy π, the age of component $i \in \mathbb{N}$ at its replacement epoch is denoted by $T_i(\pi)$, and the total maintenance cost by $C_i(\pi)$. The (infinite horizon) cost-rate of policy π under filtration $\mathcal{F}$ (or $\mathcal{O}$ for the Oracle), denoted $g(\pi)$, is defined as

$$g(\pi) := \limsup_{n\to\infty} \frac{\sum_{i=1}^n C_i(\pi)}{\sum_{i=1}^n T_i(\pi)}. \tag{5.2}$$

Let $\Pi_{\mathcal{F}}$ and $\Pi_{\mathcal{O}}$ denote the set of policies adapted to $\mathcal{F}$ and $\mathcal{O}$, respectively. The optimal cost-rate under $\mathcal{F}$ and $\mathcal{O}$, denoted $g^*(\mathcal{F})$ and $g^*(\mathcal{O})$, respectively, are given by

$$g^*(\mathcal{F}) := \inf_{\pi\in\Pi_{\mathcal{F}}} g(\pi), \qquad g^*(\mathcal{O}) := \inf_{\pi\in\Pi_{\mathcal{O}}} g(\pi). \tag{5.3}$$

When $g^*(\mathcal{F})$ ($g^*(\mathcal{O})$) can be attained by a policy, such a policy is said to be optimal and is denoted $\pi^*(\mathcal{F})$ ($\pi^*(\mathcal{O})$). The objective of the decision maker is to identify an optimal policy, $\pi^*(\mathcal{F})$, if one exists.

Examples: Compound Poisson Degradation.

Poisson cumulative damage models have been extensively employed in both the academic literature and practical applications to characterize the progressive deterioration of components or systems under stochastic shocks. In these models, the operating component experiences shocks that occur according to a Poisson process with rate ν. Each shock induces an i.i.d. random damage, and the damages accumulate additively over time. Thus, $X_{i,\tau}$, $i \in \mathbb{N}, \tau \in \mathbb{N}_0$, follows a compound Poisson distribution. Here, we omit the index i, for notational simplicity.

The compound Poisson process is defined as follows. Let $M_{(\tau_1,\tau_2]}$, with $\tau_1 < \tau_2$, denote the number of shocks occurring during the time interval between ages $\tau_1, \tau_2 \in \mathbb{N}_0$, denoted by $(\tau_1, \tau_2]$, and let $Z_{m,\tau}$ denote the damage from the m-th shock in the interval $(\tau - 1, \tau]$. The sequence $\{Z_{m,\tau}\}_{m\in\mathbb{N},\tau\in\mathbb{N}_0}$ consists of non-negative i.i.d. random variables with a common distribution function $F_Z(\cdot \mid \theta_Z)$ with a mass or density function $f_Z(\cdot \mid \theta_Z)$, where θ_Z is the parameter of F_Z. The parameter space, Θ, is defined as the set of all possible values of (ν, θ_Z) that may be realized across the operating components, i.e., $\Theta = \mathbb{R}_{++} \times \Theta_Z$ where $\Theta_Z \subseteq \mathbb{R}^\rho, \rho \in \mathbb{N}$, denote the spaces from which ν and θ_Z are drawn according to the distributions F_ν and F_{θ_Z}, respectively. Notice that F_ν and F_{θ_Z} are continuous with densities f_ν and f_{θ_Z}, respectively. In this case $f_\theta(x, \mathbf{y}) = f_\nu(x) f_{\theta_Z}(\mathbf{y})$, where $x \in \mathbb{R}_+ +$ and $\mathbf{y} \in \Theta_Z$. By the memory-less property we have $M_{[\tau_1,\tau_2)} = \sum_{\tau=\tau_1+1}^{\tau_2} M_\tau$, where $M_\tau := M_{[\tau-1,\tau)}$. Consequently, the cumulative damage between $\tau - 1$ and τ, i.e., X_τ,

is given by $X_\tau = \sum_{m=1}^{M_\tau} Z_{m,\tau}$. For any $\tau \in \mathbb{N}_0$, the probability mass function of M_τ, is $f_M(x \mid \nu) := \mathbb{P}\left\{M_{[\tau,\tau+1)} = x \mid \nu\right\} = \nu^x e^{-\nu}/x!$, where, $x \in \mathbb{N}_0$. Furthermore, by the law of total expectation and law of total variance we have

$$\mathbb{E}\left[X_\tau \mid \nu\right] = \nu \mathbb{E}\left[Z_{1,1}\right], \qquad \mathbf{Var}\left[X_\tau \mid \nu\right] = \nu \mathbf{Var}\left[Z_{1,1}\right] + \nu\left(\mathbb{E}\left[Z_{1,1}\right]\right)^2. \tag{5.4}$$

To facilitate the analysis of our examples we consider two distributions for $Z_{m,\tau}$: (lattice) geometric distribution with parameter $\theta_Z = p \in (0,1)$, and (continuous) exponential distributions with parameter $\theta_Z = \omega \in \mathbb{R}_{++}$.

Example 1 (Geometric Compounding). Let $Z_{m,\tau}$ be supported on $\mathbb{N}_0$ and have a geometric distribution with the success probability of $p \in (0,1)$, that is, $f_Z(x \mid p) := \mathbb{P}\left\{Z_{m,\tau} = x \mid p\right\} = (1-p)^x p$, where $x \in \mathbb{N}_0$. Next, by the application of the law of total probability the distribution mass function of X_τ is expressed by

$$f_X\left(x \mid \nu, p\right) = \begin{cases} f_M\left(0 \mid \nu\right) + \sum_{j=1}^{\infty} f_M\left(j \mid \nu\right) f_{NB}(0 \mid j, p), & x = 0 \\ \sum_{j=1}^{\infty} f_M\left(j \mid \nu\right) f_{NB}(x \mid j, p), & x \in \mathbb{N}, \end{cases} \tag{5.5}$$

where $f_{NB}(\cdot \mid j, p)$ is the negative binomial probability mass function with the j number of successes and success probability of p, i.e., $f_{NB}(x \mid j, p) := \binom{x+j-1}{x}(1-p)^x p^j$, with $x \in \mathbb{N}_0$. Additionally by Equations (5.4) we have

$$\mu(\nu, p) = \frac{\nu(1-p)}{p}, \qquad \sigma^2(\nu, p) = \frac{\nu(1-p)(2-p)}{p^2}. \tag{5.6}$$

Example 2 (Exponential Compounding). Let $Z_{m,\tau}$ be supported on $\mathbb{R}_+$ and have an exponential distribution with parameter $\omega > 0$, that is, $f_Z(x \mid \omega) := \mathbb{P}\left\{Z_{m,\tau} = x \mid \omega\right\} = \omega e^{-\omega x}$, where $x \geq 0$. Accordingly, the probability mass function of X_τ can be expressed, using the law of total probability, as

$$f_X\left(x \mid \nu, \omega\right) = \begin{cases} f_M\left(0 \mid \nu\right) + \sum_{j=1}^{\infty} f_M\left(j \mid \nu\right) f_{Er}(0 \mid j, \omega), & x = 0 \\ \sum_{j=1}^{\infty} f_M\left(j \mid \nu\right) f_{Er}(x \mid j, \omega), & x > 0, \end{cases} \tag{5.7}$$

where $f_{Er}(\cdot \mid j, \omega)$ is the Erlang-j probability density function with the rate ω, i.e., $f_{Er}(x \mid j, \omega) := \omega^j x^{j-1} e^{-\omega x}/(j-1)!$. By Equations (5.4) the moments of X_τ are given by

$$\mu(\nu, \omega) = \frac{\nu}{\omega}, \qquad \sigma^2(\nu, \omega) = \frac{2\nu}{\omega^2}. \tag{5.8}$$

The degradation data observed by the decision maker at epoch $t \in \mathbb{N}_0$ includes the history of replacement times, the number of shocks, and the size of the damages induced by each shock up to t. That is,

$$\mathbf{I}_t = \left\{N(j), A(j), M_{N(j),A(j)}, Z_{N(j),1,A(j)}, \ldots, Z_{N(j),M_{N(j),A(j)},A(j)}\right\}_{j=0}^{t}.$$

This degradation data also enables the computation of $X_{N(t),A(t)}$ and $S_{N(t),A(t)}$. The functions F_ν and F_{θ_Z}, as well as F_X for every realization of (ν, θ_Z) are also known. The decision maker may sequentially estimate (ν, θ_Z) in a Bayesian manner, based on the degradation data $\mathbf{I}_t$ using the estimators $\left(\hat{\nu}\left(\mathbf{I}_t\right), \hat{\theta}_Z\left(\mathbf{I}_t\right)\right)$, with density functions $f_\nu(\cdot \mid \mathbf{I}_t)$ and $f_{\theta_Z}(\cdot \mid \mathbf{I}_t)$, respectively.

Returning to Examples 1 and 2, we notice that $f_M(\cdot \mid \nu)$, $f_Z(\cdot \mid p)$, and $f_Z(x \mid \omega)$ belong to the one-parameter exponential family. This property allows us to select conjugate priors for each of these (likelihood) functions, so that the posterior distributions remain within the same family, and the corresponding parameters can be updated efficiently, conditional on the degradation data $\mathbf{I}_t$ (cf. Section 5.1.5. Ghosh et al., 2007). The selection of the conjugate priors is as follows.

We first select that $f_\nu(\cdot) = f_\nu(\cdot \mid \mathbf{I}_0)$ corresponds to a gamma distribution with shape α_0 and scale β_0. The choice of the density function for θ_Z, i.e. $f_{\theta_Z}(\cdot) = f_{\theta_Z}(\cdot \mid \mathbf{I}_0)$, depends on $Z_{m,\tau}$: it follows a beta distribution when $Z_{m,\tau}$ is geometric, and a gamma distribution when $Z_{m,\tau}$ is exponential. The selected f_ν and f_{θ_Z} are conjugate priors corresponding to likelihood functions f_M and f_Z. In both cases, f_{θ_Z} is parametrized by a_0 and b_0. One can simply verify that the chosen f_ν and f_{θ_Z} satisfy all conditions specified in Section 5.3.1.

Then, $\hat{\nu}\left(\mathbf{I}_t\right)$ follows a gamma distribution with parameters α_t and β_t. Furthermore, $\hat{\theta}_Z\left(\mathbf{I}_t\right)$ follows a beta distribution with parameters a_t and b_t if $Z_{1,1}$ is geometrically distributed, and a gamma distribution with parameters a_t and b_t if $Z_{1,1}$ is exponential. We refer to $(\alpha_t, \beta_t, a_t, b_t)$ as hyperparameters. Let $M_{A(t)}$ and $X_{A(t)}$ be the number of shocks and the degradation increment in the interval $(t-1, t]$. Then, the hyperparameters can be updated as

$$\begin{aligned}(\alpha_t, \beta_t, a_t, b_t) &= \left(\alpha_{t-1} + M_{A(t)}, \beta_{t-1} + 1, a_{t-1} + M_{A(t)}, b_{t-1} + X_{A(t)}\right) \\ &= \left(\alpha_0 + \sum_{j=1}^{A(t)} M_j, \beta_0 + A(t), a_0 + \sum_{j=1}^{A(t)} M_j, b_0 + S_t\right) \end{aligned} \tag{5.9}$$

The entire analysis in our examples can be replicated for many choices for f_Z from the one-parameter exponential distributions if their conjugate priors satisfy the conditions outlined in Section 5.3.1.

5.3.4 Outline of Arguments and Main Results

In general, proving the existence of $\pi^*(\mathcal{F})$ and, if it exists, computing it is not trivial. Moreover, its implementation in practice may be highly challenging, as the policy may have a complex structure. In certain cases, where the state space is discrete and conjugate families are used to estimate θ, this problem under the discounted total cost criterion can be solved, tractably, using a multi-dimensional POMDP (cf. Elwany et al., 2011; Drent et al., 2023a). However, even in these specific situations, the POMDP becomes intractable

quickly as the failure level increases. To overcome these challenges, our goal is to construct policies that approximate the Oracle decisions. Notice that by construction, the Oracle has access to more information than a real decision maker, and therefore its achievable cost-rate is less than or equal to that of any other decision maker, since $\Pi_{\mathcal{F}} \subseteq \Pi_{\mathcal{O}}$, i.e., the Oracle faces an information relaxation of the problem faced by a real decision maker and so

$$g^*(\mathcal{O}) \leq g^*(\mathcal{F}), \tag{5.10}$$

(cf. Blackwell's informativeness in Blackwell, 1951, 1953). We define the regret associated with any policy $\pi \in \Pi_{\mathcal{F}}$, relative to the Oracle's optimal policy $\pi^*(\mathcal{O})$, as $\mathrm{REG}(\pi) := g(\pi) - g^*(\mathcal{O})$. Under this definition, inequality (5.10) is equivalent to stating that the regret of every policy $\pi \in \Pi_{\mathcal{F}}$ is nonnegative, i.e., $\mathrm{REG}(\pi) \geq 0$.

In the subsequent sections, we develop an efficient policy for non-Oracle decision makers. This policy is constructed by first estimating the unknown parameters of the operating component, $\theta_{N(t)}$, from degradation data, $\mathbf{I}_t$, using a consistent estimator, and then estimating the Oracle's optimal actions based on these parameter estimates. We impose no assumptions on the parameter estimator other than consistency; in particular, standard Bayesian estimators, such as the posterior mean, are consistent in our setting We term this policy the Estimated Oracle's Optimal Policy (EOP).

We further introduce a scaling regime in which the failure thresholds, L, grow large together with the cost parameters c_p and c_f, as well as the total operating cost $\sum_{\tau=1}^{T} \ell(S_\tau)$ to ensure that the cost rate of any policy remains finite and positive. Specifically, we consider a base instance with failure threshold $\tilde{L}$, cost of preventive replacement $\tilde{c}_p$, additional corrective replacement cost $\tilde{c}_f$, and operating cost function $\tilde{\ell}$. Then we create a continuum of instances for each $k > 0$ such that it has failure threshold $L(k) = k\tilde{L}$, cost of preventive replacement $c_p(k) = k\tilde{c}_p$, additional cost of corrective replacement $c_f(k) = k\tilde{c}_f$, and operating cost function $\ell(x, k) = \tilde{\ell}(x/k)$. Thus, k is a scaling parameter that increases both the maintenance costs and the failure threshold. This regime is motivated by practical settings where component lifetimes are typically much longer than the time between decision epochs, and maintenance costs constitute a significant share of the average system's operating expenses per unit time.

Under our scaling regime, the cost-rate and regret functions are parameterized by the scaling parameter k, that is $\mathrm{REG}_k(\pi) := g_k(\pi) - g_k^*$ (cf. Equation (5.2)). We now state our main theoretical result:

Main Result. *The regret associated with using the EOP instead of the Oracle's optimal policy vanishes in the limit as k approaches infinity, that is*

$$\lim_{k\to\infty} \mathrm{REG}_k(\mathrm{EOP}) = 0.$$

This result is formally stated in Theorem 5.6. Our main result, together with Equation (5.10), establishes the asymptotic optimality of the EOP within the proposed scaling regime.

We organize our theoretical exposition as follows. Section 5.4 characterizes the Oracle's optimal policy $\pi^*(\mathcal{O})$, which serves as the basis for the construction of the EOP in Section 5.5. Our main result is then rigorously presented at the end of Section 5.5. Subsequently, in Section 5.6 we explain how the results in Section 5.5 can be obtained, which also provides insights that improve the interpretability of our adaptive policy for practitioners.

5.4 Oracle's Optimal Policies

In this section, we address the optimal maintenance policy for the Oracle, i.e., $\pi^*(\mathcal{O})$, whose degradation data at each epoch $t \in \mathbb{N}_0$ contains the parameters of the degradation process for all operating components up to t, namely $\theta_{N(1)}, \ldots, \theta_{N(t)}$. Throughout this section, the discussion refers to filtration $\mathcal{O}$, and therefore we omit it from the notation. For example, we write $g^* := g^*(\mathcal{O})$ and $\pi^* := \pi^*(\mathcal{O})$. The reader may, however, interpret $\mathcal{O}$ implicitly wherever the context requires.

A replacement policy π is called stationary if and only if $\pi_t = \pi_{t+1}$ for all t. The critical observation to find the optimal policy of the Oracle is that the future evolution of degradation at epoch t depends on $\mathbf{O}_t$ only through $\theta_{N(t)}$, i.e., $\{S_{N(t),A(t)}, \theta_{N(t)}\}_{t\in\mathbb{N}}$ is a Markov process for any stationary policy π. The process $\{S_{N(t),A(t)}, \theta_{N(t)}\}_{t\in\mathbb{N}}$ is positive recurrent because any policy eventually replaces a component (in the worst case after failure) and expected time between failures are finite (see (5.40) in Appendix 5.C).

A stationary replacement policy π is a parameter-specific threshold policy if there exists a threshold function $\xi : \Theta \to [0, L]$ such that at each epoch $t \in \mathbb{N}_0$ the operating component is replaced if and only if the degradation level is strictly greater than $\xi(\theta_{N(t)})$, i.e.,

$$\pi_t(S_{N(t),A(t)}) = \begin{cases} \text{continue} & \text{if } S_{N(t),A(t)} \le \xi(\theta_{N(t)}), \\ \text{replace} & \text{if } S_{N(t),A(t)} > \xi(\theta_{N(t)}). \end{cases}$$

If an Oracle's optimal policy π^* exists, then there exists a bounded function $v : R_+ \times \Theta \to \mathbb{R}$ and the optimal cost-rate $g^* \in \mathbb{R}$ such that

$$g^* + v(s,\theta) = \begin{cases} \min\Big\{\ell(s) + \mathbb{E}[v(s + X_{1,1}, \theta) \mid \theta_1 = \theta],\ \ell(s) + c_p + \bar{v}(\theta)\Big\}, & s \le L,\ \theta \in \Theta, \\ \ell(s) + c_p + c_f + \bar{v}(\theta), & s > L,\ \theta \in \Theta, \end{cases} \tag{5.11}$$

where we use the shorthand for the expected continuation value after replacement,

$$\bar{v}(\theta) = \int_\Theta \mathbb{E}[v(X_{1,1}, \theta_1) \mid \theta_1 = \mathbf{y}]\, dF_\theta(\mathbf{y}),$$

and π^* denotes the policy that satisfies the optimality equations (5.11). It can be readily shown that a parameter-specific threshold policy is optimal under the discounted total cost criterion. The following theorem states that even under the more challenging

cost-rate criterion, there exists an optimal policy which is likewise a parameter-specific threshold policy.

Theorem 5.1. *(a) There exists a (cost-rate) optimal policy π^*, where $g(\pi^*) = g^*$,*

(b) For each $\theta \in \Theta$, there is an optimal parameter-specific threshold policy, i.e.,

$$\pi_t^*(S_{N(t),A(t)}) = \begin{cases} \textit{continue} & \textit{if } S_{N(t),A(t)} \leq \xi^O(\theta_{N(t)}), \\ \textit{replace} & \textit{if } S_{N(t),A(t)} > \xi^O(\theta_{N(t)}), \end{cases}$$

where $\xi^O : \Theta \to [0, L]$.

We henceforth restrict attention to the class of parameter-specific threshold policies, denoted by Π, as Theorem 5.1 establishes that this set contains an optimal policy.

Now, let Ξ denote the set of all threshold functions ξ. Observe that the replacement of an operating component under the parameter-specific threshold policy $\pi \in \Pi$ is determined solely via the threshold function $\xi \in \Xi$; thus, it is straightforward to see that Ξ and Π are equivalent. With slight abuse of notation, we write the cost-rate of a policy $\pi \in \Pi$ with threshold function $\xi \in \Xi$ as $g(\xi)$. Accordingly, for the optimal threshold function ξ^{O}, we have $g(\xi^{\mathrm{O}}) = g^*$.

We will now study how the cost rate of any threshold policy ξ can be evaluated. To this end let $T(x, \theta)$ be the lifetime of a component with degradation parameter θ if it is replaced when the degradation exceeds $x \in \mathbb{R}_+$:

$$T(x, \theta) := \inf\{\tau \in \mathbb{N} : S_{1,\tau} > x, \theta_1 = \theta\}.$$

Similarly, let $C(x, \theta)$ denote the costs incurred by a components with degradation parameter θ that is replaced when it degradation exceeds $x \in \mathbb{R}_+$:

$$C(x, \theta) := c_p \mathbf{1}_{S_{1,T(x,\theta)} \leq L} + (c_p + c_f)\mathbf{1}_{S_{1,T(x,\theta)} > L} + \sum_{\tau=1}^{T(x,\theta)} \ell(S_{1,\tau}), \qquad \theta_1 = \theta.$$

Then the following holds.

Lemma 5.2. *Under the policy $\pi \in \Pi$ characterized by the threshold function $\xi \in \Xi$, with probability one we have*

$$\begin{aligned} g(\xi) =& \frac{\int_\Theta \mathbb{E}[C(\xi(\mathbf{y}), \mathbf{y})] dF_\theta(\mathbf{y})}{\int_\Theta \mathbb{E}[T(\xi(\mathbf{y}), \mathbf{y})] dF_\theta(\mathbf{y})} \\ =& \frac{c_p + \int_\Theta \left(c_f \mathbb{P}\{S_{1,T(\xi(\theta_1),\theta_1)} > L \mid \theta_1 = \mathbf{y}\} + \mathbb{E}\left[\sum_{\tau=1}^{T(\xi(\theta_1),\theta_1)} \ell(S_{1,\tau}) \Big| \theta_1 = \mathbf{y} \right] \right) dF_\theta(\mathbf{y})}{\int_\Theta \mathbb{E}[T(\xi(\mathbf{y}), \mathbf{y})] dF_\theta(\mathbf{y})}. \end{aligned} \tag{5.12}$$

The intuition behind Lemma 5.2 is that any two components are probabilistically equivalent, and the corresponding stopping decisions follow the same probabilistic structure. Equation (5.12) can be equivalently written in terms of the amount by which the degradation of a component i exceeds $\xi(\theta_i)$ upon replacement. We denote this amount by the random variable $\tilde{Y}_i(\xi(\theta_i))$ where $\tilde{Y}_i(x)$ is expressed by $\tilde{Y}_i(x) := S_{i,T_i} - x$. Figure 5.1 illustrates a sample degradation path for the first component, failure level L, replacement threshold $\xi(\theta_1)$, excess degradation $\tilde{Y}_1(\xi(\theta_1))$, and $\phi\left(\xi(\theta_1)\right)$, where $\phi(x) := L - x$ represents the distance of the point $x \in [0, L]$ and L.

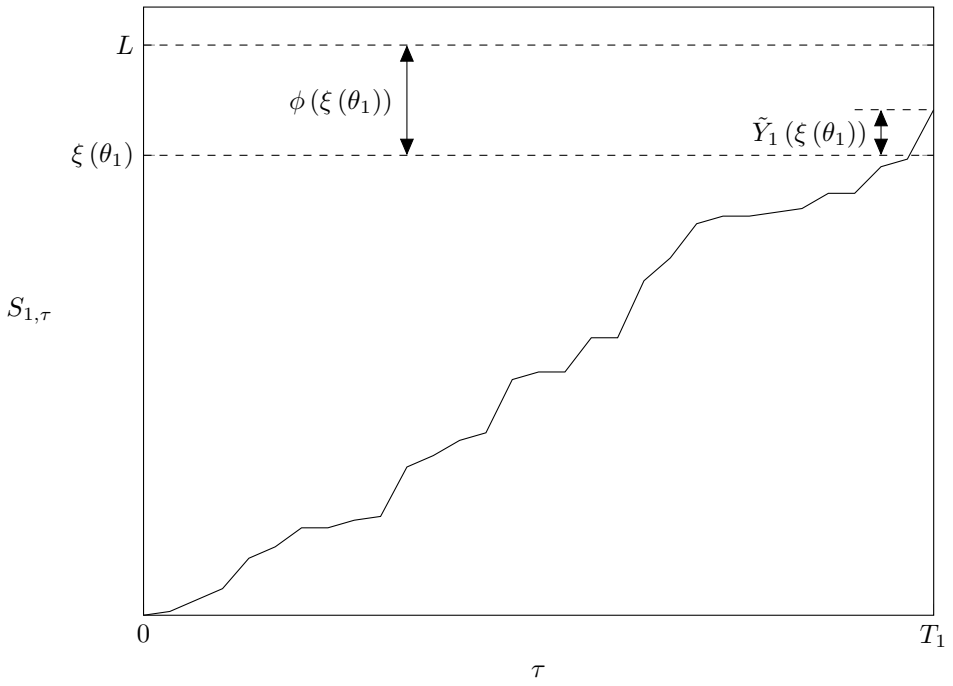

Figure 5.1: A sample degradation path.

Then we have

$$g(\xi) = \frac{c_p + \int_\Theta \left(c_f \mathbb{P}\left\{ \tilde{Y}_1(\xi(\theta_1)) > \phi\left(\xi(\theta_1)\right) \mid \theta_1 = \mathbf{y} \right\} + \mathbb{E}\left[\sum_{\tau=1}^{T(\xi(\theta_1),\theta_1)} \ell(S_{1,\tau}) \mid \theta_1 = \mathbf{y} \right] \right) dF_\theta(\mathbf{y})}{\int_\Theta \frac{1}{\mu(\mathbf{y})} \left(\xi(\mathbf{y}) + \mathbb{E}\left[\tilde{Y}_1(\xi(\theta_1)) \mid \theta_1 = \mathbf{y} \right] \right) dF_\theta(\mathbf{y})}, \tag{5.13}$$

Note that $S_{i,T_i} = \sum_{\tau=1}^{T_i} X_{i,\tau} = \xi(\theta_i) + \tilde{Y}_i(\xi(\theta_i))$ so that by Wald's identity $\mathbb{E}[T(\xi, \theta_i)] = (\xi(\theta_i) + \mathbb{E}[\tilde{Y}_i(\xi(\theta_i))])/\mu(\theta_i)$.

Equation (5.12) (and equivalently (5.13)) implies that any function $\xi \in \Xi$ satisfies

$$\int_\Theta \mathbb{E}[C(\xi(\mathbf{y}), \mathbf{y})] dF_\theta(\mathbf{y}) - g^* \int_\Theta \mathbb{E}[T(\xi(\mathbf{y}), \mathbf{y})] dF_\theta(\mathbf{y}) \geq 0,$$

where equality holds if and only if ξ corresponds to the optimal threshold function, i.e.,

$\xi = \xi^{\mathrm{O}}$. This property can be exploited to develop an iterative procedure to compute the optimal cost-rate and corresponding policies.

Define $\tilde{h} : \mathbb{R} \times [0, L] \times \Theta \to \mathbb{R}$ as

$$\tilde{h}(\lambda, x, \theta) :=$$
$$c_p + c_f \mathbb{P}\left\{\tilde{Y}_1(x) > \phi(x) \middle| \theta_1 = \theta\right\} + \mathbb{E}\left[\sum_{\tau=1}^{T(x,\theta_1)} \ell(S_{1,\tau}) \middle| \theta_1 = \theta\right] - \frac{\lambda}{\mu(\mathbf{y})}\left(\xi(\mathbf{y}) + \mathbb{E}\left[\tilde{Y}_1(x) \middle| \theta_1 = \theta\right]\right), \tag{5.14}$$

and let the threshold $\tilde{\xi}_\lambda(\theta)$ minimize $\tilde{h}(\lambda, x, \theta)$ for a given $\lambda \in \mathbb{R}$ and $\theta \in \Theta$. It follows that, for $\lambda \in \mathbb{R}$, the threshold function $\tilde{\xi} \in \Xi$ satisfies

$$\tilde{\xi}_\lambda \in \arg\min_{\xi \in \Xi} h(\lambda, \xi, \theta), \qquad \forall \theta \in \Theta.$$

Notice that $\tilde{\xi}_{g^*}$ determines the optimal threshold function ξ^{O}, and by optimality

$$\int_\Theta \tilde{h}(g^*, \tilde{\xi}_{g^*}, \mathbf{y}) dF_\theta(\mathbf{y}) = 0.$$

We build on the seminal results of Aven and Bergman (1986) to present the following lemma, which ensures sequential convergence to g^*.

Lemma 5.3. *(a) g^* is the unique solution to the equation $\lambda = g(\tilde{\xi}_\lambda)$.*

(b) If $\lambda_1 \in \mathbb{R}$, and $\lambda_{j+1} = g(\tilde{\xi}_{\lambda_j})$, for $j \in \{2, 3, \dots\}$, then

$$\lim_{j \to \infty} \lambda_j = g^*.$$

Lemma 5.3 provides a fixed-point iteration to compute g^* and, consequently, the optimal replacement thresholds. Notice that the functions in (5.13) and (5.14) can, in general, be numerically evaluated through simulation.

5.5 Asymptotically Optimal Adaptive Replacement Policies

5.5.1 Estimated Oracle's Optimal Policy (EOP)

In this section, we study a non-Oracle decision maker who seeks to determine the optimal policy under the filtration $\mathcal{F}$, i.e., $\pi^*(\mathcal{F})$. At each epoch $t \in \mathbb{N}_0$, the (true) parameter of the operating component $\theta_{N(t)}$ is unknown to the decision maker and can only be estimated conditional on the degradation data $\mathbf{I}_t$. Since only the information related to the operating component is relevant for estimation, we focus on an arbitrary component

without loss of generality. Let $\theta_{N(t_0+\tau)} \in \Theta$, denote the (true) parameter of the component installed at t_0 and aged τ, with $t_0, \tau \in \mathbb{N}_0$. The decision maker estimates $\theta_{N(t_0+\tau)}$, based on the degradation data $\mathbf{I}_{t_0+\tau}$, using a point estimator $\hat{\theta}(\mathbf{I}_{t_0+\tau})$. For instance, the decision maker first uses the prior distribution F_θ, which represents the baseline belief $F_\theta(\cdot \mid \mathbf{I}_0)$, to update her knowledge of $\theta_{N(t_0+\tau)}$ by obtaining a posterior distribution $F_\theta(\cdot \mid \mathbf{I}_{t_0+\tau})$. She then estimates $\theta_{N(t_0+\tau)}$ using the posterior mean, which is the Bayes estimator under the squared error loss function. Note that the information before t_0 is irrelevant for posterior updating and estimation. Therefore, we suppress the index of the operating component and t_0, for notational convenience, whenever the context permits. We also assume that $\hat{\theta}(\mathbf{I}_\tau) \in \Theta$ with probability 1, without loss of generality. The point estimator $\hat{\theta}(\mathbf{I}_\tau)$ is said to be (weakly) consistent if $\hat{\theta}(\mathbf{I}_\tau) \xrightarrow{p} \theta$. In the following, we restrict attention to consistent point estimators and refer to them simply as point estimators.

At each epoch, the point estimate of the parameter can be used to compute a replacement threshold based on the full-information (Oracle) optimal policy. If the estimate were equal to the true parameter, the resulting threshold would coincide with the optimal full-information threshold. In particular, the optimal threshold $\xi^{\mathrm{O}}(\theta)$ can be estimated by $\xi^{\mathrm{O}}\left(\hat{\theta}(\mathbf{I}_\tau)\right)$. This requires us first to calculate g^* using Lemma 5.3, then to obtain ξ^{O} as the minimizer of $\tilde{h}(g^*, x, \theta)$ according to Equation (5.14), and finally to compute $\xi^{\mathrm{O}}\left(\hat{\theta}(\mathbf{I}_\tau)\right)$. Unfortunately, applying Equations (5.13) and (5.14) is generally difficult, as analytical expressions for $\mathbb{P}\left\{\tilde{Y}_1(\xi(\theta_1)) \leq x \mid \theta_1 = \theta\right\}$, $\mathbb{E}\left[\sum_{\tau=1}^{T(\xi(\theta),\theta)} \ell(S_{1,\tau}) \middle| \theta_1 = \theta\right]$, and $\mathbb{E}\left[\tilde{Y}_1(\xi(\theta_1)) \middle| \theta_1 = \theta\right]$ are not readily available. However, we construct computationally efficient approximations of these functions and prove their asymptotic optimality as the failure threshold grows large.

Let $Y(\theta)$, $\theta \in \Theta$ be a random variable such that

$$\mathbb{P}\{Y(\theta) \leq x\} = \begin{cases} \dfrac{1}{\mu(\theta)} \displaystyle\int_0^x (1 - F_X(y, \theta))\, dy & \text{if } F_X \text{ is continuous,} \\ \dfrac{d}{\mu(\theta)} \displaystyle\sum_{j=0}^{\lfloor x/d \rfloor - 1} (1 - F_X(jd, \theta)) & \text{if } F_X \text{ is } d\text{−lattice,} \end{cases} \tag{5.15}$$

where empty sums are 0 by definition and

$$\mathbb{E}[Y(\theta)] = \begin{cases} \dfrac{\sigma^2(\theta) + \mu^2(\theta)}{2\mu(\theta)} & \text{if } F_X \text{ is continuous,} \\ \dfrac{\sigma^2(\theta) + \mu^2(\theta)}{2\mu(\theta)} + \dfrac{d}{2} & \text{if } F_X \text{ is } d\text{−lattice.} \end{cases} \tag{5.16}$$

Define $r : [0, L] \times \Theta \to \mathbb{R}_+$ as

$$r(x,\theta) = \begin{cases} \mathbb{E}\left[\ell(X_{1,1}) \mid \theta_1 = \theta\right] + \dfrac{1}{\mu(\theta)} \displaystyle\int_0^x \mathbb{E}\left[\ell(X_{1,1}+y) \mid \theta_1 = \theta\right] dy, & \text{if } F_x \text{ is continuous,} \\ \mathbb{E}\left[\ell(X_{1,1}) \mid \theta_1 = \theta\right] + \dfrac{d}{\mu(\theta)} \displaystyle\sum_{j=0}^{\lfloor x/d \rfloor - 1} \mathbb{E}\left[\ell(X_{1,1}+jd) \middle| \theta_1 = \theta\right], & \text{if } F_x \text{ is } d\text{-lattice.} \end{cases} \tag{5.17}$$

We will later show that Y is the asymptotic expansion of $\tilde{Y}_i$, for all $i \in \mathbb{N}$ as the replacement thresholds tend to infinity. Similarly, $r(\xi(\theta), \theta)$ provides the asymptotic expansion of $\mathbb{E}\left[\sum_{\tau=1}^{T(\xi(\theta),\theta)} \ell(S_{1,\tau}) \middle| \theta_1 = \theta\right]$ in the same limiting regime. Based on these asymptotic expansions, we approximate the cost-rate function $g(\xi)$ by $\gamma(\xi)$ given by

$$\gamma(\xi) := \frac{c_p + \displaystyle\int_\Theta \left(c_f \mathbb{P}\left\{Y(\mathbf{y}) > \phi\left(\xi(\mathbf{y})\right)\right\} + r(\xi(\mathbf{y}), \mathbf{y})\right) dF_\theta(\mathbf{y})}{\displaystyle\int_\Theta \frac{1}{\mu(\mathbf{y})} \left(\xi(\mathbf{y}) + \mathbb{E}[Y(\mathbf{y})]\right) dF_\theta(\mathbf{y})}, \tag{5.18}$$

(cf. Equation (5.13)). Morover, we approximate the function $\tilde{h}$ by h expressed as

$$h(\lambda, x, \theta) := c_p + c_f \mathbb{P}\left\{Y(\theta) > \phi\left(x\right)\right\} + r(x,\theta) - \frac{\lambda}{\mu(\theta)} \left(x + \mathbb{E}[Y(\theta)]\right), \tag{5.19}$$

(cf. Equation (5.14)).

Let ξ_λ minimize h for fixed $\lambda \in \mathbb{R}$, i.e., $\xi_\lambda(\theta) \in \arg\min_{x \in [0,L]} h(\lambda, x, \theta)$, $\theta \in \Theta$. We propose to use the heuristic threshold functions $\xi_\lambda(\theta)$ instead of $\tilde{\xi}_\lambda$. While the computation of $\tilde{h}(\lambda, x, \theta)$ as per Equation (5.14) could pose challenges, $h(\lambda, x, \theta)$ can be calculated by Equation (5.19) using the distribution function and the first two moments of $X_{1,1}$, having parameter θ. The next theorem explains that the heuristic threshold functions $\xi_\lambda(\theta)$ can be simply calculated using conventional search methods. This allows us to approximate the optimal cost-rate g^* with the approximated optimal cost-rate γ^* through an iterative approach similar to Lemma 5.3.

Define $Dh : \mathbb{R} \times [0, L] \times \Theta \to \mathbb{R}$ as

$$Dh(\lambda, x, \theta) = \begin{cases} c_f\left(1 - F_X(\phi(x) \mid \theta)\right) + \mathbb{E}\left[\ell\left(x + X_{1,1}\right) \mid \theta_1 = \theta\right] - \lambda, & \text{if } F_x \text{ is continuous,} \\ c_f d\left(1 - F_X\left(d\left\lfloor \dfrac{\phi(x)}{d} - 1 \right\rfloor \middle| \theta\right)\right) + \mathbb{E}\left[\ell\left(d\left\lfloor \dfrac{x + X_{11}}{d} - 1 \right\rfloor\right) \middle| \theta_1 = \theta\right] - \lambda, & \text{if } F_x \text{ is } d\text{-lattice.} \end{cases} \tag{5.20}$$

Dh represents the rate of change of h when x is increased by an infinitesimal amount if F_X is continuous, or by one unit if F_X is $d-$lattice. For any fixed $\lambda \in \mathbb{R}$ and $\theta \in \Theta$, let $\mathcal{D}_{\lambda,\theta}$ denote the set of the roots of $Dh(\lambda, x, \theta)$, i.e.,

$$\mathcal{D}_{\lambda,\theta} := \begin{cases} \left\{x \in [0, L] : Dh(\lambda, x, \theta) = 0\right\}, & \text{if } F_x \text{ is continuous,} \\ \left\{jd : j \in \mathbb{N}, j \le \dfrac{L}{d}, Dh(\lambda, x, \theta) = 0\right\}, & \text{if } F_x \text{ is } d\text{-lattice.} \end{cases}$$

If $Dh(\lambda, x, \theta)$ in strictly negative for all values of $x \in [0, L]$, then $\mathcal{D}_{\lambda,\theta} := \{L\}$. Conversely, if $Dh(\lambda, x, \theta)$ is strictly positive for all $\mathcal{D}_{\lambda,\theta} := \{0\}$. Clearly, $\mathcal{D}_{\lambda,\theta}$ is a closed and bounded set, and therefore it admits a well-defined minimum element.

Theorem 5.4. *(a) γ^* is the unique solution to the equation $\lambda = \gamma(\xi_\lambda)$.*

(b) Choose $\lambda_1 \in \mathbb{R}$, and $\lambda_{j+1} = \gamma(\xi_{\lambda_j}), j = 1, 2, \ldots$ Then

$$\lim_{j\to\infty} \lambda_j = \gamma^*.$$

(c) $Dh(\lambda, x, \theta)$ is non-decreasing in x; and $\xi_\lambda(\theta)$ is given by

$$\xi_\lambda(\theta) = \begin{cases} \min\{x \in [0, L] : x \in \mathcal{D}_{\lambda,\theta}\}, & \text{if } F_x \text{ is continuous,} \\ \min\left\{jd : j \in \mathbb{N}, j \le \dfrac{L}{d}, jd \in \mathcal{D}_{\lambda,\theta}\right\}, & \text{if } F_x \text{ is d-lattice.} \end{cases}$$

By Theorem 5.4, the heuristic threshold $\xi_\lambda(\theta)$ can be obtained through the application of any root finding technique, such as bisection, on Dh. Consequently, we can efficiently compute the approximated optimal cost-rate γ^* via the iterative procedure stated in Part (b) of Theorem 5.4. The resulting approximated optimal threshold function, $\xi^{\mathrm{A}}(\theta) := \xi_{\gamma^*}(\theta)$, $\theta \in \Theta$, minimizes $\gamma(\xi)$. Finally, the integrals over the parameter space in the numerator and denominator of Equation (5.18) can be efficiently evaluated using numerical techniques, such as Monte Carlo or Gaussian quadrature methods, once ξ_λ is specified.

Remark 5.5. When $\ell(x)$ is constant for all $x \in \mathbb{R}_+$, the computation of $\phi(\xi^{\mathrm{A}}_\lambda(\theta)) = L - \xi_\lambda(\theta)$ reduces to identifying a specific quantile of $F_X(\cdot \mid \theta)$, similar to the classical newsvendor problem. In particular, for continuous degradation processes we have that

$$\phi(\xi^{\mathrm{A}}_\lambda(\theta)) = F^{-1}_X(1 - \lambda/c_f \mid \theta).$$

Now, we propose our Estimated Oracle's Optimal Policy (EOP) with its adaptive thresholds expressed as

$$\xi^{\mathrm{P}}(\mathbf{I}_\tau) := \xi^{\mathrm{A}}\left(\hat{\theta}(\mathbf{I}_\tau)\right) \tag{5.21}$$

In summary, EOP thresholds can be computed as follows:

1. Compute γ^* using $F_\theta(\cdot \mid \mathbf{I}_0)$, as specified in Theorem 5.4,
2. For a component aged τ compute $\xi^{\mathrm{P}}(\mathbf{I}_\tau)$, as given in Equation (5.21),
3. Continue operating the component while $S_\tau \le \xi^{\mathrm{P}}(\mathbf{I}_\tau)$; replace otherwise.

EOP is broadly applicable across various degradation and estimation processes, as its parameter estimation and decision making steps are decoupled and operate independently.

For practical implementation, the policy can be explained to practitioners in terms of an estimator of the safety margin around the failure level, defined by $\phi\left(\xi^{\mathrm{P}}(\mathbf{I}_\tau)\right)$. This estimator provides an indicator of how close the system is to its failure level, so the decision maker can take informed replacement actions while accounting for uncertainty in both the degradation and parameter estimates.

Examples (Continued): Compound Poisson Degradation.

We continue Examples 1 and 2 from Section 5.3.3 to illustrate the practical implementation of our adaptive threshold policies. We demonstrate how to obtain closed-form expressions for the distribution function and expectation of Y based on Equations (5.15) and (5.16), as well as for ξ_λ and γ^* according to Theorem 5.4. Finally, we show how to determine $\xi^{\mathrm{P}}(\mathbf{I}_\tau)$ using the posterior mean following Equation (5.21) for both examples.

Example 1 (Continued): Geometric Compounding. We first describe the computation of $\xi_\lambda(\nu, p)$ based on Theorem 5.4(c). Let $F_{NB}(\cdot \mid l, p)$ denote the distribution function of a negative binomial random variable with parameters l and p. Then

$$\begin{aligned} F_X(x \mid \nu, p) =& \sum_{j=0}^{x} f_X(j \mid \nu, p) = f_M(0 \mid \nu) + \sum_{j=0}^{x} \sum_{l=1}^{\infty} f_M(l \mid \nu) f_{NB}(j \mid l, p) \\ =& f_M(0 \mid \nu) + \sum_{l=1}^{\infty} f_M(l \mid \nu) F_{NB}(x \mid l, p). \end{aligned} \tag{5.22}$$

We use Theorem 5.4(c) together with Equation (5.22) to derive a relation for ξ_λ.

Next, we describe the calculation of γ^* using Theorem 5.4(b). The key steps are to evaluate $\mathbb{P}\{Y(\nu, p) \leq x\}$ and $\mathbb{E}[Y(\nu, p)]$ that allow the calculation of $\gamma(\xi)$ according to Equation (5.18). By Equation (5.15), when F_X is 1−lattice and $x \in \mathbb{N}$, we have

$$\begin{aligned} \mathbb{P}\{Y(\nu, p) \leq x\} =& \frac{1}{\mu(\theta)} \sum_{j=0}^{x-1} (1 - F_X(j \mid \theta)) = \frac{1}{\mu(\theta)} \sum_{j=0}^{x-1} \left(1 - \sum_{l=0}^{j} f_X(l \mid \theta)\right) \\ =& \frac{1}{\mu(\theta)} \left(x - \sum_{j=0}^{x-1} (x - j) f_X(j \mid \theta)\right) \\ =& \frac{1}{\mu(\theta)} \left(x - x F_X(x - 1 \mid \theta) + \sum_{j=0}^{x-1} j f_X(j \mid \theta)\right). \end{aligned} \tag{5.23}$$

We use Equation (5.7) to derive a relation for the last term in (5.23). Let $\mu_{NB}(l, p) := \dfrac{l(1-p)}{p}$ denote the mean of a negative binomial random variable with parameters l and

p, then

$$\begin{aligned}\sum_{j=0}^{x-1} j f_X(j \mid \nu, p) &= \sum_{j=1}^{x-1}\sum_{l=1}^{\infty} f_M(l \mid \nu)\, j f_{NB}(j \mid l, p) = \sum_{j=1}^{x-1}\sum_{l=1}^{\infty} f_M(l \mid \nu)\, j \binom{j+l-1}{j}(1-p)^j p^l \\ &= \sum_{j=1}^{x-1}\sum_{l=1}^{\infty} f_M(l \mid \nu)\, \frac{l(1-p)}{p}\binom{j+l-1}{j-1}(1-p)^{j-1}p^{l+1} \\ &= \sum_{j=1}^{x-1}\sum_{l=1}^{\infty} f_M(l \mid \nu)\, \mu_{NB}(l,p) f_{NB}(j-1 \mid l+1, p) \\ &= \sum_{l=1}^{\infty} f_M(l \mid \nu)\, \mu_{NB}(l,p) F_{NB}(x-2 \mid l+1, p) \end{aligned} \tag{5.24}$$

Combining (5.8), (5.23), (5.22), and (5.24) gives us

$$\mathbb{P}\{Y(\nu,p) \le x\} = \frac{p}{\nu(1-p)}\left(x\left(1 - f_M(0 \mid \nu)\right) + \sum_{l=1}^{\infty} f_M(l \mid \nu)\Big(\mu_{NB}(l,p) F_{NB}(x-2 \mid l+1, p) - x F_{NB}(x-1 \mid l, p)\Big)\right) \tag{5.25}$$

By Equations (5.8) and (5.16)

$$E[Y(\nu,p)] = \frac{(\nu(1-p)/p)^2 + \nu(1-p)(2-p)/p^2}{2\nu(1-p)/p} + \frac{1}{2} = \frac{\nu(1-p)+2}{2p}. \tag{5.26}$$

Using (5.18) together with (5.25) and (5.26), one can compute $\gamma(\xi)$ efficiently. Subsequently, γ^* is determined using the iterative algorithm described in Theorem 5.4(b). Once γ^* is computed, $\xi^{\mathrm{A}}(\nu,p) = \xi_{\gamma^*}(\nu,p)$ can be calculated efficiently.

Next, we describe how to compute the EOP threshold using the available degradation data $\mathbf{I}_\tau = (S_\tau, M_\tau, \tau)$ for the first component, where the decision epoch coincides with the component's age.

We use Equation (5.9) to update the posterior and use the posterior mean as the point estimator of the parameter, that is

$$\begin{aligned}\hat{\theta}(\mathbf{I}_\tau) &= \left(\hat{\nu}(\mathbf{I}_\tau), \hat{\theta}_Z(\mathbf{I}_\tau)\right) = \left(\frac{\alpha_\tau}{\beta_\tau}, \frac{a_\tau}{a_\tau + b_\tau}\right) \\ &= \left(\frac{\alpha_{\tau-1} + M_\tau}{\beta_{\tau-1}+1}, \frac{a_{\tau-1}+M_\tau}{a_{\tau-1}+b_{\tau-1}+M_\tau+X_\tau}\right) = \left(\frac{\alpha_0 + \sum_{j=1}^{\tau} M_j}{\beta_0+\tau}, \frac{a_0 + \sum_{j=1}^{\tau} M_j}{a_0+b_0+\sum_{j=1}^{\tau} M_j + S_\tau}\right), \end{aligned} \tag{5.27}$$

and

$$\xi^{\mathrm{P}}(\mathbf{I}_\tau) = \xi^{\mathrm{A}}\left(\hat{\nu}(\mathbf{I}_\tau), \hat{\theta}_Z(\mathbf{I}_\tau)\right).$$

Example 2 (Continued): Exponential Compounding. In the following, we adapt the previous analysis by replacing the geometric compounding distribution with an

exponential distribution with a similar analytical framework. We omit some intermediate derivation steps, as the reader can follow from the geometric compounding example. Let $F_{Er}(\cdot \mid l, \omega)$ be the Erlang$-l$ distribution function with rate ω. Then $F_X(\cdot \mid \nu, \omega)$ is given by

$$F_X(x \mid \nu, \omega) = f_M(0 \mid \nu) + \sum_{l=1}^{\infty} f_M(l \mid \nu) F_{Er}(x \mid l, \omega). \tag{5.28}$$

Let the mean of the Erlang$-l$ distribution be $\mu_{Er}(l, \omega) := l/\omega$. Then

$$\mathbb{P}\{Y(\nu, \omega) \leq x\} = \\ \frac{\omega}{\nu}\left(x\left(1 - f_M(0 \mid \nu)\right) + \sum_{l=1}^{\infty} f_M(l \mid \nu)\Big(\mu_{Er}(l, \omega) F_{Er}(x \mid l+1, \omega) - x F_{NB}(x \mid l, \omega)\Big)\right). \tag{5.29}$$

and

$$E[Y(\nu, \omega)] = \frac{\nu + 2}{\omega}. \tag{5.30}$$

Having these results, one can compute ξ_{γ^*} using Theorem 5.4. Next, the hyperparameters can be updated as per (5.9), that is

$$\begin{aligned} \hat{\theta}(\mathbf{I}_\tau) &= \left(\hat{\nu}(\mathbf{I}_\tau), \hat{\theta}_Z(\mathbf{I}_\tau)\right) = \left(\frac{\alpha_\tau}{\beta_\tau}, \frac{a_\tau}{b_\tau}\right) \\ &= \left(\frac{\alpha_{\tau-1} + M_\tau}{\beta_{\tau-1} + 1}, \frac{a_{\tau-1} + M_\tau}{b_{\tau-1} + X_\tau}\right) = \left(\frac{\alpha_0 + \sum_{j=1}^{\tau} M_j}{\beta_0 + \tau}, \frac{a_0 + \sum_{j=1}^{\tau} M_j}{b_0 + S_\tau}\right), \end{aligned} \tag{5.31}$$

The difference between (5.31) and (5.27) lies in the distribution of θ_Z: when the compounding distribution is exponential, f_{θ_Z} is the density of a Gamma distribution, whereas under geometric compounding, θ_Z follows a Beta distribution. Finally, our EOP thresholds are expressed as follows

$$\xi^{\mathrm{P}}(\mathbf{I}_\tau) = \xi^{\mathrm{A}}\left(\hat{\nu}(\mathbf{I}_\tau), \hat{\theta}_Z(\mathbf{I}_\tau)\right).$$

5.5.2 Scaling Regime and Main Asymptotic Optimality Results

Our analysis so far provides a tractable policy that fully exploits the available data for decision making in an otherwise intractable maintenance problem. Beyond tractability, a key performance measure of any policy is its optimality gap, which we examine in the arguments that follow. We restrict our study to large failure thresholds for several reasons. First, as the average time to failure increases, exact methods quickly become intractable, and heuristics such as the one we propose become particularly attractive due to their tractability advantages. Second, from a practical perspective, new equipment often operates for a relatively long time before failing. Finally, in terms of learning effectiveness, short failure times leave little opportunity for accurate parameter estimation, which lead to high estimation errors. In such cases, the problem can be reduced to a homogeneous

population, as posterior distributions just before replacements are nearly identical across components.

Next, we study the EOP thresholds $\xi^{\mathrm{P}}(\mathbf{I}_\tau)$ under the scaling regime described in Section 5.3.4, and formally state our main result that $\mathrm{REG}_k(\xi^{\mathrm{P}}(\mathbf{I}_\tau))$ vanishes as the scaling parameter increases. Let, for $k > 0$, Ξ_k be the set of all threshold functions $\xi : \Theta \to [0, L(k)]$. Define $g_k : \Xi_k \to \mathbb{R}_+$ and $\gamma_k : \Xi_k \to \mathbb{R}_+$ by

$$g_k(\xi) := \frac{c_p(k) + \int_\Theta \left(c_f(k)\mathbb{P}\left\{ \tilde{Y}_1(\xi(\theta_1)) > \phi_k(\xi(\theta_1)) \middle| \theta_1 = \mathbf{y} \right\} + \mathbb{E}\left[\sum_{\tau=1}^{T(\xi(\theta_1),\theta_1)} \ell(S_{1,\tau}, k) \middle| \theta_1 = \mathbf{y} \right] \right) dF_\theta(\mathbf{y})}{\int_\Theta \frac{1}{\mu(\mathbf{y})} \left(\xi(\mathbf{y}) + \mathbb{E}\left[\tilde{Y}_1(\xi(\theta_1)) \middle| \theta_1 = \mathbf{y} \right] \right) dF_\theta(\mathbf{y})} \tag{5.32}$$

and

$$\gamma_k(\xi) := \frac{c_p(k) + \int_\Theta \left(c_f(k)\mathbb{P}\left\{ Y(\mathbf{y}) > \phi_k(\xi(\mathbf{y})) \right\} + r_k(\xi, \mathbf{y}) \right) dF_\theta(\mathbf{y})}{\int_\Theta \frac{1}{\mu(\mathbf{y})} \left(\xi(\mathbf{y}) + \mathbb{E}[Y(\mathbf{y})] \right) dF_\theta(\mathbf{y})}, \tag{5.33}$$

where $\phi_k(x) := L(k) - x$, $x \in [0, L(k)]$, and $r_k : [0, L(k)] \times \Theta \to \mathbb{R}_+$ as

$$r_k(x,\theta) = \begin{cases} \mathbb{E}[\ell(X_{1,1}, k) \mid \theta_1 = \theta] + \dfrac{1}{\mu(\theta)} \displaystyle\int_0^x \mathbb{E}[\ell(X_{1,1} + y, k) \mid \theta_1 = \theta]\, dy, & \text{if } F_x \text{ is continuous,} \\ \mathbb{E}[\ell(X_{1,1}, k) \mid \theta_1 = \theta] + \dfrac{d}{\mu(\theta)} \displaystyle\sum_{j=0}^{\lfloor x/d \rfloor} \mathbb{E}[\ell(X_{1,1} + jd, k) \mid \theta_1 = \theta], & \text{if } F_x \text{ is } d\text{-lattice.} \end{cases} \tag{5.34}$$

Define the minimizers of g_k and γ_k by

$$\xi_k^{\mathrm{O}} \in \arg\min_{\xi \in \Xi_k} g_k(\xi), \qquad \text{and,} \qquad \xi_k^{\mathrm{A}} \in \arg\min_{\xi \in \Xi_k} \gamma_k(\xi),$$

Then our EOP thresholds in the scaling regime are defined as

$$\xi_k^{\mathrm{P}}(\mathbf{I}_\tau) := \xi_k^{\mathrm{A}}\left(\hat{\theta}(\mathbf{I}_\tau)\right). \tag{5.35}$$

Next, we present the main theoretical result of this chapter.

Theorem 5.6. *As $k \to \infty$, the cost-rate of the EOP thresholds converge to the cost-rate of the Oracle's optimal policy, that is*

$$\lim_{k\to\infty} g_k(\xi_k^P) - g_k(\xi_k^O) = \lim_{k\to\infty} REG(\xi_k^P) = 0.$$

Theorem 5.6 establishes that the cost-rates of our EOP thresholds are asymptotically optimal as $k \to \infty$.

5.6 Asymptotic Analysis

We next outline the main steps of the proof of Theorem 5.6.

1. We first show that, under the scaling regime, both $\xi_k^{\mathrm{O}}(\theta)$ and $\xi_k^{\mathrm{A}}(\theta)$ diverge to infinity for all $\theta \in \Theta$ as $k \to \infty$, while the cost-rate function $g_k(\xi)$ and its approximation $\gamma_k(\xi)$ remain bounded for all threshold functions ξ. We further show that, as k increases, the gap between $g_k(\xi_k^{\mathrm{O}})$ and $g_k(\xi_k^{\mathrm{A}})$ vanishes, i.e., $g_k(\xi_k^{\mathrm{A}}) - g_k(\xi_k^{\mathrm{O}}) \to 0$. The proof proceeds by establishing the convergence of each function appearing in the expression of g_k in (5.32) to its corresponding counterpart in γ_k as defined in (5.33).

2. We next show that, under the EOP threshold function ξ_k^{P}, the lifetime of an operating component also diverges, as $k \to \infty$. Specifically, let $\mathcal{T}_k(\theta_1)$ denote the lifetime of the first component under the sequence of EOP thresholds $\{\xi_k^{\mathrm{P}}(\mathbf{I}_\tau)\}_{\tau=1}^{\mathcal{T}_k(\theta_1)}$, where k is the scaling parameter. Let $\xi_{k,\mathrm{last}}^{\mathrm{P}}(\theta_1)$ denote the random variable corresponding to the last EOP threshold at which the component is replaced by component 2, i.e., $\xi_{k,\mathrm{last}}^{\mathrm{P}}(\theta_1) := \xi_k^{\mathrm{P}}\left(\mathbf{I}_{\mathcal{T}_k(\theta_1)}\right)$. We show that $\mathcal{T}_k(\theta_1) \to \infty$ as $k \to \infty$, which implies that for large k the decision maker's error in estimating the true parameter θ_1 from the degradation data becomes negligible. The same argument holds for all components $i \in \mathbb{N}$. We use this result together with the (weak) consistency of the estimator to prove that the cost-rate of replacing at $\xi_{k,\mathrm{last}}^{\mathrm{P}}(\theta_i)$, i.e., $g(\xi_{k,\mathrm{last}}^{\mathrm{P}})$, converges to $g_k(\xi_k^{\mathrm{O}})$ as k becomes large.

While this section presents the main asymptotic results, the full mathematical proofs are detailed in the Appendix.

5.6.1 Asymptotics of the Oracle's Optimal Thresholds

Our construction of the Oracle's approximate optimal threshold function ξ^{A} relies on the asymptotic expansion of the overshoot random variable from classical renewal theory. In our setting, this result takes the following form:

Lemma 5.7. *(Theorems 2.6.1 and 2.6.2 Gut, 2009a) For all $i \in \mathbb{N}$ and $\theta_i \in \Theta$, as $\xi(\theta_i) \to \infty$,*

(a) $\mathbb{P}\{\tilde{Y}_i(\xi(\theta_i)) \le x\} \to \mathbb{P}\{Y(\theta_i) \le x\}$ *for all* $x \ge 0$,

(b) $\mathbb{E}[\tilde{Y}_i(\xi(\theta_i))] \to \mathbb{E}[Y(\theta_i)]$.

We now present the convergence results for different functions used to construct g_k in (5.32).

Lemma 5.8. *Let for the sequence of threshold functions* $\left\{\xi^{(k)} \in \Xi_k\right\}_{k>0}$, *we have* $\lim_{k\to\infty} \xi^{(k)}(\theta)/k > 0$ *for all* $\theta \in \Theta$. *Then, as* $k \to \infty$

(a) For all $\theta \in \Theta$, $\mathbb{E}\left[\sum_{\tau=1}^{T\left(\xi^{(k)}(\theta_1),\theta_1\right)} \ell(S_{1,\tau}, k)\middle|\theta_1 = \theta\right] - r_k\left(\xi^{(k)}(\theta), \theta\right) \to 0,$

(b) $\int_\Theta \mathbb{E}\left[\sum_{\tau=1}^{T\left(\xi^{(k)}(\theta_1),\theta_1\right)} \ell(S_{1,\tau}, k)\middle|\theta_1 = \mathbf{y}\right] dF_\theta(\mathbf{y}) - \int_\Theta r_k\left(\xi^{(k)}(\mathbf{y}), \mathbf{y}\right) dF_\theta(\mathbf{y}) \to 0,$

(c) $\int_\Theta \mathbb{P}\left\{\tilde{Y}_1(\xi^{(k)}(\theta_1)) > \phi_k\left(\xi^{(k)}(\theta_1)\right)\middle|\theta_1 = \mathbf{y}\right\} dF_\theta(\mathbf{y}) - \int_\Theta \mathbb{P}\left\{Y(\mathbf{y}) > \phi_k\left(\xi^{(k)}(\mathbf{y})\right)\right\} dF_\theta(\mathbf{y}) \to 0$

(d) $\frac{1}{k}\int_\Theta \frac{\xi^{(k)}(\mathbf{y}) + \mathbb{E}\left[\tilde{Y}_1(\xi^{(k)}(\theta_1))\middle|\theta_1 = \mathbf{y}\right]}{\mu(\mathbf{y})} dF_\theta(\mathbf{y}) - \frac{1}{k}\int_\Theta \frac{\xi^{(k)}(\mathbf{y}) + \mathbb{E}\left[Y(\mathbf{y})\right]}{\mu(\mathbf{y})} dF_\theta(\mathbf{y}) \to 0$

Lemma 5.8 indicates that as the scaling parameter k increases, each function in (5.32) converges to its counterpart in (5.33), when the threshold functions scale at least linearly with respect to k. Intuitively, this leads to the convergence of the gap between g_k and γ_k, which is formalized in the next lemma.

Lemma 5.9. *For the sequence of* $\left\{\xi^{(k)} \in \Xi_k\right\}_{k>0}$ *that, for all* $\theta \in \Theta$, *satisfies* $\lim_{k\to\infty} \xi^{(k)}(\theta)/k > 0$

$$\lim_{k\to\infty}\left(g_k\left(\xi^{(k)}\right) - \gamma_k\left(\xi^{(k)}\right)\right) = 0.$$

Lemmas 5.8 and 5.9 apply only to threshold functions whose growth is asymptotically linear in k. A natural question is whether the minimizers of g_k and γ_k, i.e.,

$$\xi_k^{\mathrm{O}} \in \arg\min_{\xi\in\Xi_k} g_k(\xi), \qquad \text{and,} \qquad \xi_k^{\mathrm{A}} \in \arg\min_{\xi\in\Xi_k} \gamma_k(\xi),$$

fall into this class. The following lemma establishes that both threshold functions indeed satisfy this property.

Lemma 5.10. *For* $\theta \in \Theta$

(a) $\lim_{k\to\infty} \xi_k^O(\theta)/k \in (0, \tilde{L}]$,

(b) $\lim_{k\to\infty} \xi_k^A(\theta)/k \in (0, \tilde{L}]$,

We can now apply Lemma 5.9 and 5.10 to establish the following convergence result for the cost-rates of g_k and γ_k.

Theorem 5.11. *As* $k \to \infty$,

(a) $g_k(\xi_k^A) - g_k(\xi_k^O) \to 0$.

(b) $g_k(\xi_k^A) - \gamma_k(\xi_k^A) \to 0$.

Theorem 5.11(a) shows that the optimality gap resulting from the use of the threshold function ξ_k^{A} in place of the optimal threshold ξ_k^{O} vanishes as k increases. Part (b) shows that if $g_k\left(\xi_k^{\mathrm{A}}\right)$ is predicted by $\gamma_k\left(\xi_k^{\mathrm{A}}\right)$, the prediction gap tends to zero for sufficiently large k.

5.6.2 Asymptotic Analysis of the EOP Thresholds

We begin our analysis with the first component; the same reasoning applies to any operating component. By Lemma 5.10 together with Equations (5.35) one can verify that $\xi_k^{\mathrm{P}}(\mathbf{I}_\tau)$ diverge as k increases, i.e., for all $\tau \in \mathbb{N}_0$,

$$\lim_{k\to\infty} \xi_k^{\mathrm{P}}(\mathbf{I}_\tau) = \lim_{k\to\infty} \xi_k^{\mathrm{A}}\left(\hat{\theta}(\mathbf{I}_\tau)\right) = \infty, \qquad \text{almost surely.} \tag{5.36}$$

By combining the divergence result in (5.36) with the strong law for counting processes (Theorem 2.5.1 in Gut, 2009a), we obtain that, for the true parameter $\theta_1 \in \Theta$,

$$\lim_{k\to\infty} \mathcal{T}_k(\theta_1) = \infty \qquad \textit{almost surely.} \tag{5.37}$$

Equation (5.37) implies that the number of observations also diverges as k approaches infinity. Consequently, the errors in estimating the true parameter, θ_1, by the point estimators $\hat{\theta}(\mathbf{I}_\tau)$ vanish due to consistency. The following lemmas state this observation in detail and establish the corresponding convergence results for the EOP thresholds, the total operating cost, and the probability of failure at replacement as k becomes large. Let $\hat{\theta}(\mathbf{I}_\tau) \xrightarrow{p} \theta_1$, for $\theta_1 \in \Theta$.

Lemma 5.12. *As $k \to \infty$,*

(a) *For $\theta_1 \in \Theta$, $\hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta_1)}\right) \xrightarrow{p} \theta_1$;*

(b) *For $\theta_1 \in \Theta$, $\frac{1}{k}\xi_{k,last}^P(\theta_1) \xrightarrow{p} \frac{1}{k}\xi_k^A(\theta_1)$;*

(c) $\frac{1}{k}\int_\Theta \frac{1}{\mu(\mathbf{y})}\left(\xi_{k,last}^P(\mathbf{y}) + \mathbb{E}\left[Y(\xi_{k,last}^P(\mathbf{y}))\right]\right) dF_\theta(\mathbf{y}) \xrightarrow{p} \frac{1}{k}\int_\Theta \frac{1}{\mu(\mathbf{y})}\left(\xi_k^A(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi_k^A(\mathbf{y})\right)\right]\right) dF_\theta(\mathbf{y})$;

(d) $\frac{1}{k}\int_\Theta r_k\left(\xi_{k,last}^P(\mathbf{y}), \mathbf{y}\right) dF_\theta(\mathbf{y}) \xrightarrow{p} \frac{1}{k}\int_\Theta r_k\left(\xi_k^A(\mathbf{y}), \mathbf{y}\right) dF_\theta(\mathbf{y})$;

(e) $\int_\Theta \mathbb{P}\left\{Y(\mathbf{y}) > \phi_k\left(\xi_{k,last}^P(\mathbf{y})\right)\right\} dF_\theta(\mathbf{y}) \xrightarrow{a.s.} \int_\Theta \mathbb{P}\left\{Y(\mathbf{y}) > \phi_k\left(\xi_k^A(\mathbf{y})\right)\right\} dF_\theta(\mathbf{y})$.

Lemma 5.12 leads to Theorem 5.6.

5.7 Simulation Study Results

In this section, we present the results of an extensive simulation study. The objectives of this study are threefold.

1. To evaluate the performance of the asymptotically optimal adaptive threshold policy relative to the optimal policy under the same filtration, when computation of the latter is feasible;
2. To compare the performance of the asymptotically optimal adaptive threshold policy with the Oracle optimal policy;
3. To analyze the optimality and prediction gaps associated with the Oracle's asymptotically optimal policy.

In this study, we assume that decision makers have access to the true hyperparameters $(\alpha_0, \beta_0, a_0, b_0)$. This assumption allows us to isolate the impact of the policies of interest on the cost rates from the effect of estimation errors in the prior, which would arise if the prior were learned from historical data. The latter effect is examined in Section 5.8, where we observed that the EOP performs remarkably well under such conditions.

Our study addresses two settings in terms of degradation state spaces: discrete, i.e. $1-$lattice, and continuous. The discrete setting is characterized by a compound Poisson process with geometric compounding, whereas the continuous setting is characterized by a compound Poisson process with exponential compounding. Example 1 details the computation of the EOP thresholds in the discrete setting, and Example 2 extends this to the continuous setting. The optimal policy $g^*(\mathcal{F})$ can be computed in a tractable manner using the integrated Bayes approach introduced in Drent et al. (2023a); however, this is feasible only for the discrete setting. Hence, we do not compute $g^*(\mathcal{F})$ for the continuous state space.

A tractable implementation of the integrated Bayes approach requires truncating both the cumulative number of shocks and the ages of the components, whereas the asymptotically optimal adaptive threshold policy does not. This truncation, however, affects the resulting updated posteriors. The cost-rate of the EOP is calculated under both the full observations, denoted $g(\xi^{\mathrm{P}})$, and the truncated observations, denoted $g(\xi^{\mathrm{P,trunc}})$, to assess the impact of truncation.

We are interested in the regret (relative to the Oracle) of the adaptive policies $\xi \in \{\xi^{\mathrm{P}}, \xi^{\mathrm{P,trunc}}, \xi^{\mathrm{IB}}\}$, defined as $\%\mathrm{REG}(\xi) := 100 \cdot (g(\xi) - g^*(\mathcal{O}))\,/g^*(\mathcal{O})$, which serves as a common performance metric across all policies and both discrete and continuous settings. As noted earlier, $\%\mathrm{REG}\,(\xi^{\mathrm{IB}})$ is only available for discrete state space. We are also interested in the optimality gap and prediction gap of ξ^{A} with respect to ξ^{O}, defined as $\%OPT := 100 \cdot (g\,(\xi^{\mathrm{A}}) - g^*(\mathcal{O}))\,/g^*(\mathcal{O})$ and $\%PRED := 100 \cdot (\gamma^* - g\,(\xi^{\mathrm{A}}))\,/g\,(\xi^{\mathrm{A}})$,

respectively. $\%OPT$ measures the relative additional cost incurred by using ξ^{A} instead of ξ^{O} as the Oracle's optimal decision, while $\%PRED$ quantifies how far the approximated Oracle's optimal cost-rate γ^* deviates from the true value $g\left(\xi^{\mathrm{A}}\right)$.

We construct a comprehensive testbed for our simulation study, shown in Table 5.1, and conduct experiments across all combinations of the parameter values, with $c_p = 1$. The shock arrival process follows a gamma prior with $\alpha_0 = \beta_0$, yielding a mean of 1. For instances with a discrete state space, the damage process follows a Beta prior, whereas for instances with a continuous state space, it follows a Gamma prior, with $a_0 = b_0$ in both cases, which gives means of 0.5 and 1, respectively. We vary the coefficients of variation to assess the impact of the degree of heterogeneity across instances. We consider three operating cost functions $\ell_0(x) = 0$, $\ell_1(x) = \dfrac{c_p \min\{x, L\}}{L^2/2 + L + 1}$, and $\ell_2(x) = \dfrac{c_p(\min\{x, L\})^2}{L^3/3 + L^2 + L + 1}$. The selection of ℓ_1 and ℓ_2 is such that the expected total operating cost at failure is approximately the same value as c_p when $\mu(\theta) = 1$, and $\sigma^2(\theta)$ is negligible.

Table 5.1: Input parameters for numerical studies.

	Input parameter	No. of Choices	Values
1	Coefficient of variation for the prior of the shock arrival process, cv_ν^2	2	0.3, 0.6
2	Coefficient of variation for the prior of the damage process, $cv_{\theta_Z}^2$	2	0.04, 0.08
3	Failure threshold, L,	3	10,20,30
4	Additional corrective maintenance cost, c_f	2	4,9
5	Operating cost, ℓ	3	ℓ_0, ℓ_1, ℓ_2

We compute γ^* using Theorem 5.4(b) with Monte Carlo integration using 10^5 samples per integral. The value of g^* is obtained from Lemma 5.3(b), with Monte Carlo integration using 10^5 samples per integral and 4×10^3 simulated components per expectation function. This ensures that the confidence intervals for the computed γ^* and g^* remain small. We simulate 10^4 components for each policy $\xi \in \{\xi^{\mathrm{P}}, \xi^{\mathrm{P,trunc}}, \xi^{\mathrm{IB}}\}$ and compute the cost-rate across all components. Each simulation is repeated 10 times to keep the confidence intervals for the performance metric small. The policy ξ^{IB} is obtained by solving the optimality Equations (10) in Drent et al. (2023a), modified to account for the operating cost, with a discount factor of 0.99 and a truncation of the number of shocks and components' age at 40, 50, and 60 when the failure level is 10, 20, and 30, respectively.

The results of the simulation study are reported in Tables 5.2 and 5.3. The results in Tables 5.2 and 5.3 indicate that EOP thresholds perform well in both discrete and continuous settings, with a mean $\%\mathrm{REG}(\xi^{\mathrm{P}})$ of 6.59% and 3.3%, respectively. The regret of EOP relative to the integrated Bayes method is negligible, implying that EOP performs as well as the optimal policy in our experiments. As the failure threshold L increases from 10 to 30, mean $\%\mathrm{REG}(\xi^{\mathrm{P}})$ decreases from 11.05% to 3.06% in the discrete setting and from 6.1% to 1.4% in the continuous setting, consistent with our main convergence result (Theorem 5.6). Moreover, $\%\mathrm{REG}(\xi^{\mathrm{P}})$ increases with the corrective maintenance cost c_f, because higher c_f leads to smaller optimal thresholds, which in turn increase

Table 5.2: Results of the simulation study - discrete state space

		%REG $\left(\xi^{\mathrm{P}}\right)$			%REG $\left(\xi^{\mathrm{P,trunc}}\right)$			%REG $\left(\xi^{\mathrm{IB}}\right)$			$\%OPT$			$\%PRED$		
Input	Values	Min	Mean	Max	Min	Mean	Max	Min	Mean	Max	Min	Mean	Max	Min	Mean	Max
L	10	3.76	11.05	24.93	3.19	11.54	24.20	2.69	11.16	26.75	0.03	3.03	8.69	-13.88	-8.04	-0.11
	20	1.14	5.68	15.53	0.85	6.19	16.24	0.97	6.17	16.97	0.02	1.44	4.53	-8.47	-4.87	-0.27
	30	0.47	3.06	9.55	0.35	3.50	10.46	1.06	4.00	11.31	0.03	0.90	3.04	-5.79	-3.37	0.24
cv_{ν}^2	0.3	0.75	7.55	24.93	1.27	7.61	22.83	1.06	6.82	21.89	0.14	2.04	8.69	-11.12	-4.64	-0.09
	0.6	0.47	5.64	19.53	0.35	6.55	24.20	0.97	7.40	26.75	0.02	1.54	6.14	-13.88	-6.21	0.24
$cv_{\theta Z}^2$	0.04	0.47	5.82	19.74	0.35	6.51	20.42	0.97	6.45	20.23	0.18	1.73	6.74	-13.88	-5.28	-0.09
	0.08	0.69	7.37	24.93	0.59	7.65	24.20	1.06	7.77	26.75	0.02	1.84	8.69	-13.58	-5.57	0.24
c_f	4	0.47	4.99	16.55	0.35	5.46	16.66	0.97	5.86	20.59	0.02	1.58	7.07	-13.88	-5.44	0.06
	9	0.83	8.20	24.93	0.66	8.7	24.20	1.07	8.36	26.75	0.05	2.00	8.69	-13.58	-5.41	0.24
ℓ	ℓ_0	4.53	11.36	24.93	5.66	12.76	24.20	5.31	13.09	26.75	1.89	3.97	8.69	-1.76	-0.63	0.24
	ℓ_1	0.47	3.88	14.06	0.35	3.98	13.20	0.97	3.81	12.26	0.02	0.52	2.45	-13.29	-7.51	-3.96
	ℓ_2	0.65	4.54	14.98	0.53	4.50	13.88	1.09	4.43	11.79	0.18	0.87	2.70	-13.88	-8.14	-4.58
Total		0.47	6.59	24.93	0.35	7.08	24.20	0.97	7.11	26.75	0.02	1.79	8.69	-13.88	-5.43	0.24

Table 5.3: Results of the simulation study - continuous state space

		%REG $\left(\xi^{\mathrm{P}}\right)$			$\%OPT$			$\%PRED$		
Input	Values	Min	Mean	Max	Min	Mean	Max	Min	Mean	Max
L	10	2.42	6.11	12.72	0.28	0.77	1.82	-1.34	-0.58	-0.11
	20	0.70	2.51	6.81	0.11	0.66	2.55	-0.48	-0.18	0.19
	30	0.08	1.38	4.55	0.00	0.48	1.74	-0.73	-0.12	0.39
cv_{ν}^2	0.3	0.08	3.53	12.72	0.00	0.50	1.63	-1.19	-0.30	0.09
	0.6	0.28	3.14	12.61	0.11	0.77	2.55	-1.34	-0.29	0.39
$cv_{\theta_Z}^2$	0.04	0.08	3.00	10.05	0.00	0.79	2.55	-0.74	-0.19	0.39
	0.08	0.13	3.67	12.72	0.00	0.49	1.74	-1.34	-0.40	0.19
c_f	4	0.08	2.45	7.69	0.00	0.47	1.99	-0.78	-0.21	0.28
	9	0.29	4.22	12.72	0.02	0.81	2.55	-1.34	-0.37	0.39
ℓ	ℓ_0	2.12	5.71	12.72	0.05	1.14	2.55	-1.30	-0.24	0.39
	ℓ_1	0.28	2.17	8.02	0.00	0.41	0.93	-1.34	-0.38	-0.09
	ℓ_2	0.08	2.12	8.04	0.00	0.37	0.86	-1.17	-0.26	0.00
Total		0.08	3.33	12.72	0.00	0.64	2.55	-1.34	-0.29	0.39

the approximation gap between ξ^{A} and ξ^{O} and amplify estimation errors of $\hat{\theta}(\mathbf{I}_\tau)$. Both tables show that the EOP thresholds effectively adapt to situations where the system incurs operating costs. We observe from the results that the optimality gap of using ξ^{A} instead of ξ^{O} is small, indicating that ξ^{A} provides an excellent approximation to the Oracle's optimal thresholds. Finally, the prediction gaps are more significant when the state space is discrete than in the continuous one, where they are negligible.

5.8 Case Study Results

In this section, we evaluate the performance of our EOP thresholds using real degradation data from the case study on Interventional X-ray (IXR) systems presented in Drent et al. (2023a). Below, we provide a brief summary of this case study and refer the interested reader to Drent et al. (2023a) for a more detailed discussion.

IXR systems are advanced imaging platforms used in hospitals to support minimally invasive procedures, such as cardiac or vascular interventions. Their most critical and costly replacement components are X-ray tubes, which gradually lose efficiency over time as the tungsten filaments inside them evaporate through high-temperature heating. During operation, these filaments emit electrons that are directed toward a target to generate X-ray images. Unexpected tube failures can force procedures to stop or be rescheduled, delay critical patient care, and generate substantial corrective maintenance expenses.

Filament condition monitoring is performed primarily through real-time measurements of electrical resistance. The measurements are then processed to generate a one-dimensional degradation indicator for the filaments, which is stored in the database of the

IXR system. Major manufacturers of IXR systems, such as Philips Healthcare, require maintenance strategies that exploit these real-time measurements to prevent tube failure while maximizing their useful life.

Drent et al. (2023a) provide time series degradation data for 52 filaments obtained from lab experiments by Philips Healthcare. The set of these 52 time series is denoted $\mathcal{J}$ and its cardinality $|\mathcal{J}| = 52$. They describe how the data are preprocessed, including statistical tests to assess filament heterogeneity and to verify that electric shocks follow a Poisson process. They also justify modeling the damage process as geometric. Next, they show how, under this assumption, the policy adapted to $\mathcal{F}$ can be computed as the optimal solution of a Partially Observable Markov Decision Process (POMDP). The resulting policy is optimal within the assumed model, though not necessarily for the specific case data. This policy is called the Integrated Bayes Policy (IBP). The IBP specifies adaptive replacement thresholds that depend on the information available at time t, denoted by $\xi^{\text{IB}}(\mathbf{I}_t)$.

The condition-based maintenance problem in this case study aligns with Example 1, for which our EOP thresholds ξ^{P} can be efficiently applied. Notice that $\xi^{\text{P}}(\mathbf{I}_t)$ is significantly more computationally efficient than $\xi^{\text{IB}}(\mathbf{I}_t)$, often by more than an order of magnitude. In fact, $\xi^{\text{IB}}(\mathbf{I}_t)$ becomes numerically intractable as the failure level increases. The only computationally intensive step in applying $\xi^{\text{P}}(\mathbf{I}_t)$ is the calculation of γ^* via Theorem 5.4(b), which requires numerical integration. We perform this using Monte Carlo methods. However, γ^* needs to be computed only when the hyperparameters are updated The runtime to update the posterior and to evaluate $\xi^{\text{P}}(\mathbf{I}_t)$ is negligible. In contrast, solving a POMDP is NP-hard, so exact solutions might become computationally intractable for some realistic problem sizes.

In the following, we evaluate the performance of ξ^{P} relative to ξ^{IB}, which is the current state-of-the-art strategy, using the 52 degradation time-series provided by Drent et al. (2023a). This assessment automatically benchmarks our approach against the other state-of-the-art methods presented in Drent et al. (2023a). We focus on cost savings percentages, defined as $\%SAV = \dfrac{\bar{g}(\xi^{\text{IB}}) - \bar{g}(\xi^{\text{P}})}{\bar{g}(\xi^{\text{IB}})} \times 100$ where $\bar{g}(\cdot)$ denotes the average cost function. Note that $(\alpha_0, \beta_0, a_0, b_0)$ is not known a priori in this problem. We therefore use a sample from $\mathcal{J}$, called the training set, to estimate these hyperparameters via maximum likelihood estimation (MLE), as described in Appendix C of Drent et al. (2023a). Since $\%SAV$ depends on the estimated $(\alpha_0, \beta_0, a_0, b_0)$, we generate multiple training sets from $\mathcal{J}$, denoted by $\mathcal{J}_{\text{training}}$. We choose the sizes of the training sets $\left|\mathcal{J}_{\text{training}}\right|$ from $\{5, 10, 15\}$, and for each size, we randomly draw 150 subsets from $\mathcal{J}$, each of which provides an estimate of the hyperparameters. This procedure gives us 450 bootstrapping instances. For each instance, we compute $\%SAV$ based on the training set, denoted $\mathcal{J}_{\text{test}} := \mathcal{J} \setminus \mathcal{J}_{\text{training}}$. c_p and c_f are chosen as 1 and 4, respectively. We compute the integrals in Equation (5.18) using Monte Carlo sampling with 5×10^4 samples, which gives narrow 95% confidence

intervals. The results of our study are presented in Table 5.4.

Table 5.4: Results of the case study.

$\left\lvert\mathcal{J}_{\text{training}}\right\rvert$	$\%SAV$
5	-0.35
10	1.09
15	1.17
Total	0.63

Table 5.4 shows that our asymptotically optimal adaptive threshold policy slightly reduces the average maintenance cost-rate by 0.63% compared to the integrated Bayes policy across the bootstrapping instances. The results indicate that the integrated Bayes policy performs slightly better for small training sets, i.e., when $\left\lvert\mathcal{J}_{\text{training}}\right\rvert = 5$, whereas the asymptotically optimal adaptive threshold policy slightly outperforms the integrated Bayes policy for larger training sets, i.e., when $\left\lvert\mathcal{J}_{\text{training}}\right\rvert = 10$ or 15. We observe that these results remain consistent when using smaller sample sizes to compute the integrals in Equation (5.18), for example, as small as 10^4. These results demonstrate that our asymptotically optimal adaptive threshold policy is not only substantially faster than the integrated Bayes policy but also achieves strong cost-rate performance on real data. From a managerial perspective, this combination of computational and cost-rate efficiency makes it highly suitable for real-time maintenance decision making.

5.9 Conclusion

In this chapter, we have studied the condition-based maintenance (CBM) problem for a heterogeneous population of components that undergo non-negative i.i.d. degradation processes with parameters unknown to the decision maker. At equidistant time epochs, the decision maker receives real-time degradation information and must choose between replacing the operating component or allowing it to continue functioning. The cost of replacing a failed component is considerably higher than that of a healthy component, and the system experiences a non-negative operating cost that rises with the degradation level. At each time epoch, the degradation data is collected and used to learn the unknown parameters of the degradation process via a consistent estimator. The decision maker then determines the maintenance action to minimize the cost-rate (long-run average cost). In general, this problem can be solved using a partially observable Markov decision process (POMDP). It is well known that POMDPs suffer from the curse of dimensionality.

We have proposed the Estimated Oracle's Optimal Policy (EOP) that chooses the CBM actions based on estimates of the optimal decision of an Oracle who has the full information of the true parameters of the degradation processes. Furthermore, we have introduced a scaling regime in which the failure level, cost parameters, and total oper-

ating costs grow jointly. This scaling regime captures the practical settings where the components' lifetime and the maintenance costs are much larger than the frequency of degradation measurements. This common setting also represents conditions under which conventional POMDP approaches become intractable as the state space grows large. We have proved that regret of the EOP, i.e., the difference between its cost-rate and that of the Oracle, converges to zero in our scaling regime, which also implies its asymptotic optimality. We have evaluated the performance of the EOP on real degradation data of IXR systems relative to the state-of-the-art Integrated Bayes policy (Drent et al., 2023a), which computes replacement policies using a POMDP model. Our results show that employing the EOP, on average, achieves a 0.63% cost reduction compared to the Integrated Bayes policy. We have conducted an extensive numerical experiment including discrete and continuous state spaces to test the performance of the EOP the Oracle's optimal policy. We have observed that the relative regret of our EOP is small in both types of state spaces, with the average relative regret of 6.59% in discrete state space and 3.3% in continuous state space. We have not been able to statically distinguish the regret of the EOP from the POMDP's (on average 7.11%) which could only be computed in the discrete setting. To the best of our knowledge, our EOP is the first data-driven CBM policy that can accommodate a wide range of degradation and parameter learning processes, provides an asymptotic performance guarantee, and has achieved remarkable results on real-life degradation data as well as numerical experiments across both discrete and continuous state spaces.

5.A Lemma 5.13: On the Finite Expected Lifetime of Components

Lemma 5.13. *Under any policy π, the expected lifetime and the expected maintenance cost of an operating component $i \in \mathbb{N}$ are finite, i.e., $\mathbb{E}[T_i] < \infty$, and $\mathbb{E}[C_i] < \infty$.*

Proof of Lemma 5.13. We upperbound $\mathbb{E}[T_i]$ using the expected lifetime of the component under the policy that only replaces upon failure. Note that $T_i(L, \theta_i)$ is the time until failure of the i-th component. Let U_i be the renewal measure of $F_X(\cdot \mid \theta_i)$ defined by $U_i(x) := \sum_{\tau=0}^{\infty} F_X^{*\tau}(x \mid \theta_i)$ with $F_X^{*\tau}(\cdot \mid \theta_i)$ denoting the τ-fold convolution of $F_X(\cdot \mid \theta_i)$, and F_X^{*0} the distribution function degenerate at the origin. Next, For all $x \geq 0$,

$$\frac{1}{\mu(\theta_i)}x \leq U_i(x) \leq \frac{1}{\mu(\theta_i)}x + \frac{2\mathbb{E}\left[Y(\theta_i)\right]}{\mu(\theta_i)}, \tag{5.38}$$

where $\mathbb{E}[Y(\theta_i)]$ is given by (5.16). The left inequality in (5.38) follows from the non-negativity of the residual life after x (cf. V.6. Asmussen, 2003), and the right inequality in (5.38) is Lorden's inequality (Lorden, 1970). Moreover, from renewal theory,

$$\mathbb{E}[T\left(L, \theta_i\right)] = U_i(L) + 1 \tag{5.39}$$

(cf. Theorem V.2.4. Asmussen, 2003). Equation (5.38) together with (5.39) imply

$$0 \leq \mathbb{E}[\mathbb{E}[T(L,\theta_i) \mid \theta_i]] \leq L\mathbb{E}\left[\frac{1}{\mu(\theta_i)}\right] + 2\mathbb{E}\left[Y(\theta_i)\right] + 1 < \infty, \tag{5.40}$$

where the last inequality follows from Assumption A.7. Next,

$$\mathbb{E}[C_i] \leq c_p + c_f + \mathbb{E}\left[\sum_{\tau=1}^{T(L,\theta_i)} \ell(L)\right] = c_p + c_f + \ell(L)\mathbb{E}[T(L,\theta_i)] < \infty. \tag{5.41}$$

The first inequality in 5.41 follows from the definition of C_i according to Equation (5.1) and the upperbound on $\ell(\cdot)$ as per Assumption A.9, and the second inequality from (5.40). $\square$

5.B Proof of Theorem 5.1

Consider the following Bellman optimality equations under the discounted total cost criterion with discount factor $\alpha \in (0,1)$

$$v_\alpha(s,\theta) = \begin{cases} \min\left\{\ell(s) + \alpha\mathbb{E}[v_\alpha(s + X_{1,1},\theta) \mid \theta_1 = \theta], \ell(s) + c_p + \alpha\overline{v}_\alpha\right\}, & \forall s \leq L, \theta \in \Theta \\ \ell(s) + c_p + c_f + \alpha\overline{v}_\alpha, & \forall s > L, \theta \in \Theta, \end{cases} \tag{5.42}$$

where

$$\overline{v}_\alpha = \int_\Theta \mathbb{E}[v_\alpha(X_{1,1},\theta_1) \mid \theta_1 = \mathbf{y}] dF_\theta(\mathbf{y})\Big\}$$

Define $m_\alpha := \inf_{(s,\theta)\in[0,L]\times\Theta} v_\alpha(s,\theta)$, and $w_\alpha := v_\alpha(s,\theta) - m_\alpha$. We intend to use the results of Schäl (1993) to show that

$$\lim_{\alpha\to 1}(1-\alpha)m_\alpha = \lim_{\alpha\to 1}(1-\alpha)v_\alpha(s,\theta) = g^*, \qquad \forall (s,\theta) \in [0,L]\times\Theta. \tag{5.43}$$

In particular, our aim is to use Theorem 3.8 in Schäl (1993). We restate this Theorem as the following lemma, in a form that is more convenient for our analysis.

Lemma 5.14. *(adapted from Theorem 3.8. Schäl, 1993) For all states $(s,\theta) \in [0,L]\times\Theta$, assume that (A) $\sup_{\alpha<1} w_\alpha(s,\theta) < \infty$, and (B) the decision space is finite. Then there exists a stationary policy π^* that is average optimal in the sense that $g(\pi^*) = g^*$. Furthermore, the convergence result* (5.43) *holds.*

The critical step is to prove that condition (A) in Lemma 5.14 holds for our problem. To do this, we use Lemma 4.1 in Schäl (1993) which is restated (a customized version) as follows. For some $\eta > 0$, let $\mathcal{A}_{S,\eta} \subseteq [0,L]$ and $\mathcal{A}_{\theta,\eta} \subseteq \Theta$ be such that $v(s,\theta) \leq m_\alpha + \eta$ for all $s \in \mathcal{A}_{S,\eta}$, and $\theta \in \mathcal{A}_{\theta,\eta}$. Notice that $v(s,\theta)$ is non-decreasing in s and that $\mathcal{A}_{S,\eta}$ and $\mathcal{A}_{\theta,\eta}$ cannot be empty as η is positive and m_α is the infimum of $v(s,\theta)$ over the state space. Define the stopping time $n_\alpha := \inf_{i\in\mathbb{N}}\{i : S_{T_i+1} \in \mathcal{A}_{s,\eta}, \theta_{T_i+1} \in \mathcal{A}_{\theta,\eta}\}$.

Lemma 5.15. *(adapted from Lemma 4.1. Schäl, 1993) For* $\eta \geq 0$, $\alpha < 1$, $(s, \theta) \in [0, L] \times \Theta$*:* $w_\alpha(s, \theta) \leq \eta + \inf_\pi \mathbb{E}[\sum_{i=1}^{n_\alpha} C(S_{T_i(\pi)+1}, \theta_{T_i(\pi)+1})]$.

Observe that the sequences $\{S_{T_i+1}\}_{i\in\mathbb{N}}$ and $\{\theta_{N(T_i+1)}\}_{i\in\mathbb{N}}$ are i.i.d. Therefore, n_α is a geometric random variable and has a finite expectation. Define $t_\alpha = \sum_{i=1}^{n_\alpha} T_i$. By Wald's identity and Lemma 5.13, $\mathbb{E}[t_\alpha] = \mathbb{E}[n_\alpha]\mathbb{E}[T_i] < \infty$. Now, By Lemma 5.15, $w_\alpha(s, \theta) \leq \eta + (c_p + c_f + \ell(L))\mathbb{E}[t_\alpha] < \infty$ which proves that condition (A) in Lemma 5.14 holds. Condition (B) in Lemma 5.14 obviously holds, since we have binary decisions for each state. Thus, Part (a) of Theorem 5.1 follows from Lemma 5.14.

Finally, there exists a threshold policy that satisfies the optimality Equations (5.42) that complete the proof for Part (b). □

5.C Proof of Lemma 5.2

Let π be a policy that determines the replacement actions based on the available information of the operating component. Under this structure, maintenance interventions disregard any history of actions or observations associated with previously installed components. Consequently, the sequences $\{C_i\}_i \in \mathbb{N}$ and $\{T_i\}_{i\in\mathbb{N}}$ are i.i.d., as each component is drawn from the same population. The result then follows directly from the renewal–reward theorem [cf. Theorem VI.3.1.][](Asmussen, 2003) combined with Lemma 5.13.

Remark 5.16. This argument is more general than the Lemma's statement and remains valid for any replacement policy that depends only on the information of the operating component. In that case, $g(\xi)$ can be expressed as

$$g(\xi) = \frac{\int_\Theta \mathbb{E}[\mathbb{E}[C(\xi(\mathbf{y}), \mathbf{y}) \mid \xi(\mathbf{y})]] dF_\theta(\mathbf{y})}{\int_\Theta \mathbb{E}[\mathbb{E}[T(\xi(\mathbf{y}), \mathbf{y}) \mid \xi(\mathbf{y})]] dF_\theta(\mathbf{y})}, \tag{5.44}$$

(cf. Equation 5.12). We will use this result later when we introduce our adaptive policies, in which the replacement decisions are made based on the available degradation data of the current component. □

5.D Proof of Lemma 5.3

The proof of Lemma 5.3 is adapted from the Appendix of Aven and Bergman (1986). We present a tailored version for completeness.

Lemma 5.17. *(a) For* $\lambda > g^*$, $\int_\Theta \tilde{h}(\lambda, \xi_\lambda, \mathbf{y}) dF_\theta(\mathbf{y}) < 0$*, and equivalently,* $g(\xi_\lambda) < \lambda$.

(b) $\int_\Theta \tilde{h}(\lambda, \xi_\lambda, \mathbf{y}) dF_\theta(\mathbf{y})$ *is non-increasing and concave in* λ.

(c) For $\lambda_1, \lambda_2 \in \mathbb{R}$ *satisfying* $g^* < \lambda_1 < \lambda_2$, *we have*

$$(\lambda_2 - \lambda_1) \int_\Theta \mathbb{E}\left[T(\xi_{\lambda_1}, \mathbf{y})\right] dF_\Theta(\mathbf{y}) < -\int_\Theta \mathbb{E}\left[\tilde{h}\left(\lambda_2, \xi_{\lambda_2}, \mathbf{y}\right)\right] dF_\Theta(\mathbf{y}) < 0.$$

By the arguments in Appendix 5.C and the cost structure, for all $\xi \in \Xi$

$$0 < \int_\Theta \mathbb{E}\left[C(\xi(\mathbf{y}), \mathbf{y}) dF_\Theta(\mathbf{y})\right] < \infty, \qquad \text{and} \qquad 1 \leq \int_\Theta \mathbb{E}\left[T(\xi(\mathbf{y}), \mathbf{y}) dF_\Theta(\mathbf{y})\right] < \infty.$$

Therefore, $0 < g(\xi) < \infty$, and for all $\lambda \in \mathbb{R}$ we have $\left| \int_\Theta \mathbb{E}\left[\tilde{h}\left(\lambda, \xi, \mathbf{y}\right)\right] dF_\Theta(\mathbf{y}) \right| < \infty$.

Proof of Lemma 5.17.

(a) Let $\xi \in \Xi$ satisfy $g^* \leq g(\xi) < \lambda$. Then,

$$\int_\Theta \tilde{h}(\lambda, \xi_\lambda, \mathbf{y}) dF_\theta(\mathbf{y}) \leq \int_\Theta \tilde{h}(\lambda, \xi, \mathbf{y}) dF_\theta(\mathbf{y}) < \int_\Theta \tilde{h}(g(\xi), \xi_\lambda, \mathbf{y}) dF_\theta(\mathbf{y}) = 0.$$

By Equation (5.13), $g(\xi_\lambda) < \lambda$.

(b) We first prove the monotonicity. Let $\lambda_1 \leq \lambda_2$, then

$$\int_\Theta \tilde{h}(\lambda_2, \xi_{\lambda_2}, \mathbf{y}) dF_\theta(\mathbf{y}) \leq \int_\Theta \tilde{h}(\lambda_2, \xi_{\lambda_1}, \mathbf{y}) dF_\theta(\mathbf{y}) \leq \int_\Theta \tilde{h}(\lambda_1, \xi_{\lambda_1}, \mathbf{y}) dF_\theta(\mathbf{y}).$$

Next we show the concavity. Let for $\lambda_1, \lambda_2 \in \mathbb{R}$ and $\alpha \in [0, 1]$, $\lambda = \alpha\lambda_1 + (1-\alpha)\lambda_2$, then

$$\begin{aligned}\int_\Theta \tilde{h}(\lambda, \xi_\lambda, \mathbf{y}) dF_\theta(\mathbf{y}) =& \alpha \int_\Theta \tilde{h}(\lambda_1, \xi_\lambda, \mathbf{y}) dF_\theta(\mathbf{y}) + (1-\alpha) \int_\Theta \tilde{h}(\lambda_2, \xi_\lambda, \mathbf{y}) dF_\theta(\mathbf{y}) \\ \geq& \alpha \int_\Theta \tilde{h}(\lambda_1, \xi_{\lambda_1}, \mathbf{y}) dF_\theta(\mathbf{y}) + (1-\alpha) \int_\Theta \tilde{h}(\lambda_2, \xi_{\lambda_2}, \mathbf{y}) dF_\theta(\mathbf{y}).\end{aligned}$$

(c) Let $g^* < \lambda_1 < \lambda_2$, then

$$\begin{aligned}&\int_\Theta \tilde{h}(\lambda_2, \xi_{\lambda_2}, \mathbf{y}) dF_\theta(\mathbf{y}) \leq \\ &\int_\Theta \tilde{h}(\lambda_2, \xi_{\lambda_1}, \mathbf{y}) dF_\theta(\mathbf{y}) = \int_\Theta \mathbb{E}[C(\lambda_1, \xi_{\lambda_1}, \mathbf{y})] - \lambda_2 \mathbb{E}[T(\lambda_1, \xi_{\lambda_1}, \mathbf{y})] dF_\theta(\mathbf{y}) \\ &= \int_\Theta \mathbb{E}[C(\lambda_1, \xi_{\lambda_1}, \mathbf{y})] - \lambda_1 \mathbb{E}[T(\lambda_1, \xi_{\lambda_1}, \mathbf{y})] - (\lambda_2 - \lambda_1)\mathbb{E}[T(\lambda_1, \xi_{\lambda_1}, \mathbf{y})] dF_\theta(\mathbf{y}) \\ &= \int_\Theta \tilde{h}(\lambda_1, \xi_{\lambda_1}, \mathbf{y}) dF_\theta(\mathbf{y}) - \int_\Theta (\lambda_2 - \lambda_1)\mathbb{E}[T(\lambda_1, \xi_{\lambda_1}, \mathbf{y})] dF_\theta(\mathbf{y}) \\ &< -\int_\Theta (\lambda_2 - \lambda_1)\mathbb{E}[T(\lambda_1, \xi_{\lambda_1}, \mathbf{y})] dF_\theta(\mathbf{y})\end{aligned}$$

Notice that by Part (a), both $\int_\Theta \tilde{h}(\lambda_2, \xi_{\lambda_2}, \mathbf{y}) dF_\theta(\mathbf{y})$ and $\int_\Theta \tilde{h}(\lambda, \xi_\lambda, \mathbf{y}) dF_\theta(\mathbf{y})$ are strictly positive which completes the proof. □

Proof of Lemma 5.3. Let $\lambda_1 \in \mathbb{R}$, then $\lambda_2 = g(\xi_{\lambda_1}) \geq g^*$. We can verify by induction and using part (a) of Lemma 5.17 that for all $j \geq 2$, $g^* \leq \lambda_{j+1} = g(\xi_{\lambda_j}) < \lambda_j$. If $\lambda_j = g^*$, then $\lambda_{j+1} = \lambda_j = g^*$. Therefore, there exists a $\lambda_\infty \geq g^*$ such that $\lambda_\infty \leq \lambda_j$ for all $j \in \mathbb{N}$ and $\lim_{j\to\infty} \lambda_j - \lambda_\infty = 0$. Notice that by Equation (5.40), $\int_\Theta \mathbb{E}\left[T(\xi_{\lambda_j}, \mathbf{y})\right] dF_\Theta(\mathbf{y})$ is bounded for all $j \in \mathbb{N}$. Thus, by Parts (b) and (c) of Lemma 5.17,

$$\lim_{j\to\infty} \int_\Theta \mathbb{E}\left[\tilde{h}\left(\lambda_j, \xi_{\lambda_j}, \mathbf{y}\right)\right] dF_\Theta(\mathbf{y}) = 0,$$

which establishes $\lambda_\infty = g^*$. $\square$

5.E Proof of Lemma 5.8

We provide the proof for the continuous degradation process; the argument for the d-lattice case is similar, with integration over the degradation level replaced by summation.

Proof of Part (a): For notational convenience, we omit the conditioning on θ_1, θ, and the index 1. We have,

$$\begin{aligned}
\mathbb{E}\left[\sum_{\tau=1}^{T(\xi)} \ell(S_\tau, k)\right] =& \sum_{\tau=1}^{\infty} \mathbb{E}\left[\ell(S_\tau, k); T\left(\xi\right) \geq \tau\right] \\
=& \int_0^\infty \ell(x,k) dF_X(x) + \sum_{\tau=1}^{\infty} \int_0^\xi \int_0^\infty \ell(s+x,k) dF_X(x) dF_X^{*\tau}(s) \\
=& \mathbb{E}[\ell(X,k)] + \int_0^\xi \mathbb{E}[\ell(X+s,k)] dU(s) \qquad (5.45)
\end{aligned}$$

Define for $k > 0$

$$w_k(s) := \mathbb{E}\left[\tilde{\ell}\left(\frac{X+s}{k}\right)\right], \qquad (5.46)$$

and observe that $w_k(\cdot) \leq \tilde{\ell}\left(\tilde{L}\right)$. Then Equation (5.45) gives

$$\mathbb{E}\left[\sum_{\tau=1}^{T\left(\xi^{(k)}\right)} \ell(S_\tau, k)\right] = w_k(0) + \int_0^{\xi^{(k)}} w_k(s) dU(s). \qquad (5.47)$$

Note that $w_k(\cdot)$ is not directly Riemann integrable, since $\int_0^\infty w_k(s) ds$ diverges, for every $k > 0$. Thus the key renewal theorem (cf. Theorem V.4.7. Asmussen, 2003) cannot be applied immediately, on the RHS of Equation (5.47). We, therefore, adopt a more elaborate approach to establish the result.

By renewal theory there exists a function $U_c : \mathbb{R}_+ \to \mathbb{R}_+$ such that

$$U(s) = \frac{s}{\mu} + \frac{\mathbb{E}[Y]}{\mu} + U_c(s),$$

with

$$|U_c(s)| \leq \frac{\mathbb{E}[Y]}{\mu}, \tag{5.48a}$$

$$\lim_{s\to\infty} U_c(s) = 0, \tag{5.48b}$$

(cf. Propositions V.6.1 and V.6.2 Asmussen, 2003). Then we have

$$\int_0^{\xi^{(k)}} w_k(s)dU(s) - \frac{1}{\mu}\int_0^{\xi^{(k)}} w_k(s)ds = \int_0^{\xi^{(k)}} w_k(s)dU_c(s). \tag{5.49}$$

We use integration by part to obtain

$$\int_0^{\xi^{(k)}} w_k(s)dU_c(s) = w_k\left(\xi^{(k)}\right)U_c\left(\xi^{(k)}\right) - w_k(0)U_c(0) - \int_0^{\xi^{(k)}} U_c(s)dw_k(s). \tag{5.50}$$

By the lemma's assumption as $k \to \infty$, $\xi^{(k)} \to \infty$. Thus, by (5.48b), and boundedness of $w_k(\cdot)$

$$\lim_{k\to\infty} w_k\left(\xi^{(k)}\right)U_c\left(\xi^{(k)}\right) = 0 \tag{5.51}$$

Next, observe that, $\tilde{\ell}\left(\frac{X}{k}\right) \leq \tilde{\ell}(\tilde{L})$, and

$$\lim_{k\to\infty} \tilde{\ell}\left(\frac{X}{k}\right)U_c(0) = 0,$$

since $U_c(0)$ is bounded by (5.48a). Therefore by Lebesgue's dominated convergence theorem

$$\lim_{k\to\infty} w_k(0)U_c(0) = \lim_{k\to\infty} \mathbb{E}\left[\tilde{\ell}\left(\frac{X}{k}\right)\right]U_c(0) = 0 \tag{5.52}$$

(cf. Theorem 16.4 Billingsley, 1995).

The remainder of the proof is devoted to showing that the final term on the RHS of (5.50) vanishes as k approaches infinity. We organize the remainder of the proof into two steps.

Step I: We construct a sequence $\left\{s_k \in \left[0, \tilde{L}\right]\right\}_{k\gg 0}$ such that as $k \to \infty$, we have $s_k \to 0$, and $ks_k \to \infty$:
By continuity of $\tilde{\ell}$, for every $\epsilon_0 > 0$ there exists $s_0 \in (0, \tilde{L}]$ such that for all $s \in [0, s_0)$, we have $\tilde{\ell}(s) < \epsilon_0$. Additionally, by Assumption A.8, for every $\epsilon_r > 0$ there exists $s_r \in (0, \tilde{L}]$ such that for all $s \in [0, s_r)$, we have $\tilde{\ell}(s)/s < \epsilon_r/2$. Choose $k > 0$ such that $1/k < \min\{\epsilon_0, \epsilon_r/2\}$. Then there exist $s_k \in \left(0, \frac{\min\{s_0, s_r\}}{2}\right)$ such that $\tilde{\ell}(2s_k) = 1/k$ and $ks_k = s_k/\tilde{\ell}(2s_k) > 1/\epsilon_r$.

This construction implies, also, that $\xi^{(k)}/ks_k \to \infty$ as $k \to \infty$, because by the lemma's assumption $\lim_{k\to\infty} \xi^{(k)}/k > 0$. Hence, we restrict attention to sufficiently large k with $\xi^{(k)} \geq ks_k$.

Step II: We partition the integration domain of the RHS of (5.50) into $[0, ks_k]$ and $(ks_k, \xi^{(k)}]$. First, for any $s \in [0, ks_k)$ we find an upper bound for $w_k(s)$ as follows

$$\begin{aligned} w_k(s) =& \int_0^\infty \tilde{\ell}\left(\frac{x+s}{k}\right) dF_X(x) = \int_0^{ks_k} \tilde{\ell}\left(\frac{x+s}{k}\right) dF_X(x) + \int_{ks_k}^\infty \tilde{\ell}\left(\frac{x+s}{k}\right) dF_X(x) \\ \leq& \tilde{\ell}(2s_k) F_X(ks_k) + \tilde{\ell}\left(\tilde{L}\right)(1 - F_X(ks_k)) \leq \frac{1}{k} + \tilde{\ell}\left(\tilde{L}\right)(1 - F_X(ks_k)). \end{aligned}$$

This together with (5.48a) leads us to

$$\left| \int_0^{ks_k} U_c(s) dw_k(s) \right| \leq \kappa_1(k), \tag{5.53}$$

where

$$\kappa_1(k) := \frac{\mathbb{E}[Y]}{\mu} \left(\frac{1}{k} + \tilde{\ell}\left(\tilde{L}\right)(1 - F_X(ks_k)) \right). \tag{5.54}$$

For the interval $\left(ks_k, \xi^{(k)}\right]$ we have

$$\left| \int_{ks_k}^{\xi^{(k)}} U_c(s) dw_k(s) \right| \leq \kappa_2(k), \tag{5.55}$$

where

$$\kappa_2(k) := \left| \max_{s \in \left(ks_k, \xi^{(k)}\right]} U_c(s) \right| \tilde{\ell}\left(\tilde{L}\right). \tag{5.56}$$

Now by (5.53) through (5.56)

$$\left| \int_0^{\xi^{(k)}} U_c(s) dw_k(s) \right| \leq \left| \int_0^{ks_k} U_c(s) dw_k(s) \right| + \left| \int_{ks_k}^{\xi^{(k)}} U_c(s) dw_k(s) \right| \leq \kappa_1(k) + \kappa_2(k). \tag{5.57}$$

By construction in Part I along with (5.48b),

$$\lim_{k\to\infty} \kappa_1(k) = \lim_{k\to\infty} \kappa_2(k) = 0. \tag{5.58}$$

Equations (5.49), (5.51), (5.52), and (5.58) give

$$\lim_{k\to\infty} \left(\int_0^{\xi^{(k)}} w_k(s) dU(s) - \frac{1}{\mu} \int_0^{\xi^{(k)}} w_k(s) ds \right) = 0,$$

which completes the proof for Part (a) of the lemma.

Proof of Part (b): By (5.47), (5.49), and (5.50) it is sufficient to show that the following limit exists and equals zero:

$$\lim_{k\to\infty} \int_\Theta \left(\left| w_{k,\mathbf{y}}\left(\xi^{(k)}(\mathbf{y})\right) U_c\left(\xi^{(k)}(\mathbf{y}), \mathbf{y}\right) \right| + \left| w_{k,\mathbf{y}}(0,\mathbf{y}) U_c(0,\mathbf{y}) \right| + \left| \kappa_1(k,\mathbf{y}) + \kappa_2(k,\mathbf{y}) \right| \right) dF_\theta(\mathbf{y}),$$

where $w_{k,\theta}(s)$, $U_c(\cdot,\theta)$, $\kappa_1(k,\theta)$, and $\kappa_2(k,\theta)$, are defined as in the proof of Part (a), except here they are parameterized by $\theta \in \Theta$. Define $K : \Theta \to \mathbb{R}_+$ by

$$K(\theta) := \frac{\mathbb{E}\left[Y(\theta)\right]}{\mu(\theta)} = \frac{\sigma^2(\theta) + \mu^2(\theta)}{2\mu^2(\theta)}.$$

For $k \gg 0$ and $\theta \in \Theta$:

$$\left| w_{k,\theta}\left(0\right) U_c\left(0,\theta\right) \right| \leq \tilde{\ell}\left(\tilde{L}\right) K(\theta), \tag{5.59a}$$

$$\left| w_{k,\theta}\left(\xi^{(k)}(\theta)\right) U_c\left(\xi^{(k)}(\theta),\theta\right) \right| \leq \tilde{\ell}\left(\tilde{L}\right) K(\theta), \tag{5.59b}$$

$$\left| \kappa_1(k,\theta) + \kappa(n,\theta) \right| \leq 2\tilde{\ell}\left(\tilde{L}\right) K(\theta). \tag{5.59c}$$

By Assumption A.7, K is integrable, i.e. $\int_\Theta K(\mathbf{y}) dF_\theta(\mathbf{y})$ is finite. Hence, the right-hand sides of (5.59a) through (5.59c) are integrable. Moreover, by the arguments in the proof of Part (a), the left-hand sides of (5.59a) through (5.59c) converge to zero almost everywhere as n grows to infinity. Thus by Lebesgue's dominated convergence theorem we have

$$\lim_{k\to\infty} \int_\Theta \left(\left| w_{k,\mathbf{y}}\left(\xi^{(k)}(\mathbf{y})\right) U_c\left(\xi^{(k)}(\mathbf{y}),\mathbf{y}\right) \right| + \left| w_{k,\mathbf{y}}(0,\mathbf{y}) U_c(0,\mathbf{y}) \right| + \left| \kappa_1(k,\mathbf{y}) + \kappa_2(k,\mathbf{y}) \right| \right) dF_\theta(\mathbf{y}) =$$
$$\int_\Theta \lim_{k\to\infty} \left(\left| w_{k,\mathbf{y}}\left(\xi^{(k)}(\mathbf{y})\right) U_c\left(\xi^{(k)}(\mathbf{y}),\mathbf{y}\right) \right| + \left| w_{k,\mathbf{y}}(0,\mathbf{y}) U_c(0,\mathbf{y}) \right| + \left| \kappa_1(k,\mathbf{y}) + \kappa_2(k,\mathbf{y}) \right| \right) dF_\theta(\mathbf{y}) = 0.$$

The last equality follows from the arguments in the proof of Part (a).

Proof of Part (c): Since any probability distribution function is bounded, the result follows directly from the bounded convergence theorem (cf. Theorem 16.5. Billingsley, 1995) and Lemma 5.7.

Proof of Part (d): Note that for any $k > 0$ and $\theta \in \Theta$, we have $\xi^{(k)} \in [0, k\tilde{L}]$. Thus, by Lorden's inequality (Lorden, 1970) we have for $k > 1$ and $\theta \in \Theta$

$$\left| \frac{1}{\mu(\theta)} \left(\frac{\xi^{(k)}(\theta)}{k} + \frac{\mathbb{E}\left[\tilde{Y}_1\left(\xi^{(k)}(\theta_1)\right) \middle| \theta_1 = \theta \right]}{k} \right) \right| \leq \frac{1}{\mu(\theta)} \left(\tilde{L} + 2\mathbb{E}[Y(\theta)] \right). \tag{5.60}$$

By Assumption A.7, the RHS of (5.60) is integrable. Therefore, applying Lebesgue's dominated convergence theorem together with Lemma 5.7 provides the result. □

5.F Proof of Lemma 5.9

Observe that for $k > 0$ and $\xi^{(k)} \in \Xi^{(k)}$

$$0 < \frac{\tilde{c}_p}{\tilde{L}\int_\Theta \frac{1}{\mu(\mathbf{y})} dF_\theta(\mathbf{y}) + 2\int_\Theta \frac{\mathbb{E}\left[\tilde{Y}(\mathbf{y})\right]}{\mu(\mathbf{y})} dF_\theta(\mathbf{y})} \leq g_k\left(\xi^{(k)}\right) \leq \frac{\tilde{c}_p + \tilde{c}_f + \tilde{\ell}(\tilde{L})}{\tilde{L}\int_\Theta \frac{1}{\mu(\mathbf{y})} dF_\theta(\mathbf{y})} < \infty \tag{5.61}$$

The lower bound on g_k in (5.61) results from minimizing the numerators and maximizing the denominators in (5.32) and applying Lorden's inequality Lorden (1970). The upper bound follows from the cost-rate of replacing at failure. The upper bound is finite and the lower bound is strictly positive under Assumptions A.7 and A.5, respectively. It follows from (5.61) that $\lim_{k\to\infty} g_k\left(\xi^{(k)}\right)$ exists and is positive for any sequence $\left\{\xi^{(k)}\right\}_{k>0}$. By Equation (5.32) we have the following.

$$g_k\left(\xi^{(k)}\right) = \frac{\tilde{c}_p + \int_\Theta \left(\tilde{c}_f \mathbb{P}\left\{\tilde{Y}_1\left(\xi^{(k)}(\theta_1)\right) > \phi_k\left(\xi^{(k)}(\theta_1)\right) \middle| \theta_1 = \mathbf{y}\right\} + \frac{1}{k}\mathbb{E}\left[\sum_{\tau=1}^{T\left(\xi^{(k)}(\theta_1),\theta_1\right)} \ell(S_{1,\tau},k) \middle| \theta_1 = \mathbf{y}\right]\right) dF_\theta(\mathbf{y})}{\int_\Theta \frac{1}{\mu(\mathbf{y})}\left(\frac{\xi^{(k)}(\mathbf{y})}{k} + \frac{\mathbb{E}\left[\tilde{Y}_1(\xi^{(k)}(\theta_1)) \middle| \theta_1 = \mathbf{y}\right]}{k}\right) dF_\theta(\mathbf{y})} \tag{5.62}$$

The equality in (5.62) follows from dividing both the numerator and denominator by k. By Lemma 5.8 and the growth rate of $\xi^{(k)}$ from Lemma's conditions, the limit of the numerator and denominator of are finite and positive as k approaches infinity. Therefore, $\lim_{k\to\infty} g_k\left(\xi^{(k)}\right)$ equals the ratio of the limits of the numerator and denominators in the RHS of 5.62. Hence

$$\begin{aligned}
\lim_{k\to\infty} g_k\left(\xi^{(k)}\right) &= \frac{\lim_{k\to\infty}\left(\tilde{c}_p + \int_\Theta \left(\tilde{c}_f \mathbb{P}\left\{Y(\mathbf{y}) > \phi_k\left(\xi^{(k)}(\mathbf{y})\right)\right\} + \frac{1}{k} r_k\left(\xi^{(k)},\mathbf{y}\right)\right) dF_\theta(\mathbf{y})\right)}{\lim_{k\to\infty}\left(\int_\Theta \frac{1}{\mu(\mathbf{y})}\left(\frac{\xi^{(k)}(\mathbf{y})}{k} + \frac{\mathbb{E}\left[Y(\mathbf{y})\right]}{k}\right) dF_\theta(\mathbf{y})\right)} \\
&= \lim_{k\to\infty} \frac{c_p(k) + \int_\Theta \left(c_f(k)\mathbb{P}\left\{Y(\mathbf{y}) > \phi_k\left(\xi^{(k)}(\mathbf{y})\right)\right\} + r_k\left(\xi^{(k)},\mathbf{y}\right)\right) dF_\theta(\mathbf{y})}{\int_\Theta \frac{1}{\mu(\mathbf{y})}\left(\xi^{(k)}(\mathbf{y}) + \mathbb{E}\left[Y(\mathbf{y})\right]\right) dF_\theta(\mathbf{y})} \\
&= \lim_{k\to\infty} \gamma_k\left(\xi^{(k)}\right)
\end{aligned} \tag{5.63}$$

The first equality in (5.63) follows from Lemma 5.8. □

5.G Proof of Lemma 5.10

We present the proof for Part (a); the argument for Part (b) is analogous.

Observe that $0 \le \liminf_{k\to\infty} \xi_k^{\mathrm{O}}/k \le \tilde{L}$. Therefore, it is sufficient to show that $\liminf_{k\to\infty} \xi_k^{\mathrm{O}}/k \neq 0$. In the remainder of the argument, we fix $\theta \in \Theta$, and suppress notation involving θ or conditioning on it, as well as the index for the component number. For $k > 0$ define

$$\tilde{h}_k(\lambda,\xi) := \tilde{c}_f \mathbb{P}\left\{\tilde{Y}(\xi^{(k)}) > \phi_k\left(\xi^{(k)}\right)\right\} + \frac{1}{k}\mathbb{E}\left[\sum_{\tau=1}^{T(\xi^{(k)})} \tilde{\ell}\left(\frac{S_\tau}{k}\right)\right] - \frac{\lambda}{\mu}\left(\frac{\xi^{(k)}}{k} + \frac{\mathbb{E}\left[\tilde{Y}(\xi^{(k)})\right]}{k}\right).$$

By (5.45) we have

$$\tilde{h}_k(\lambda,\xi) := \tilde{c}_f \mathbb{P}\left\{\tilde{Y}(\xi^{(k)}) > k\tilde{L} - \xi^{(k)}\right\} + \frac{1}{k}\mathbb{E}\left[\tilde{\ell}\left(\frac{X}{k}\right)\right] + \frac{1}{k}\int_0^{\xi^{(k)}} \mathbb{E}\left[\tilde{\ell}\left(\frac{X+s}{k}\right)\right] dU(s) - \frac{\lambda}{\mu}\left(\frac{\xi^{(k)}}{k} + \frac{\mathbb{E}\left[\tilde{Y}(\xi^{(k)})\right]}{k}\right). \quad (5.64)$$

By Lemma 5.3, ξ_n^{O} minimizes $\tilde{h}_k(g_k^*,\xi)$. In order to reach a contradiction, assume the contrary of the lemma, i.e., $\liminf_{k\to\infty} \xi_n^{\mathrm{O}}/k = 0$. This implies $\lim_{k\to\infty} \tilde{h}_k(g_k^*,\xi_n^{\mathrm{O}}) = 0$. The convergence of $\frac{1}{k}\int_0^{\xi^{(k)}} \mathbb{E}\left[\tilde{\ell}\left(\frac{X+s}{k}\right)\right] dU(s) \to \frac{1}{k\mu}\int_0^{\xi^{(k)}} \mathbb{E}\left[\tilde{\ell}\left(\frac{X+s}{k}\right)\right] ds$ follows from the argument in the proof of Lemma 5.8(a). Now consider the threshold function $\xi^{(k,c)}$ that satisfies

$$0 < \lim_{k\to\infty} \mathbb{E}\left[\tilde{\ell}\left(\frac{X+\xi^{(k,c)}}{k}\right)\right] < \lim_{k\to\infty} \frac{g_k^*}{\mu}.$$

In this case $0 < \lim_{k\to\infty} \xi^{(k,c)}/k \le \tilde{L}$, as $\tilde{L}(0) = 0$. Then,

$$\begin{aligned}\lim_{k\to\infty} \tilde{h}_k(g_k^*,\xi^{(k,c)}) &= \lim_{k\to\infty} \frac{1}{k\mu}\int_0^{\xi^{(k,c)}} \mathbb{E}\left[\tilde{\ell}\left(\frac{X+s}{k}\right)\right] ds - \lim_{k\to\infty}\frac{g_k^*}{\mu}\lim_{k\to\infty}\frac{\xi^{(k,c)}}{k} \\ &\le \lim_{k\to\infty}\frac{\xi^{(k,c)}}{k\mu}\left(\lim_{k\to\infty}\mathbb{E}\left[\tilde{\ell}\left(\frac{X+\xi^{(k,c)}}{k}\right)\right] - \lim_{k\to\infty} g_k^*\right) < 0,\end{aligned}$$

which contradicts the assumption and therefore completes the proof. □

5.H Proof of Theorem 5.4

The proofs of Parts (a) and (b) proceed analogously to the argument in Appendix 5.D, with the functions g, $\tilde{h}$, and $\tilde{\xi}_\lambda$ replaced by γ, h, and ξ_λ, respectively.

The monotonicity of Dh in x follows from the monotonicity of $F_X(x \mid \cdot)$ and $\ell(x)$. If F_X is continuous, the last part of Theorem 5.4 follows from the first-order condition. In case F_X is $d-$lattice, note that

$$Dh(\lambda,x,\theta) = h(\lambda, d\left\lfloor\frac{x}{d}\right\rfloor,\theta) - h(\lambda, d\left\lfloor\frac{x}{d}-1\right\rfloor,\theta).$$

Thus, the minimum of h can be achieved at the roots of $Dh(\lambda,x,\theta)$, or, when no root exists, at one of the boundary points. □

5.I Proof of Theorem 5.11

By Lemma 5.9 and 5.10, for any sufficiently small $\epsilon > 0$, there exists $k_{l_1}, k_{l_2} > 0$, such that,

$$-\frac{\epsilon}{2} < g_k\left(\xi_k^{\mathrm{O}}\right) - \gamma_k\left(\xi_k^{\mathrm{O}}\right) < \frac{\epsilon}{2}, \qquad \text{with} \quad \forall k > k_{l_1},$$

and

$$-\frac{\epsilon}{2} < \gamma_k\left(\xi_k^{\mathrm{A}}\right) - g_k\left(\xi_k^{\mathrm{A}}\right) < \frac{\epsilon}{2}, \qquad \text{with} \quad \forall k > k_{l_2}.$$

Letting $k_l = \max(k_{l_1}, k_{l_2})$, it follows that,

$$g_k\left(\xi_k^{\mathrm{A}}\right) - g_k(\xi_k^*) - \epsilon < \gamma_k\left(\xi_k^{\mathrm{A}}\right) - \gamma_k(\xi_k^*), \qquad k > k_l.$$

We notice that $g_k\left(\xi_k^{\mathrm{A}}\right) - g_k(\xi_k^*) \geq 0$ and $\gamma_k\left(\xi_k^{\mathrm{A}}\right) - \gamma_k(\xi_k^*) \leq 0$, by optimality. This implies that,

$$0 \leq g_k\left(\xi_k^{\mathrm{A}}\right) - g_k(\xi_k^*) < \epsilon, \qquad k > k_l,$$

which gives us Theorem 5.11(a). Part (b) of Theorem 5.11 follows directly from Lemmas 5.9 and 5.10. □

5.J Proof of Lemma 5.12

Part (a). By the weak consistency of $\hat{\theta}_\tau$, we have for every $\epsilon, \eta > 0$, there exists $\tau_{\epsilon,\eta}^l > 0$ such that for all $\tau > \tau_{\epsilon,\eta}^l$

$$\mathbb{P}\left\{\left\|\hat{\theta}(\mathbf{I}_\tau) - \theta\right\| > \epsilon\right\} < \eta. \tag{5.65}$$

Moreover, by (5.37) we have for all $\mathcal{T}^l > 0$

$$\mathbb{P}\left\{\exists k_{\mathcal{T}^l}(\omega) \text{ s.t. } \mathcal{T}_k(\theta) > \mathcal{T}^l \quad \forall k > k_{\mathcal{T}^l}(\omega)\right\} = 1, \tag{5.66}$$

where ω is a sample path. Note that $k_{\mathcal{T}^l}(\omega)$ is a random variable. Now, choose $\mathcal{T}^l > \tau_{\epsilon,\eta}^l$. Then, combining (5.65) and (5.66), we obtain that for all $k > k_{\mathcal{T}^l}(\omega)$

$$\mathbb{P}\left\{\left\|\hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right) - \theta\right\| > \epsilon \middle| \mathcal{T}_k\left(\theta\right)\right\} < \eta, \qquad \text{almost surely.}$$

Using law of total probability (condition on $\mathcal{T}_k(\theta)$) completes the proof for Part (a).

Part (b). By Assumption A.4, we have that for all $\lambda \in \mathbb{R}$, $\theta \in \Theta$, $\max \mathcal{D}_{\lambda,\theta} - \min \mathcal{D}_{\lambda,\theta} \leq x_{\text{flat}}$. Consequently, for all $\lambda \in \mathbb{R}$, $\theta \in \Theta$,

$$\limsup_{\theta' \to \theta} \xi_\lambda(\theta') - \liminf_{\theta' \to \theta} \xi_\lambda(\theta') \leq x_{\text{flat}}. \tag{5.67}$$

For $\epsilon > 0$, define $k_\epsilon = 2x_{\text{flat}}/\epsilon$. By (5.67), for all $\epsilon > 0$, there exists a neighborhood of θ, $\Theta_{\theta,\epsilon} := \{\theta' \in \Theta : \|\theta' - \theta\| < \epsilon/2\}$ such that for all $k > k_\epsilon$ and $\theta' \in \Theta_{\theta,\epsilon}$

$$\frac{1}{k}\left|\xi_k^{\mathrm{A}}(\theta') - \xi_k^{\mathrm{A}}(\theta)\right| \leq \frac{x_{\text{flat}}}{k} + \frac{\epsilon}{2k} < \frac{x_{\text{flat}}}{k_\epsilon} + \frac{\epsilon}{2} = \epsilon.$$

Choose $0 < \eta < 1$. Part (a) implies that there exists a $k_{\theta,\eta}$ such that $\mathbb{P}\left\{\hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right) \in \Theta_{\theta,\epsilon}\right\} > 1 - \eta$ for each $k > k_{\theta,\eta}$. Therefore for all $k > \max\{k_\epsilon, k_{\theta,\eta}\}$ we have

$$\begin{aligned}
&\mathbb{P}\left\{\frac{1}{k}\left|\xi_{k,\text{last}}^{\mathrm{P}}(\theta) - \xi_k^{\mathrm{A}}(\theta)\right| < \epsilon\right\} = \mathbb{P}\left\{\frac{1}{k}\left|\mathbb{E}\left[\xi_k^{\mathrm{A}}\left(\hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right)\right)\right] - \xi_k^{\mathrm{A}}(\theta)\right| < \epsilon\right\} \\
&\geq \mathbb{P}\left\{\frac{1}{k}\left|\mathbb{E}\left[\xi_k^{\mathrm{A}}\left(\hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right)\right)\right] - \xi_k^{\mathrm{A}}(\theta)\right| < \epsilon \middle| \hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right) \in \Theta_{\theta,\epsilon}\right\} \mathbb{P}\left\{\hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right) \in \Theta_{\theta,\epsilon}\right\} \\
&= \mathbb{P}\left\{\hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right) \in \Theta_{\theta,\epsilon}\right\} > 1 - \eta.
\end{aligned}$$

which establishes Part (b).

Part (c). Observe that for any $k > 0$ and $\theta \in \Theta$, $\frac{1}{k}\left|\xi^{\mathrm{P}}_{k,\mathrm{last}}(\theta) - \xi^{\mathrm{A}}_{k}(\theta)\right|$ is upper bounded by $\tilde{L}$, and consequently, uniformly integrable. Thus, by Part (b) it converges to zero in expectation, i.e., for all $\theta \in \Theta$,

$$\lim_{k\to\infty} \mathbb{E}\left[\frac{1}{k\mu(\theta)}\left|\xi^{\mathrm{P}}_{k,\mathrm{last}}(\theta) - \xi^{\mathrm{A}}_{k}(\theta)\right|\right] = 0.$$

We can extend this relation to

$$\lim_{k\to\infty} \mathbb{E}\left[\frac{1}{k\mu(\theta)}\left|\xi^{\mathrm{P}}_{k,\mathrm{last}}(\theta) + \mathbb{E}\left[Y\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\theta)\right)\right] - \left(\xi^{\mathrm{A}}_{k}(\theta) + \mathbb{E}\left[Y\left(\xi^{\mathrm{A}}(\theta)\right)\right]\right)\right|\right] = 0.$$

since $\mathbb{E}[Y(\cdot)]$ is bounded by Lorden's inequality. This implies that,

$$\begin{aligned}
&\lim_{k\to\infty} \mathbb{E}\left[\int_{\Theta} \frac{1}{k\mu(\mathbf{y})}\left|\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right)\right] - \left(\xi^{\mathrm{A}}_{k}(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi^{\mathrm{A}}(\mathbf{y})\right)\right]\right)\right| dF_\theta(\mathbf{y})\right] = \\
&\lim_{k\to\infty} \int_{\Theta} \mathbb{E}\left[\frac{1}{k\mu(\mathbf{y})}\left|\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right)\right] - \left(\xi^{\mathrm{A}}_{k}(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi^{\mathrm{A}}(\mathbf{y})\right)\right]\right)\right|\right] dF_\theta(\mathbf{y}) = 0.
\end{aligned} \tag{5.68}$$

The first equality in (5.68) follows from Fubini's theorem, which allows exchanging the order of integration and expectation. Next, for every $\epsilon > 0$ we have

$$\begin{aligned}
&\mathbb{P}\left\{\left|\int_{\Theta} \frac{1}{k\mu(\mathbf{y})}\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right)\right] - \left(\xi^{\mathrm{A}}_{k}(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi^{\mathrm{A}}(\mathbf{y})\right)\right]\right) dF_\theta(\mathbf{y})\right| > \epsilon\right\} \leq \\
&\mathbb{P}\left\{\int_{\Theta} \frac{1}{k\mu(\mathbf{y})}\left|\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right)\right] - \left(\xi^{\mathrm{A}}_{k}(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi^{\mathrm{A}}(\mathbf{y})\right)\right]\right)\right| dF_\theta(\mathbf{y}) > \epsilon\right\} \leq \\
&\frac{\mathbb{E}\left[\int_{\Theta} \frac{1}{k\mu(\mathbf{y})}\left|\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right)\right] - \left(\xi^{\mathrm{A}}_{k}(\mathbf{y}) + \mathbb{E}\left[Y\left(\xi^{\mathrm{A}}(\mathbf{y})\right)\right]\right)\right| dF_\theta(\mathbf{y})\right]}{\epsilon}.
\end{aligned}$$

The second inequality follows from Markov's inequality. Combining this with (5.68) completes the proof for Part (c).

Proof of Part (d). We prove the result for continuous F_X. The proof for $d-$lattice F_X is analogous, with integration replaced by summation. By continuity of F_X in $\theta \in \Theta$ and Part (a), for any $\epsilon\ , x > 0$

$$\lim_{k\to\infty} \mathbb{P}\left\{\left|\mathbb{E}\left[\ell(X_{1,1} + x, k) \mid \theta_1 = \hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right)\right] - \mathbb{E}\left[\ell(X_{1,1} + x, k) \mid \theta_1 = \theta\right]\right| < \epsilon\right\} = 1 \tag{5.69}$$

Now by the definition of r_k as in (5.34)

$$\begin{aligned}
&\frac{1}{k}\left|r_k\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\theta),\theta\right) - r_k(\xi^{\mathrm{A}}_k(\theta),\theta)\right| \\
&\leq \frac{1}{k\mu(\theta)}\int_0^{\xi^{\mathrm{A}}_k(\theta)}\left|\mathbb{E}\left[\ell(X_{1,1}+x,k)\mid\theta_1=\hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right)\right]-\mathbb{E}\left[\ell(X_{1,1}+x,k)\mid\theta_1=\theta\right]\right|dx+ \\
&\frac{1}{k\mu(\theta)}\left|\int_{\xi^{\mathrm{A}}_k(\theta)}^{\xi^{\mathrm{P}}_{k,\mathrm{last}}(\theta)}\max\left\{\mathbb{E}\left[\ell(X_{1,1}+x,k)\mid\theta_1=\hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right)\right],\mathbb{E}\left[\ell(X_{1,1}+x,k)\mid\theta_1=\theta\right]\right\}dx\right| \\
\leq&\frac{1}{k\mu(\theta)}\int_0^{\xi^{\mathrm{A}}_k(\theta)}\left|\mathbb{E}\left[\ell(X_{1,1}+x,k)\mid\theta_1=\hat{\theta}\left(\mathbf{I}_{\mathcal{T}_k(\theta)}\right)\right]-\mathbb{E}\left[\ell(X_{1,1}+x,k)\mid\theta_1=\theta\right]\right|dx+ \\
&\frac{1}{k\mu(\theta)}\ell(\tilde{L})\left|\xi^{\mathrm{P}}_{k,\mathrm{last}}(\theta)-\xi^{\mathrm{A}}_k(\theta)\right|.
\end{aligned}$$

Combining this result with Part (b) and (5.69) give us, for every $\epsilon>0$

$$\lim_{k\to\infty}\mathbb{P}\left\{\frac{1}{k}\left|r_k\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\theta),\theta\right)-r_k\left(\xi^{\mathrm{A}}_k(\theta),\theta\right)\right|<\epsilon\right\}=1. \tag{5.70}$$

The rest of the proof follows from an analogous argument as Part (c) with the result in Part (b) is replaced by (5.70), and by noticing that for all $k>0$ and $\theta\in\Theta$

$$\frac{1}{k}r_k\left(\xi^{\mathrm{A}}_k(\theta),\theta\right)\leq\ell(\tilde{L})\left(\frac{1}{k}+\frac{\tilde{L}}{\mu(\theta)}\right)<\infty.$$

Proof of Part (e). Let $\xi^{\mathrm{A}}_{k,0}$ denote the Oracle's asymptotically optimal threshold function when $\ell(x)=0$ for all $x\in\mathbb{R}_+$. Then for $k>0$,

$$\xi^{\mathrm{A}}_{k,0}(\theta)\geq\xi^{\mathrm{A}}_k(\theta), \tag{5.71}$$

since by assumption $\ell(x)\geq 0$. Furthermore, let $\{\xi^{\mathrm{L}}_k\}_{k>0}$ be a sequence of threshold functions that satisfy $\lim_{k\to\infty}\dfrac{\xi^{\mathrm{L}}_k(\theta)}{L(k)}=1$ and $\lim_{k\to\infty}\phi_k\left(\xi^{\mathrm{L}}_k(\theta)\right)=\infty$, for all $\theta\in\Theta$, when $\ell(x)=0$. This also implies $\lim_{k\to\infty}\dfrac{\phi_k\left(\xi^{\mathrm{L}}_k(\theta)\right)}{k}=0$. Hence

$$\lim_{k\to\infty}\gamma_k\left(\xi^{\mathrm{L}}_k\right)=\frac{\tilde{c}_p}{\tilde{L}\displaystyle\int_\Theta\frac{dF_\theta(\mathbf{y})}{\mu(\mathbf{y})}}. \tag{5.72}$$

We notice, from (5.33), that for any function $\xi^{(k)}\in\Xi_k$ a natural lower bound for the limit of $\gamma_k(\xi^{(k)})$ as $k\to\infty$ is as follows

$$\lim_{k\to\infty}\gamma_k\left(\xi^{(k)}(\theta)\right)\geq\frac{\tilde{c}_p}{\tilde{L}\displaystyle\int_\Theta\frac{dF_\theta(\mathbf{y})}{\mu(\mathbf{y})}}.$$

Combining this result with (5.71) and (5.72) gives $\lim_{k\to\infty}\left(\xi^{\mathrm{A}}_k(\theta)-\xi^{\mathrm{L}}_k(\theta)\right)\leq 0$. Consequently, for all $\theta\in\Theta$

$$\lim_{k\to\infty}\phi_k\left(\xi^{\mathrm{A}}_k(\theta)\right)=\infty,$$

which immediately implies that both sides of Part (e) converge to zero as $k\to\infty$, almost surely. $\square$

5.K Proof of Theorem 5.6

We notice that Lemma 5.12(c) can be extended to

$$\frac{1}{k}\int_\Theta \frac{1}{\mu(\mathbf{y})}\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})+\mathbb{E}\left[\tilde{Y}_1\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right)\middle|\theta_1=\mathbf{y},\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right]\right)dF_\theta(\mathbf{y}) \xrightarrow{p} \frac{1}{k}\int_\Theta \frac{1}{\mu(\mathbf{y})}\left(\xi^{\mathrm{A}}_k(\mathbf{y})+\mathbb{E}\left[\tilde{Y}_1\left(\xi^{\mathrm{A}}_k(\theta_1)\right)\middle|\theta_1=\mathbf{y}\right]\right)dF_\theta(\mathbf{y}). \quad (5.73)$$

by applying Lemma 5.8(d) to both sides of Lemma 5.12(c). Similarly, we use Parts (b) and (c) of Lemma 5.8 to extend Parts (c) and (e) of Lemma 5.12, respectively

$$\frac{1}{k}\int_\Theta \mathbb{E}\left[\sum_{\tau=1}^{T\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}),\theta_1\right)}\ell(S_{1,\tau},k)\middle|\theta_1=\mathbf{y},\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right]dF_\theta(\mathbf{y}) \xrightarrow{p} \frac{1}{k}\int_\Theta \mathbb{E}\left[\sum_{\tau=1}^{T\left(\xi^{\mathrm{A}}(\theta_1),\theta_1\right)}\ell(S_{1,\tau},k)\middle|\theta_1=\mathbf{y}\right]dF_\theta(\mathbf{y}), \quad (5.74)$$

$$\int_\Theta \mathbb{P}\left\{\tilde{Y}_1(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}))>\phi_k\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right)\middle|\theta_1=\mathbf{y},\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right\}dF_\theta(\mathbf{y}) \xrightarrow{p} \int_\Theta \mathbb{P}\left\{\tilde{Y}_1(\xi^{\mathrm{A}}(\theta_1))>\phi_k\left(\xi^{(k)}(\theta_1)\right)\middle|\theta_1=\mathbf{y}\right\}dF_\theta(\mathbf{y}). \quad (5.75)$$

Observe that the RHS of (5.73), (5.74), and (5.75) are uniformly integrable, which yields the stronger mean convergence results. This observation together with Fubini's theorem give us

$$\frac{1}{k}\int_\Theta \frac{1}{\mu(\mathbf{y})}\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})+\mathbb{E}\left[\tilde{Y}_1\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right)\middle|\theta_1=\mathbf{y}\right]\right)dF_\theta(\mathbf{y}) \to \frac{1}{k}\int_\Theta \frac{1}{\mu(\mathbf{y})}\left(\xi^{\mathrm{A}}_k(\mathbf{y})+\mathbb{E}\left[\tilde{Y}_1\left(\xi^{\mathrm{A}}_k(\theta_1)\right)\middle|\theta_1=\mathbf{y}\right]\right)dF_\theta(\mathbf{y}), \quad (5.76)$$

$$\frac{1}{k}\int_\Theta \mathbb{E}\left[\sum_{\tau=1}^{T\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}),\theta_1\right)}\ell(S_{1,\tau},k)\middle|\theta_1=\mathbf{y}\right]dF_\theta(\mathbf{y}) \to \frac{1}{k}\int_\Theta \mathbb{E}\left[\sum_{\tau=1}^{T\left(\xi^{\mathrm{A}}(\theta_1),\theta_1\right)}\ell(S_{1,\tau},k)\middle|\theta_1=\mathbf{y}\right]dF_\theta(\mathbf{y}), \quad (5.77)$$

$$\int_\Theta \mathbb{P}\left\{\tilde{Y}_1(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}))>\phi_k\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right)\middle|\theta_1=\mathbf{y}\right\}dF_\theta(\mathbf{y}) \to \int_\Theta \mathbb{P}\left\{\tilde{Y}_1(\xi^{\mathrm{A}}(\theta_1))>\phi_k\left(\xi^{(k)}(\theta_1)\right)\middle|\theta_1=\mathbf{y}\right\}dF_\theta(\mathbf{y}). \quad (5.78)$$

Next, by (5.44) as $k \to \infty$

$$
\begin{aligned}
&g_k(\xi^{\mathrm{P}}) = \\
&\frac{c_p + \displaystyle\int_\Theta \left(c_f \mathbb{P}\left\{ \tilde{Y}_1(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})) > \phi_k\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right) \middle| \theta_1 = \mathbf{y} \right\} + \frac{1}{k} \left[\sum_{\tau=1}^{T\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}),\theta_1\right)} \ell(S_{1,\tau}, k) \middle| \theta_1 = \mathbf{y} \right] \right) dF_\theta(\mathbf{y})}{\frac{1}{k}\displaystyle\int_\Theta \frac{1}{\mu(\mathbf{y})} \left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y}) + \mathbb{E}\left[\tilde{Y}_1\left(\xi^{\mathrm{P}}_{k,\mathrm{last}}(\mathbf{y})\right) \mid \theta_1 = \mathbf{y} \right] \right) dF_\theta(\mathbf{y})} \\
&\to \frac{c_p + \displaystyle\int_\Theta \left(c_f \mathbb{P}\left\{ \tilde{Y}_1(\xi^{\mathrm{A}}(\theta_1)) > \phi_k\left(\xi^{(k)}(\theta_1)\right) \middle| \theta_1 = \mathbf{y} \right\} + \frac{1}{k}\mathbb{E} \left[\sum_{\tau=1}^{T\left(\xi^{\mathrm{A}}(\theta_1),\theta_1\right)} \ell(S_{1,\tau}, k) \middle| \theta_1 = \mathbf{y} \right] \right) dF_\theta(\mathbf{y})}{\frac{1}{k}\displaystyle\int_\Theta \frac{1}{\mu(\mathbf{y})} \left(\xi^{\mathrm{A}}_k(\mathbf{y}) + \mathbb{E}\left[\tilde{Y}_1\left(\xi^{\mathrm{A}}_k(\theta_1)\right) \middle| \theta_1 = \mathbf{y} \right] \right) dF_\theta(\mathbf{y})} \\
&= g_k(\xi^{\mathrm{A}}) \to g_k(\xi^{\mathrm{O}}).
\end{aligned} \tag{5.79}
$$

The first equality in (5.79) follows from Kolmogorov's law of large numbers as per (5.44) together with the expression of g_k, i.e., Equation (5.32), the first convergence from Equations (5.77) through (5.78) and the last convergence from Theorem 5.11. □

Bibliography

K. Abnett. World's largest carbon market faces revamp under draft eu plan. *Reuters*, 2020. URL `https://www.reuters.com/article/us-climate-change-eu-target/worlds-largest-carbon-market-faces-revamp-under-draft-eu-plan,urldate={2021-05-12}`.

U. E. P. Agency. Emission facts: Average carbon dioxide emissions resulting from gasoline and diesel fuel (epa420-f-05-001), 2 2005.

G. Allon and J. Van Mieghem. Global dual sourcing: Tailored base-surge allocation to near-and offshore production. *Management Science*, 56(1):110–124, 2010.

E. Alvarez, M. Van der Heijden, and W. Zijm. The selective use of emergency shipments for service-contract differentiation. *International Journal of Production Economics*, 143(2): 518–526, 2013.

J. F. Andersen, A. R. Andersen, M. Kulahci, and B. F. Nielsen. A numerical study of markov decision process algorithms for multi-component replacement problems. *European Journal of Operational Research*, 299(3):898–909, 2022.

O. Andersen, C. S. Jensen, K. Torp, and B. Yang. Ecotour: Reducing the environmental footprint of vehicles using eco-routes. In *Proceedings - IEEE international conference on mobile data management*, pages 338–340, 2013.

A. Angelus and O. Özer. Knowledge you can act on: Optimal policies for assembly systems with expediting and advance demand information. *Operations Research*, 64(6):1338–1371, 2016.

J. Arts and G. Kiesmüller. Analysis of a two-echelon inventory system with two supply modes. *European Journal of Operational Research*, 225(2):263–272, 2013.

J. Arts, M. Van Vuuren, and G. Kiesmüller. Efficient optimization of the dual-index policy using markov chains. *IIE Transactions*, 43(8):604–620, 2011.

J. Arts, R. N. Boute, S. Loeys, and H. E. van Staden. Fifty years of maintenance optimization: Reflections and perspectives. *European Journal of Operational Research*, 2024.

E. Arıkan and W. Jammernegg. The single period inventory model under dual sourcing and product carbon footprint constraint. *International Journal of Production Economics*, 157: 15–23, 2014.

M. Asghari and S. M. J. M. Alehashem. Green vehicle routing problem: A state-of-the-art review. *International Journal of Production Economics*, 231, 1 2021.

S. Asmussen. A probabilistic look at the wiener-hopf equation. *SIAM Review*, 40:189–201, 1998.

S. Asmussen. *Applied Probability and Queues*. Springer, 2nd edition, 2003. ISBN 0-387-00211-1.

T. Aven and B. Bergman. Optimal replacement times: A general set-up. *Source: Journal of Applied Probability*, 23:432–442, 1986.

X. Bai, X. Chen, M. Li, and A. Stolyar. Asymptotic optimality of semi-open-loop policies in markov decision processes with large lead times. *Operations Research*, 71:2061–2077, 11 2023.

A. Barbosa-Póvoa, C. da Silva, and A. Carvalho. Opportunities and challenges in sustainable supply chain: An operations research perspective. *European Journal of Operational Research*, 268(2):399–431, 2018.

R. Barlow and L. Hunter. Optimum preventive maintenance policies. *Operations Research*, 8: 90–100, 2 1960.

M. Barth, T. Younglove, and G. Scora. Development of a heavy-duty diesel modal emissions and fuel consumption model. 2005.

M. Behnke and T. Kirschstein. The impact of path selection on ghg emissions in city logistics. *Transportation Research Part E: Logistics and Transportation Review*, 106:320–336, 10 2017.

T. Bektaş and G. Laporte. The pollution-routing problem. *Transportation Research Part B: Methodological*, 45:1232–1250, 2011.

P. Berling and V. Martínez de Albéniz. Dynamic speed optimization in supply chains with stochastic demand. *Transportation Science*, 50(3):1114–1127, 2016.

M. Bijvank and I. F. Vis. Lost-sales inventory theory: A review. *European Journal of Operational Research*, 215:1–13, 11 2011.

M. Bijvank, W. T. Huh, G. Janakiraman, and W. Kang. Robustness of order-up-to policies in lost-sales inventory systems. *Operations Research*, 62:1040–1047, 9 2014.

M. Bijvank, W. T. Huh, and G. Janakiraman. Lost-sales inventory systems. In J.-S. J. Song, editor, *Research Handbook on Inventory Management*. Edward Elgar Publishing, Inc., 2023.

P. Billingsley. *Probability and measure*. John Wiley & Sons, Inc., 3. ed edition, 1995. ISBN 0471007102.

D. Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102. University of California Press, 1951.

D. Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 26: 265–272, 6 1953.

K. Boriboonsomsin and M. Barth. Impacts of road grade on fuel consumption and carbon dioxide emissions evidenced by use of advanced navigation systems. *Transportation Research Record*, pages 21–30, 2009.

K. Boriboonsomsin, M. J. Barth, W. Zhu, and A. Vu. Eco-routing navigation system based on multisource historical and real-time traffic information. *IEEE Transactions on Intelligent Transportation Systems*, 13:1694–1704, 2012.

M. Brandenburg, K. Govindan, J. Sarkis, and S. Seuring. Quantitative models for sustainable supply chain management: Developments and directions. *European Journal of Operational Research*, 233(2):299–312, 2014.

C. Brunner, R. Giesen, M. A. Klapp, and L. Flórez-Calderón. Vehicle routing problem with steep roads. *Transportation Research Part A: Policy and Practice*, 151:1–17, 9 2021.

J. Bu, X. Gong, and D. Yao. Technical note-constant-order policies for lost-sales inventory models with random supply functions: Asymptotics and heuristic. *Operations Research*, 68: 1063–1073, 7 2020.

J. Bu, X. Gong, and X. Chao. Asymptotic scaling of optimal cost and asymptotic optimality of base-stock policy in several multidimensional inventory systems. *Operations Research*, 72: 1765–1774, 9 2024.

J. Bu, X. Gong, and H. Yin. Managing perishable inventory systems with positive lead times: Inventory position vs. projected inventory level. *SSRN*, 2025a.

J. Bu, H. Zhang, S. M. Ross, and S. Jasin. Asymptotic optimality of simple policies for stochastic inventory systems with delivery lead time and purchase returns. *SSRN*, 2025b.

C. Bérenguer, A. Grall, L. Dieulle, and M. Roussignol. Maintenance policy for a continuously monitored deteriorating system. *Probability in the Engineering and Informational Sciences*, 17:235–250, 2003. ISSN 02699648. doi: 10.1017/S0269964803172063.

B. Chen, X. Chao, and C. Shi. Nonparametric learning algorithms for joint pricing and inventory control with lost sales and censored demand. *Mathematics of Operations Research*, 46:726–756, 5 2021.

N. Chen, Z.-S. Ye, Y. Xiang, and L. Zhang. Condition-based maintenance using the inverse gaussian degradation model. *European Journal of Operational Research*, 243(1):190–199, 2015.

X. Chen and X. Wang. Effects of carbon emission reduction policies on transportation mode selections with stochastic demand. *Transportation Research Part E: Logistics and Transportation Review*, 90:196–205, 2016.

E. Commission. State of the union: Q & a on the 2030 climate target plan, 2020.

S. Dabia, E. Demir, and T. Van Woensel. An exact approach for a variant of the pollution-routing problem. *Transportation Science*, 51:607–628, 2017.

G. Dantzig and P. Wolfe. Decomposition principle for linear programs. *Operations Research*, 8 (1):101–111, 1960.

U. N. C. Database. `https://comtrade.un.org`, 2020.

B. de Jonge and P. A. Scarf. A review on maintenance optimization. *European Journal of Operational Research*, 285(3):805–824, 2020.

R. Dekker, J. Bloemhof, and I. Mallidis. Operations research for green logistics – an overview of aspects, issues, contributions and challenges. *European Journal of Operational Research*, 219(3):671–679, 2012.

E. Demir, T. Bektaş, and G. Laporte. A comparative analysis of several vehicle emission models for road freight transportation. *Transportation Research Part D: Transport and Environment*, 16:347–357, 2011.

E. Demir, T. Bektaş, and G. Laporte. An adaptive large neighborhood search heuristic for the pollution-routing problem. *European Journal of Operational Research*, 223:346–359, 12 2012.

E. Demir, T. Bektaş, and G. Laporte. A review of recent research on green road freight transportation. *European Journal of Operational Research*, 237:775–793, 9 2014.

C. Derman. *On Optimal Replacement Rules When Changes of State Are Markovian*, pages 201–210. The Rand Corporation, 1963a.

C. Derman. Optimal replacement and maintenance under markovian deterioration with probability bounds on failure. *Management Science*, 9:478–481, 4 1963b.

E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.

C. Dong, S. Transchel, and K. Hoberg. An inventory control model for modal split transport: A tailored base-surge approach. *European Journal of Operational Research*, 264(1):89–105, 2018.

Q. Dong, C.and Li, B. Shen, and X. Tong. Sustainability in supply chains with behavioral concerns. *Sustainability*, 11(15), 2019.

C. Drent, M. Drent, J. Arts, and S. Kapodistria. Real-time integrated learning and decision making for cumulative shock degradation. *Manufacturing and Service Operations Management*, 25:235–253, 1 2023a.

M. Drent and J. Arts. Expediting in two-echelon spare parts inventory systems. *Manufacturing & Service Operations Management*, 23(6):1431–1448, 2021.

M. Drent and J. Arts. Effective dual-sourcing through inventory projection. *arXiv preprint arXiv:2207.12182*, 2022.

M. Drent, P. Moradi, and J. Arts. Efficient emission reduction through dynamic supply mode selection. *European Journal of Operational Research*, 311:925–941, 12 2023b.

H. Dündar, M. Soysal, M. Ömürgönülşen, and A. Kanellopoulos. A green dynamic tsp with detailed road gradient dependent fuel consumption estimation. *Computers and Industrial Engineering*, 168, 6 2022.

J. F. Ehmke, A. M. Campbell, and B. W. Thomas. Data-driven approaches for emissions-minimized paths in urban areas. *Computers and Operations Research*, 67:34–47, 3 2016a.

J. F. Ehmke, A. M. Campbell, and B. W. Thomas. Vehicle routing to minimize time-dependent emissions in urban areas. *European Journal of Operational Research*, 251:478–494, 6 2016b.

J. F. Ehmke, A. M. Campbell, and B. W. Thomas. Optimizing for total costs in vehicle routing in urban areas. *Transportation Research Part E: Logistics and Transportation Review*, 116: 242–265, 8 2018.

A. H. Elwany, N. Z. Gebraeel, and L. M. Maillart. Structured replacement policies for components with complex degradation processes and dedicated sensors. *Operations Research*, 59: 684–695, 5 2011.

E. Engebrethsen and S. Dauzére-Pérés. Transportation mode selection in inventory models: A literature review. *European Journal of Operational Research*, 279(1):1–25, 2019.

E. Ericsson, H. Larsson, and K. Brundell-Freij. Optimizing route choice for lowest fuel consumption - potential effects of a new driver support tool. *Transportation Research Part C: Emerging Technologies*, 14:369–383, 2006.

European Environment Agency. *Trends and projections in Europe 2020: Tracking progress towards Europe's climate and energy targets*. 2020.

European Environment Agency. Trends and projections in europe 2021. 2021. doi: 10.2800/80374. URL `http://www.eea.europa.eu,`.

R. M. Feldman. Optimal replacement with semi-markov shock models. *Journal of Applied Probability*, 13:108–117, 1976.

W. Feller. *An introduction to probability theory and its applications, Volume 2*. Wiley Series in Probability and Statistics. Wiley, 1991. ISBN 9780471257097.

Q. Feng, S. Sethi, H. Yan, and H. Zhang. Are base-stock policies optimal in inventory problems with multiple delivery modes? *Operations Research*, 54(4):801–807, 2006.

R. W. Fox, A. T. McDonald, and J. W. Mitchell. *Fox and McDonald's introduction to fluid mechanics*. John Wiley & Sons, 2020.

A. Franceschetti, D. Honhon, T. Van Woensel, T. Bektaş, and G. Laporte. The time-dependent pollution-routing problem. *Transportation Research Part B: Methodological*, 56:265–293, 2013.

Y. Fukuda. Optimal policies for the inventory problem with negotiable leadtime. *Management Science*, 10(4):690–708, 1964.

GE HealthCare. GE Healthcare introduces insite onwatch to its assurepoint services, helping enable greater efficiencies in nuclear medicine and PET/CT, 2010. URL `https://www.gehealthcare.com/about/newsroom/press-releases/ge-healthcare-introduces-insite-onwatch-its-assurepoint-services-helping-enable?npclid=botnpclid&srsltid=AfmBOorYTvfqrbujFtvj6ZvQSAjEg0SP_cjYIFecwVj3l-JcGzcRvTLT`.

J. K. Ghosh, M. Delampady, and T. Samanta. *An Introduction to Bayesian Analysis: Theory and Methods*. Springer Texts in Statistics. Springer New York, 2007. ISBN 9780387400846.

M. A. Girshick and H. Rubin. Institute of mathematical statistics a bayes approach to a quality control model. *Annals of Mathematical Statistics*, 23:114–125, 3 1952.

D. A. Goldberg, D. A. Katz-Rogozhnikov, Y. Lu, M. Sharma, and M. S. Squillante. Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. *Mathematics of Operations Research*, 41:898–913, 8 2016.

D. A. Goldberg, M. I. Reiman, and Q. Wang. A survey of recent progress in the asymptotic analysis of inventory systems. *Production and Operations Management*, 30:1718–1750, 6 2021.

S. Goldfarb and S. Patterson. Why Are Gasoline Prices So High? Ukraine-Russia War Sparks Increases Across U.S. *The Wall Street Journal*, 4 2022. URL `https://www.wsj.com/articles/why-gas-prices-expensive-11646767172`.

M. Goodier and D. Campbell. Thousands harmed and 87 dead after nhs equipment failures in england. *The Guardian*, 6 2025. URL `https://www.theguardian.com/society/2025/jun/09/thousands-harmed-and-87-dead-after-nhs-equipment-failures-in-england`.

A. Gut. *Stopped Random Walks: Limit Theorems and Applications*. Springer New York, 2009a. ISBN 978-0-387-87834-8.

A. Gut. *An Intermediate Course in Probability*. Springer, second edition, 2009b.

J. Hansen and S. Lebedeff. Global trends of measured surface air temperature. *Journal of Geophysical Research*, 92:13345–13358, 1987.

C. Haubitz and U. Thonemann. How to change a running system—controlling the transition to optimized spare parts inventory policies. *Production and Operations Management*, 30(5): 1386–1405, 2021.

K. Hoen, T. Tan, J. Fransoo, and G. Van Houtum. Effect of carbon emission regulations on transport mode selection under stochastic demand. *Flexible Service Manufacturing Journal*, 26:170–195, 2014a.

K. Hoen, T. Tan, J. Fransoo, and G. Van Houtum. Switching transport modes to meet voluntary carbon emission targets. *Transportation Science*, 48(4):592–608, 2014b.

Y. Huang, L. Zhao, T. Van Woensel, and J. P. Gross. Time-dependent vehicle routing problem with path flexibility. *Transportation Research Part B: Methodological*, 95:169–195, 1 2017.

W. T. Huh, G. Janakiraman, J. A. Muckstadt, and P. Rusmevichientong. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Science*, 55(3):404–420, 2009.

W. T. Huh, G. Janakiraman, and M. Nagarajan. Average cost single-stage inventory models: An analysis using a vanishing discount approach. *Operations Research*, 59:143–155, 1 2011.

Intergovernmental Panel on Climate Change. Climate change 2021: The physical science basis. contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change, 2021.

International Energy Agency. Material efficiency in clean energy transitions, 2019. URL `https://www.iea.org/reports/material-efficiency-in-clean-energy-transitions`.

International Maritime Organization. Resolution MEPC.203(62) amendments to the annex of the protocol of 1997 to amend the international convention for the prevention of pollution from ships, 1973, as modified by the protocol of 1978 relating thereto (inclusion of regulations on energy efficiency for ships in MARPOL annex vi). 2011.

S. Johansen and A. Thorstenson. Pure and restricted base-stock policies for the lost-sales inventory system with periodic review and constant lead times. In *15th International Symposium on Inventories*, 2008. Conference date: 22-08-2008 Through 26-08-2008.

H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack problems.* Springer, Berlin, 2004.

A. Khaleghei and M. J. Kim. Optimal control of partially observable semi-markovian failing systems: An analysis using a phase methodology. *Operations Research*, 69:1282–1304, 7 2021.

M. Kim. Robust control of partially observable failing systems. *Operations Research*, 64(4): 999–1014, 2016.

M. J. Kim and V. Makis. Joint optimization of sampling and control of partially observable failing systems. *Operations Research*, 61:777–790, 5 2013.

J. F. Kingman. The single server queue in heavy traffic. *Mathematical Proceedings of the Cambridge Philosophical Society*, 57:902–904, 1961.

C. Koc, T. Bektaş, O. Jabali, and G. Laporte. The fleet size and mix pollution-routing problem. *Transportation Research Part B: Methodological*, 70:239–254, 12 2014.

P. Kolesar. Minimum cost replacement under Markovian deterioration. *Management Science*, 12:694–706, 5 1966.

D. Konur, J. Campbell, and S. Monfared. Economic and environmental considerations in a stochastic inventory control model with order splitting under different delivery schedules among suppliers. *Omega*, 71:46–65, 2017.

D. Lai, Y. Costa, E. Demir, A. M. Florio, and T. Van Woensel. The pollution-routing problem with speed optimization and uneven topography. *Computers & Operations Research*, 164: 106557, 2024.

J. Larminie and J. Lowry. *Electric vehicle technology explained, Second Edition.* John Wiley & Sons, 2012. ISBN 9781119942733.

D. Lawson and E. Porteus. Multistage inventory management with expediting. *Operations Research*, 48(6):878–893, 2000.

N. Lemmens, J. Gijsbrechts, and R. Boute. Synchromodality in the physical internet – dual sourcing and real-time switching between transport modes. *European Transport Research Review*, 1(10):518–526, 2019.

G. Lorden. On excess over the boundary. *The Annals of Mathematical Statistics*, 41:520–527, 1970.

M. Lübbecke and J. Desrosiers. Selected topics in column generation. *Operations Research*, 53 (6):1007–1023, 2005.

C. Lyu, H. Zhang, and L. Xin. Ucb-type learning algorithms with kaplan–meier estimator for lost-sales inventory models with lead times. *Operations Research*, 72:1317–1332, 7 2024.

L. M. Maillart. Maintenance policies for systems with condition monitoring and obvious failures. *IIE Transactions (Institute of Industrial Engineers)*, 38:463–475, 6 2006.

L. M. Maillart and L. Zheltova. Structured maintenance policies on interior sample paths. *Naval Research Logistics*, 54:645–655, 9 2007.

R. Moghdani, K. Salimifard, E. Demir, and A. Benyettou. The green vehicle routing problem: A systematic literature review. *Journal of Cleaner Production*, 279, 1 2021.

P. Moradi, J. Arts, and J. C. Velázquez-Martínez. Load asymptotics and dynamic speed optimization for the greenest path problem: a comprehensive analysis. *OR Spectrum*, 47:477–524, 2024.

R. Nelsen. *An Introduction to Copulas*. Springer, second edition, 2006.

New York City Department of Transportation. Off-hour deliveries. `https://ohdnyc.com`, 2024. Accessed: 25 July 2024.

NTM. Methods and manuals. `https://www.transportmeasures.org/en/wiki/manuals/`, 2015.

NTM. NTMCalc Basic 4.0. `https://www.transportmeasures.org/ntmcalc/v4/basic/index.html#/`, n.d. Accessed: 2021-7-20.

OpenStreetMap contributors. Planet dump retrieved from https://planet.osm.org . `https://www.openstreetmap.org`, 2017.

OpenStreetMap contributors. Osmnx user reference, 2022. URL `https://osmnx.readthedocs.io/en/stable/osmnx.html`.

G. Palak, S. Ekşioğlu, and J. Geunes. Analyzing the impacts of carbon regulatory mechanisms on supplier and mode selection decisions: An application to a biofuel supply chain. *International Journal of Production Economics*, 154:198–216, 2014.

D. Pamučar, L. Gigović, G. Ćirović, and M. Regodić. Transport spatial model for the definition of green routes for city logistics centers. *Environmental Impact Assessment Review*, 56:72–87, 1 2016.

J. Poppe, R. N. Boute, and M. R. Lambrecht. A hybrid condition-based maintenance policy for continuously monitored components with two degradation thresholds. *European Journal of Operational Research*, 268:515–532, 7 2018.

Radiology Business. Siemens debuts new equipment service program, 2010. URL `https://radiologybusiness.com/topics/medical-imaging/molecular-imaging/siemens-debuts-new-equipment-service-program?utm_source=chatgpt.com`.

R. Raeesi and K. G. Zografos. The multi-objective steiner pollution-routing problem on congested urban road networks. *Transportation Research Part B: Methodological*, 122:457–485, 4 2019.

W. Rao, F. Liu, and S. Wang. An efficient two-objective hybrid local search algorithm for solving the fuel consumption vehicle routing problem. *Applied Computational Intelligence and Soft Computing*, 2016:1–16, 2016.

H. Rosič and W. Jammernegg. The economic and environmental performance of dual sourcing: A newsvendor approach. *International Journal of Production Economics*, 143(1):109–119, 2013.

S. M. Ross. A Markovian replacement model with a generalization to include stocking. *Management Science*, 15:702–715, 1969.

S. M. Ross. Quality control under markovian deterioration. *Management Science*, 17:587–596, 1971.

M. Schäl. Average optimality in dynamic programming with general state space. *Mathematics of Operations Research*, 18:163–172, 1993.

M. Schröder and P. Cabral. Eco-friendly 3d-routing: A gis based 3d-routing-model to estimate and reduce co2-emissions of distribution transports. *Computers, Environment and Urban Systems*, 73:40–55, 1 2019.

G. Scora and M. Barth. Comprehensive modal emissions model (cmem), version 3.01: User's guide. `https://www.cert.ucr.edu/sites/default/files/2019-07/CMEM_User_Guide_v3.01d.pdf`, 6 2006.

G. Scora, K. Boriboonsomsin, and M. Barth. Value of eco-friendly route choice for heavy-duty trucks. *Research in Transportation Economics*, 52:3–14, 10 2015.

A. Sheopuri, G. Janakiraman, and S. Seshadri. New policies for the stochastic inventory control problem with two supply sources. *Operations Research*, 58(3):734–745, 2010.

Siemens. The true cost of downtime 2024, 2024. URL `https://assets.new.siemens.com/siemens/assets/api/uuid:1b43afb5-2d07-47f7-9eb7-893fe7d0bc59/TCOD-2024_original.pdf`.

J. Sun and J. Van Mieghem. Robust dual sourcing inventory management: Optimality of capped dual index policies and smoothing. *Manufacturing & Service Operations Management*, 21(4): 912–931, 2019.

J. Svoboda, S. Minner, and M. Yao. Typology and literature review on multiple supplier inventory control models. *European Journal of Operational Research*, 293(1):1–23, 2021.

D. Thomas and J. Tyworth. Pooling lead-time risk by order splitting: A critical review. *Transportation Research Part E: Logistics and Transportation Review*, 42(4):245–257, 2006.

D. S. Thomas. Cost-effective environmental sustainability: A focus on the circular economy. Technical report, National Institute of Standards and Technology, 2023. URL `https://www.nist.gov/el/applied-economics-office/manufacturing/circular-economy/product-longevity`.

H. C. Tijms. *A First Course in Stochastic Models*. John Wiley & Sons Ltd, 2003. ISBN 0-471-49881-5.

UN Environment Programme International Resource Panel. Re-defining value - the manufacturing revolution: Remanufacturing, refurbishment, repair and direct reuse in the circular economy, 2018. URL `https://www.resourcepanel.org/reports/re-defining-value-manufacturing-revolution`.

United Nations Conference on Trade and Development. Review of maritime transport. 2020.

United Nations Environment Programme. *Emissions Gap Report 2020*. UNEP DTU Partnership, 2020.

US Environmental Protection Agency. Inventory of u.s. greenhouse gas emissions and sinks: 1990-2020 (epa 430-r-22-003), 2022.

U.S. Geological Survey. Digital Elevation Model - SRTM 1 Arc-Second 30m (NASA, NGA). `https://lta.cr.usgs.gov/citation`, 2000.

W. van Jaarsveld and J. Arts. Projected inventory-level policies for lost sales inventory systems: Asymptotic optimality in two regimes. *Operations Research*, 72(5):1790–1805, 2024.

C. van Oosterom, H. Peng, and G. J. V. Houtum. Maintenance optimization for a markovian deteriorating system with population heterogeneity. *IISE Transactions*, 49:96–109, 2017.

H. E. van Staden and R. N. Boute. The effect of multi-sensor data on condition-based maintenance policies. *European Journal of Operational Research*, 290:585–600, 4 2021.

S. Veeraraghavan and A. Scheller-Wolf. Now or later: A simple policy for effective dual sourcing in capacitated systems. *Operations Research*, 56(4):850–864, 2008.

M. Vieweg, D. Bongardt, C. Hochfeld, A. Jung, E. Scherer, R. Adib, and F. Guerra. *Towards Decarbonising Transport – A 2018 Stocktake on Sectoral Ambition in the G20*. 2018.

C. Waltho, S. Elhedhli, and F. Gzara. Green supply chain network design: A review focused on policy adoption and emission quantification. *International Journal of Production Economics*, 208:305–318, 2 2019.

A. Whittemore and S. Saunders. Optimal inventory under stochastic demand with two supply options. *SIAM Journal on Applied Mathematics*, 32(2):293–305, 1977.

World Bank Group. World development indicators, 2025. URL `https://data.worldbank.org`.

Y. Xiao, X. Zuo, J. Huang, A. Konak, and Y. Xu. The continuous pollution routing problem. *Applied Mathematics and Computation*, 387, 12 2020.

L. Xin. Technical note—understanding the performance of capped base-stock policies in lost-sales inventory models. *Operations Research*, 69:61–70, 1 2021.

L. Xin. 1.79-approximation algorithms for continuous review single-sourcing lost-sales and dual-sourcing inventory models. *Operations Research*, 70:111–128, 1 2022.

L. Xin and D. A. Goldberg. Optimality gap of constant-order policies decays exponentially in the lead time for lost sales models. *Operations Research*, 64:1556–1565, 11 2016.

H. Zhang and W. Zhang. Analytical solution to a partially observable machine maintenance problem with obvious failures. *Management Science*, 69:3993–4015, 7 2023.

L. Zhang, Y. Lei, and H. Shen. How heterogeneity influences condition-based maintenance for gamma degradation process. *International Journal of Production Research*, 54(19):5829–5841, 2016.

N. Zhang, M. Fouladirad, A. Barros, and J. Zhang. Condition-based maintenance for a k-out-of-n deteriorating system under periodic inspection with failure dependence. *European Journal of Operational Research*, 287(1):159–167, 2020.

M. Zhou, H. Jin, and W. Wang. A review of vehicle fuel consumption models to evaluate eco-driving and eco-routing. *Transportation Research Part D: Transport and Environment*, 49: 203–218, 12 2016.