Indian School of Business

Essays on Freight Traffic Management in Indian Railways

by

Himanshu Arha

A dissertation submitted to the

Operations Management Department

Indian School of Business

September 2025

Committee Members

Name	**Signature**
Milind Sohoni (Co-chair)	_____
Sumit Kunnumkal (Co-chair)	_____
Vishwakant Malladi (Member)	_____
Raja Gopalakrishnan (Member)	_____

Abstract

Indian Railways (IR) is the world's fourth-largest railway network. In 2023–24 alone, IR transported more than 7 billion passengers and over 1,400 million tonnes of freight. Freight operations contribute to the majority ($\sim$70%) of the IR's revenue. Despite being the most cost-effective and energy-efficient mode of transporting freight, the share of IR in the freight transportation market has been steadily declining. One of the primary reasons for this decline is the low speed of freight trains, which has remained constant at around 25 kmph for the last five decades. Unlike passenger trains, freight trains typically operate without fixed schedules, with their movement decisions made in real time by traffic controllers (also known as *section controllers*). Thus, to better understand the causes of low speed, it is important to study the decision-making process that governs freight train movements.

In the first essay, we empirically study how section controllers make decisions for freight trains. Using a detailed dataset from the Deen Dayal Upadhyay division of IR, we construct minute-by-minute snapshots of network congestion caused by freight and passenger train movements. Utilizing this, we estimate two discrete choice models corresponding to the section controllers' hold and stop decisions. We explicitly account for spatial dependence across these decisions by using a copula-based approach and address potential endogeneity concerns using network instruments. Our results show that passenger trains are given high priority, which significantly affects the hold and stop decisions for freight trains. Even when the subsequent track is free, freight trains are held in the anticipation of an upcoming passenger train (*strategic idling effect*). Additionally, the presence of trailing freight trains increases the probability of the lead freight train being released (*push effect*). Further, we find that section controllers are more likely to hold freight trains at high-capacity stations, such as junctions.

In the second essay, we use estimates from the fitted models to conduct a detailed simulation-based counterfactual analysis aimed at improving freight train speeds. Specifically, we study two sets of counterfactuals: capacity-based and non-capacity-based. In the capacity-based counter-

factual, we simulate Freight Only Corridor (FOC), a dedicated corridor reserved exclusively for freight trains. We find that FOC leads to a 37% reduction in average dwell time, and a 31% reduction in stoppages, resulting in a 16% increase in average speed. Given the high cost of capacity interventions, we also examine two operational, non-capacity-based counterfactuals: threshold-based releases and capacity consolidation. As cost-effective alternatives, we find that a 60-minute threshold-based release policy increases average speed by 13%, while a moderate capacity consolidation of 30% results in a 13% improvement in average speed. Together, these findings highlight practical and effective interventions for improving freight train speeds.

In the third essay, motivated by station-level freight queues in IR, we analyze a finite-buffer, two-class priority-queue model with class-specific deterministic service times. The objective is to characterize optimal service-control policies for non-priority customers with the aim of minimizing the waiting costs. We analyze two cases: visibility and no-visibility, where visibility refers to the information of the time and type of the next customer arrival. For the no-visibility case, we formulate the problem under both long-run average cost and infinite-horizon discounted cost objectives, and derive structural results. Under the long-run average cost objective, we find that the optimal policy is to always serve the non-priority customer. For the infinite-horizon discounted cost objective, we identify conditions under which a threshold policy, based on the number of non-priority customers waiting in the queue, is optimal. In the visibility case, our simulation study shows that the optimal policy follows a monotone switching curve, defined by the number of non-priority customers waiting in the queue and the expected time to the next priority arrival. This curve characterizes the threshold at which the policy shifts from holding to serving a non-priority customer. We further empirically validate the results using data from IR. We find support for a threshold policy under which freight trains are held, in anticipation of passenger train arrival, up to a threshold. We hope these results can serve as practical guidelines for managing freight train movements while minimizing delays and maintaining passenger train priorities.

Acknowledgments

This dissertation would not have been possible without the support of numerous individuals and institutions. First and foremost, I am deeply grateful to Professor Milind Sohoni for his tremendous guidance throughout this journey. His unwavering mentorship shaped every stage of this dissertation, and I remain deeply indebted to him. I am thankful to Professor Kashish Arora, Professor Vishwakant Malladi, and Professor Sumit Kunnumkal for their support and advice in nurturing my research skills. Their constructive feedback and push for excellence have been instrumental in shaping this dissertation. Additionally, I would like to thank Dr. Raja Gopalakrishnan for his motivation and for sharing valuable insights into Indian Railways' operations. I would like to thank the Indian School of Business, particularly the FPM office, for providing the financial and administrative support necessary for my research and travel to conferences. I am also grateful to the Indian Railways for providing the necessary data. Above all, I am thankful for the unwavering support of my family, friends, and colleagues.

Contents

Abstract ii

Acknowledgments iii

1 Structural Estimation of Freight Delays in the Indian Railways Network **1**

1.1 Introduction 1

1.2 Literature Review 4

1.3 Data and Empirical Context 8

1.3.1 Indian Railways 8

1.3.2 Section Controllers' Role in Freight Movement 8

1.3.3 Data 9

1.3.4 Model-free Evidence 12

1.4 Model 14

1.4.1 Decision 1 (Hold/Release) 16

1.4.2 Decision 2 (Stop/Pass-through) 17

1.4.3 Modeling Dependency in Decisions 18

1.5 Estimation 22

1.5.1 Endogeneity 22

1.5.2 Network Instrumental Variables 23

1.5.3 Control Function Approach to Estimation 26

1.6 Estimation Results and Model Validation 28

1.6.1 Estimation Results . . . 28
1.6.2 Model Validation . . . 31
1.7 Conclusion . . . 32
1.8 Appendix . . . 34
1.8.1 Data and Model Appendix . . . 34
1.8.2 Estimation Appendix . . . 37
1.8.3 Results Appendix . . . 40

2 Improving Freight Train Speed in Indian Railways: Counterfactual Analysis 48
2.1 Introduction . . . 48
2.2 Capacity Counterfactual: Freight-Only Corridors . . . 51
2.2.1 Partial Sharing of FOC Track . . . 58
2.2.2 DFC-like Network Enhancements . . . 59
2.3 Threshold Based Releases . . . 60
2.3.1 Potential Interference to Passenger Trains . . . 62
2.4 Capacity Consolidation: Vertically Stacked Trains . . . 63
2.4.1 Consolidation Losses . . . 67
2.5 Costs, Benefits and Feasibility Analysis . . . 68
2.6 Conclusion . . . 71
2.7 Appendix . . . 72
2.7.1 Simulator Schema . . . 72
2.7.2 Results Appendix . . . 73

3 Dynamic Control of Priority Queues 76
3.1 Introduction . . . 76
3.2 Literature Review . . . 79
3.3 Queue Control Problem . . . 82

3.4 No Visibility Case 83
3.4.1 Long Run Average Cost Formulation 84
3.4.2 Infinite Horizon Discounted Cost Formulation 96
3.4.3 Discussion 111
3.5 Visibility Case 112
3.6 Numerical Analysis 116
3.7 Empirical Evidence 119
3.8 Conclusion 123
3.9 Appendix 124
3.9.1 Sub-cycle Analysis for Discounted Cost Formulation 124

Bibliography **128**

Chapter 1

Structural Estimation of Freight Delays in the Indian Railways Network

1.1 Introduction

Railways play a crucial role in freight transportation, a market that is expected to more than double by the year 2050.[1] Rail transportation stands out for its cost-effectiveness, particularly for transporting large volumes of freight over long distances. In the United States, for example, the cost of freight transportation per ton-mile by railways is estimated at $0.04, compared to $0.12 by truck and $1.36 by air shipping.[2] In addition to their cost-effectiveness, railways are one of the most sustainable modes of transportation. They account for 7% of freight and 8% of passenger traffic globally, yet represent just 2% of total transport energy demand.[3] Despite their cost efficiency and sustainable advantage, railways worldwide have been losing market share in inland freight transportation. From 2007 to 2021, the market share of railways in inland freight transportation globally has dropped from 54% to 43%.[4]

A major challenge for railway freight globally is the slow speed of freight trains. The average

[1] https://www.itf-oecd.org/worldwide-transport-activity-double-emissions-rise-further
[2] https://assets.publishing.service.gov.uk/media/5c73fc7340f0b603d87fe977/automation_in_freight.pdf
[3] https://www.iea.org/reports/the-future-of-rail
[4] https://uic.org/com/IMG/pdf/the_modal_share_of_rail_in_inland_transport_and_infrastructure_investment.pdf

speed of freight trains in many parts of the world, like in the central and eastern European countries and India, has remained low (between 20 and 30 kmph).[5] One of the reasons for the slow speed is the heavier weight of freight, which makes these trains harder to accelerate and decelerate. Additionally, in many regions, including Europe and Asia, freight trains share railway infrastructure with passenger trains and often get de-prioritized over passenger traffic, causing freight delays. Not only do these delays (and consequently slower speeds) make freight transportation less attractive to the customers, they also lead to a substantial monetary loss for the railways. Lovett et al. [2015] quantify the cost of these delays, ranging from $834 for bulk freight trains to $2214 for intermodal freight trains per hour per train.[6]

In India, a similar pattern of declining railway freight market share emerges, but with unique challenges and dynamics. The Indian railways are the fourth-largest globally, and freight delays incur proportionally larger monetary costs, leading to inflated logistics costs in the country. Specifically, India's logistics costs are estimated at 14% of the GDP, among the highest globally. These costs are unsustainable given that the Indian logistics market is one of the fastest-growing in the world, projected to reach $484 billion by 2029. Given India's immense scale and the importance of railways in the Indian logistics industry, freight delays have far-reaching implications for the Indian economy and transportation network.

The unique aspect of Indian railways is its operational model – unlike passenger trains, which adhere to fixed schedules, freight trains lack a set timetable, and their movements are subject to manual control by traffic managers (also known as *section controllers*). This is in contrast to the U.S., where, despite regulations favoring passenger trains, freight train operators often prioritize freight trains, resulting in faster average speeds for freight trains compared to Europe and India.[7] The priority given to passenger trains adds to the complexity of the problem and often leads to extended delays. It is noteworthy that despite generating over 70% of its 30 billion USD annual

[5] https://www.eca.europa.eu/lists/ecadocuments/sr16_08/sr_rail_freight_en.pdf
[6] These costs are broken down into crew, locomotives, fuel, railcars, and lading costs.
[7] https://specialty-freight.com/2021/06/us-requires-railroads-to-prioritize-passengers-over-freight/

revenue from freight transportation, the average freight train speed in the Indian Railways is only around 25 kmph and has remained constant for the past few decades.[8]

Section controllers are pivotal in managing freight traffic. They make real-time operational decisions for efficient freight movement across the network. Specifically, they determine when and how long a freight train halts at a station and thus exert significant influence over freight trains' movement. The stop and hold decisions made by the section controllers ultimately determine the overall freight delays and freight train speeds in the network. The goal of this chapter is to understand how railway controllers make stop and hold decisions for freight trains. Our work provides one of the first empirical analyses of freight train flows by studying the railway controller's decision-making. Even though we study the Indian setting, similar control mechanisms are used by railway organizations across the globe for efficiently handling train dispatches. Hence, the methodology can generally be applied to study freight delays in other railway networks worldwide.

To study the freight delays, we construct a detailed dataset that captures high-frequency changes in network congestion. The dataset comprises minute-by-minute snapshots of train flows, enabling us to monitor network congestion closely. Since we can observe both passenger and freight train counts, the empirical setting provides a unique opportunity to study section controllers' decisions directly as a function of the number of freight and passenger trains in the network. A preliminary analysis of this data substantiates the significant dwell times experienced by freight trains. We use this high-frequency data to estimate the parameters underlying the section controller's decisions. In order to do so, we estimate two discrete choice models corresponding to the section controllers' stop and hold decisions. The hold decisions are made instantaneously and frequently when the train is dwelling. The stop decisions are made after a train leaves the current station and before it stops at the next station (to allow the train to decelerate).

There are two main challenges in the estimation. First, decisions made at a particular station may be influenced by, and in turn influence, the decisions made at neighboring stations. Thus, the

[8] https://timesofindia.indiatimes.com/city/hubballi/many-passenger-trains-delayed-as-swr-gives-priority-to-freight/articleshow/71447387.cms

section controller's decisions are not independent. We build this dependency in decisions using Copula. Second, as in many discrete choice settings, observed covariates like the freight congestion variables (counts of freight trains in different parts of the network) are correlated with unobservables such as crew changes and unobserved maintenance schedules, and hence, are endogenous. Using recent literature, we construct network instrumental variables to correct for the bias caused by the endogeneity problem. Since our endogenous variables are count variables, we use the negative binomial distribution to model these and employ a control function approach for estimation.

Our estimation results reveal several insights. First, they confirm that passenger trains are accorded high priority in movement, significantly affecting freight trains' hold and stop decisions. We find that the likelihood of a freight train being held at a station increases with more passenger trains at the focal and neighboring stations. Second, we document a *strategic idling effect* [Baron et al., 2014] where a freight train is kept dwelling if a passenger train is approaching from behind. Third, we also document a *push effect* in the freight train queue where freight trains queued up behind a train nudge it to clear the occupied capacity (similar to queuing systems with finite buffers and no abandonment). Finally, we find that the section controllers strategically hold freight trains at stations with substantial capacities, such as junctions or stations with more tracks. The remainder of this chapter is organized as follows: In Section 1.2, we review the relevant literature; in Section 1.3, we describe the data and the empirical context; in Section 1.4, we present the model; in Section 1.5, we describe the estimation strategy; in Section 1.6, we discuss the results; and in Section 1.7, we conclude.

1.2 Literature Review

Our work is closely related to five streams of literature: (i) railway operations, (ii) queueing networks and their control, (iii) choice models, and (iv) structural estimation in operations management/research. We review the recent work in each of these streams and discuss our contributions

relative to these streams below.

Railway operations: The railway operations literature has got significant attention in the transportation and operations research communities. A body of work in this area focuses on the network design and expansion aspect [Barnhart et al., 2000, Bärmann et al., 2017, Wang et al., 2022]. A second literature stream focuses on the scheduling and routing of trains and crew planning in the railway setting [Caprara et al., 2002, Davarnia et al., 2019, Högdahl and Bohlin, 2023]. Finally, a body of work also focuses on real-time traffic control in railways [Lamorgese and Mannino, 2019, Dal Sasso et al., 2022]. Our work most closely aligns with the latter literature on real-time control. It is also closely related to the research on freight train operations. Caprara et al. [2011] study the problem of creating profitable freight routes in a railway corridor. Zhu et al. [2014] study the scheduled service network design problem for freight and integrate scheduling, car blocking, and routing decisions in an optimization model. In this literature, our work mostly relates to Borndörfer et al. [2016], which studies a freight traffic management setting in a congested network where both passenger and freight trains operate.

In contrast to the works that use optimization or game theoretic models, we use a data-driven approach to understand the section controller's preferences in managing congestion. To our knowledge, our research is the first empirical investigation of trains' dwell times via the section controller's decision-making. Our high-frequency micro-level data helps us estimate a structural model of the controllers' decision-making.

Queueing networks and their control: Our work also contributes to the literature on queueing networks that explores two specific subareas: priority queueing and strategic decision-making in queueing networks. Recently, emerging literature has explored priority queues like the ones found in amusement parks and order deliveries [Wang et al., 2015, Li et al., 2017, Gurvich et al., 2019]. Li et al. [2023] study patient prioritization in a hospital's emergency department and find that decision-makers use urgency-specific and delay-dependent prioritization rules instead of selecting

patients on a first-come, first-served basis.

Strategic queuing in network settings is another emerging research area to which our work contributes [Gopalakrishnan et al., 2016, Zhan and Ward, 2019, Arlotto et al., 2019, Zhong et al., 2023]. A strand in this literature focuses on the behavioral and strategic aspects of queueing control. Hathaway et al. [2023] analyze the behavior of frontline workers in service encounters, analyzing their decisions to either address customer requests themselves or escalate. Hathaway et al. [2021] use a structural model to study callback policies in a call center setting. They find that offering callbacks as a demand postponement strategy reduces average online waiting time. Unlike the majority of these studies, we empirically examine a controller's decisions for non-priority class customers (freight trains) within a network setting while focusing on quantifying the delays due to the priority structure.

Choice models: Choice models have been widely used by researchers in transportation science [McFadden, 1974, Ben-Akiva and Lerman, 1985, De Jong and Ben-Akiva, 2007] to model the transport mode chosen by an individual in both revealed and stated preferences studies. These studies have focused on estimating the parameters underlying an individual's decision-making, often described using a choice model. In the operations community, a body of literature has focused on optimizing assortments under the multinomial logit model [Gallego and Phillips, 2004, Talluri and van Ryzin, 2004, Rusmevichientong and Topaloglu, 2012, Gao et al., 2021]. This set of papers goes beyond estimating the parameters and uses the estimated parameters for a downstream optimization problem. More recently, papers in the operations research and management science community have used choice models for estimating spatial demand for bike-share systems and taxis in cities like Paris, London, and New York City [Kabra et al., 2020, He et al., 2019, Ata et al., 2019, Rosaia, 2022]. These studies use high-frequency granular data instead of survey/stated preferences data used in the early literature to estimate demand parameters, where demand is shaped by an underlying model of customer decision-making. Similar to these studies,

we use high-frequency granular data to estimate the parameters underlying a section controller's decision process. Methodologically, we use a choice model with a control function approach for estimation. The estimation procedure closely follows the work of Petrin and Train [2010], which introduces a control function approach to address endogeneity in consumer choice models. Similar methods have also been used in the empirical operations management literature [Guajardo et al., 2012, Feldman et al., 2022, Arora et al., 2023]. The method also involves using generalized residuals for discrete endogenous variables, which is standard in count data models [Wooldridge, 2014, Ragasa and Mazunda, 2018].

Structural estimation in OM/OR: We also add to the increasing literature on structural estimation in OM/OR with application areas in airlines, retail [Moon et al., 2018], healthcare [Kim et al., 2023], supply chain management [Olivares et al., 2008, Musalem et al., 2010], etc. In the airline setting, Deshpande and Arıkan [2012] use structural estimation techniques to analyze the deteriorating on-time performance of airlines, pointing to the underestimation of flight delay costs by airlines when setting scheduled block times. Arıkan et al. [2013] investigate the factors contributing to flight delays in the U.S. and identify bottleneck airports. The authors develop a stochastic model and use empirical data to study delay propagation. Barnhart et al. [2014] model the historical delays for U.S. domestic airline passengers. The authors use a multinomial logit model to estimate historical passenger travel and refine a previous benchmark heuristic to determine resulting passenger delays. Li et al. [2014] estimate the degree of strategic behavior by customers when buying airline tickets. They find that customers strategically delay purchasing airline tickets in anticipation of lower prices in the future.

1.3 Data and Empirical Context

1.3.1 Indian Railways

Indian Railways, India's largest public sector enterprise, is the backbone of the country's logistics sector. Annually, it efficiently transports over 1.4 billion tonnes of freight and boasts a vast network that extends to virtually every corner of the nation. This extensive reach is vital in propelling the country's social and economic development. The Indian Railway network is the fourth-largest globally, behind the United States, Russia, and China. It is also the world's largest network under a single management and has over 7,000 stations with a track length of more than 68,000 kilometers. The network is organized hierarchically into 17 geographic zones, which are further divided into 68 divisions and numerous sections, which are stretches of railway lines comprising multiple stations. As of 2019-20, the Indian Railways ran around 13,100 passenger trains daily and transported an average of 8.1 billion passengers annually.[9] In addition to passenger trains, Indian Railways also operates freight trains that play a crucial role in facilitating the movement of goods across the country.[10] In fact, freight is the biggest revenue-generating segment for Indian Railways, accounting for nearly 70% of total revenue in FY22.[11] During the same fiscal year, Indian Railways transported over 1,400 million tonnes of freight.[12]

1.3.2 Section Controllers' Role in Freight Movement

Railway organizations across the globe use a centralized control mechanism for managing traffic movement. The broad objective of the control mechanism is to handle train dispatches by making hold and stop decisions. Two scenarios where these decisions are important are the *meet* and

[9] https://indianrailways.gov.in/railwayboard/uploads/directorate/stat_econ/Annual-Reports-2019-2020/Year-Book-2019-20-English_Final_Web.pdf

[10] Freight operations cater to a diverse range of commodities like coal, agricultural products, minerals, etc.

[11] https://www.ibef.org/industry/indian-railways

[12] https://tinyurl.com/re2623ks

overtake scenarios. The meet scenario involves deciding which trains to stop when trains traveling in opposite directions cross each other. The overtake scenario involves stopping one train to let another train overtake the former. The necessity for such decisions arises due to the differences in speeds and a priority structure across different types of trains.

Like railways worldwide, a central controller manages traffic in railway networks. In the context of Indian Railways, these traffic controllers are called *Section Controllers*. The overall network is divided into different sections, each managed by a specific controller who makes decisions for trains in their sections. What makes the role of section controller even more important in Indian Railways is that freight trains are unscheduled, unlike passenger trains that run on predetermined schedules. While the section controller oversees passenger train movements, controllers have limited control due to their predefined schedules. On the contrary, freight train movements are precisely managed by section controllers, who may hold or stop a freight train if a railway section is congested or if passenger trains are in the queue that take priority in dispatch.

The section controllers exert significant influence over freight movement in the railway networks. They coordinate with station managers at different train stations to ensure the smooth movement of trains in and out of the station. Their job is highly demanding as they need to coordinate with multiple station managers and control the flow of various trains. To manage the train movements, section controllers use a software application (known as the *control office application*) that provides real-time information about trains.[13]

1.3.3 Data

We use two different sources of data for our study: (i) railway-network characteristics data, and (ii) passenger and freight train flow data. We briefly describe these data sources below.

[13] https://cris.org.in/crisweb/design1/proCOA.jsp

Railway-Network Characteristics Data

First, we obtained the railway-network characteristics data from the Deen Dayal Upadhyay division of the East Central Railway zone (see Figure 1.1 right panel) from two sources: rbs.indianrail.gov.in & ecr.indianrailways.gov.in. The data contain information on various station level characteristics: (i) station identifiers and locations, (ii) type of station (regular or junction), and (ii) number of railway tracks. It also contains information on the track segments between two stations, also known as *block sections*: (i) block section length and (ii) number of lines.

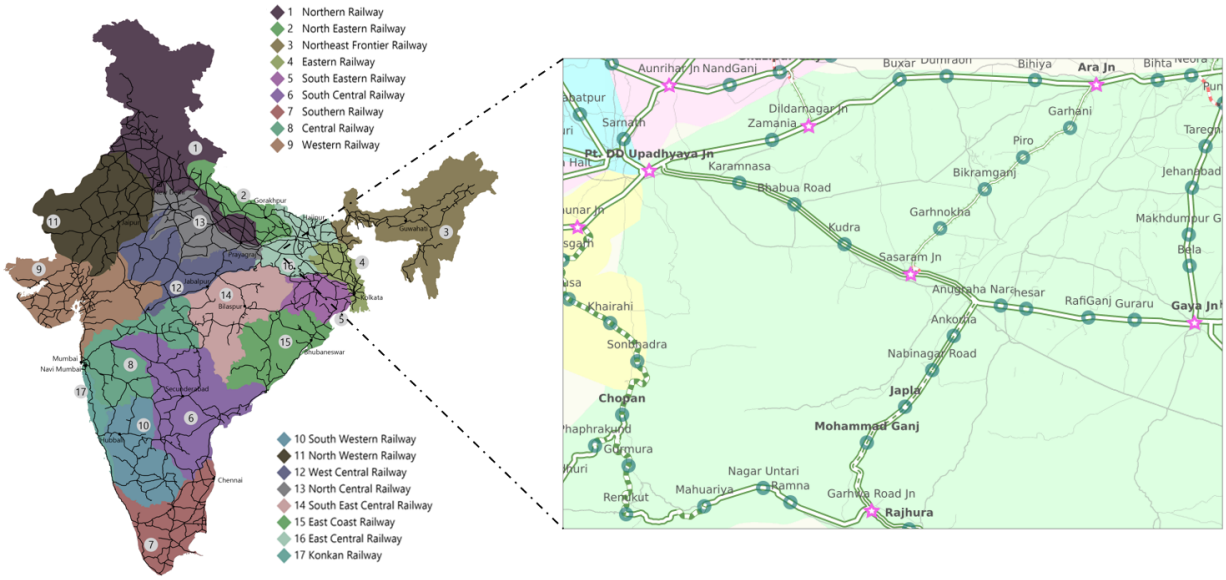

Figure 1.1: Indian Railways network and the different divisions (left) & Deen Dayal Upadhyay Division zoomed in (right). The stars depict junction stations.

The Deen Dayal Upadhyay division has 73 stations, of which 66 are regular stations and 7 are junction stations. The average number of tracks for these stations is 4.35. The division is divided into five administrative sections, each managed by a different section controller. The mean inter-station distance in this section is around 5.38 kilometers. There are 25 block sections on the triple line, 51 on the double line, and 5 on the single line corridors.

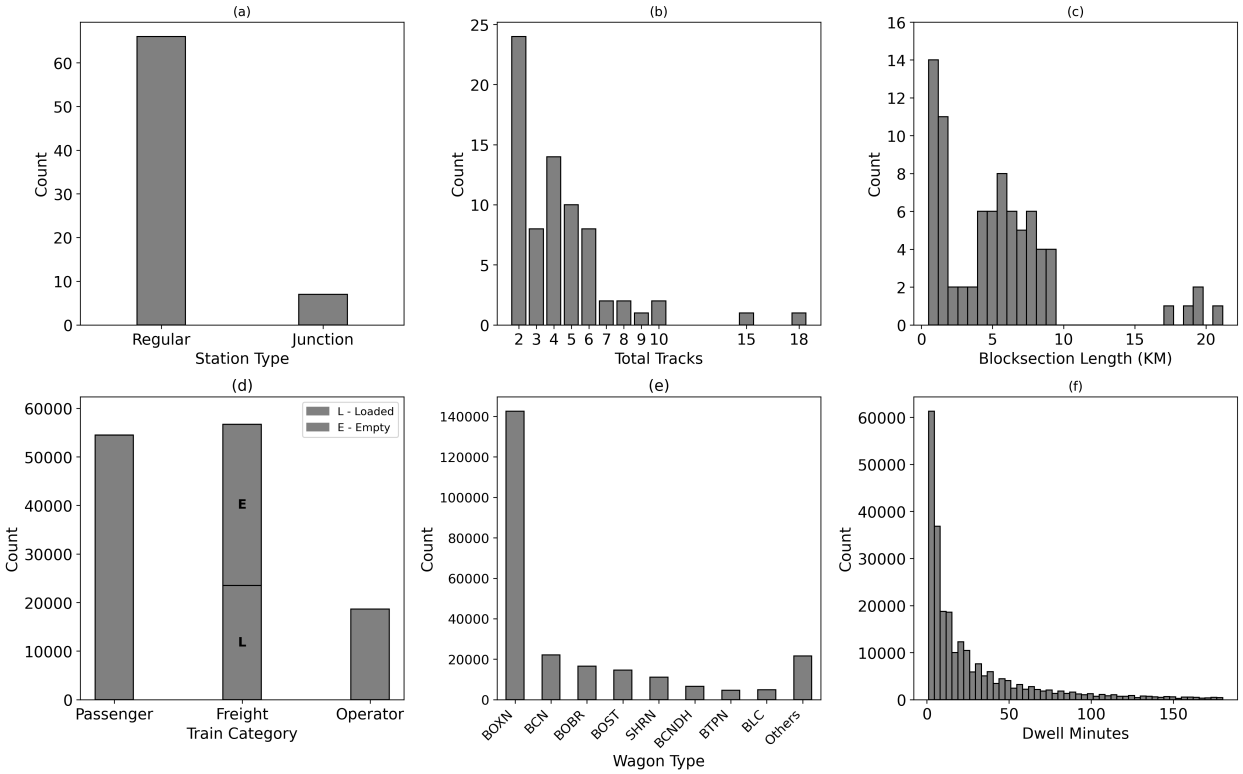

Figure 1.2: Station and train level characteristics: (a) Station types (b) Number of tracks (c) Distribution of block section length (km) (d) Train categories (e) Wagon types (The bottom 10 percentile are tagged as 'Others') & (f) Distribution of dwell time (conditional on a freight train stopping).

Passenger and Freight Train Flow data

We also obtained passenger and freight train flow data for the same division from Indian Railways. The dataset spans from 1^{st} September 2021 to 30^{th} August 2022 and covers over 54,000 passenger and 56,000 freight trains. The dataset contains information on (i) timestamps for each train's arrival and departure on/from a station, (ii) scheduled time of arrival and departure for passenger trains, and (iii) train level characteristics (type of train - freight/passenger, whether the train was loaded/empty, etc.).

We calculate a freight train's dwell time at a particular station from the arrival and departure timestamps. We remove observations with a dwell time greater than 180 minutes.[14] We observe

[14]We decided the threshold of 180 minutes after consultation with Indian Railways. This threshold was set to exclude extreme cases, such as data entry errors or rare disruptions (e.g., locomotive failures causing multi-day delays) that do not reflect normal operations.

an average dwell time of 4.74 minutes per freight train per station. We note that 83.5% of our data have zero dwell time (when the train does not stop at a station). Conditional on a freight train stopping, the average dwell time is around 27.54 minutes. Figure 1.2 shows the summary statistics for the station and train characteristics data. In calculating the statistics, we drop the freight train data on one single-line section due to negligible freight traffic, and one section because it only contains a high percentage of missing arrival and departure timestamps. We report our data cleaning procedure in Appendix 1.8.1.

High-frequency congestion data: We are interested in modeling a freight train i's dwell at a particular station s as a function of the congestion in the network. To do so, we construct measures of network congestion by counting the number of passenger and freight trains at all stations and block sections, *every minute*. The constructed congestion vector Z contains elements of the type s_d_f, where $s \in$ {Current, Previous, Next} denotes the current, previous and next station; $d \in$ {Up, Down} denotes the upstream ($\rightarrow$) and downstream ($\leftarrow$) directions; and $f \in$ {P,G} denotes passenger or freight trains. For example, the congestion variable Current_Up_G encodes the count of freight trains at the focal station that are traveling upstream. Similarly, the variable Next_Up_P represents the total number of passenger trains in the next upstream station's direction.[15] For example, in the scenario depicted in Figure 1.3, the congestion variable Previous_Up_P takes the value one and the variable Next_Down_G takes the value two. All other variables in the congestion vector are zero.

1.3.4 Model-free Evidence

Before introducing our structural model, we provide some model-free evidence that motivates our structural model. First, we show that the network-level congestion variables significantly affect freight trains' dwell times. Second, we show that junctions and stations at section ends have much

[15] For computing the next (previous) variables, we count the total number of trains at the next (previous) station and the corresponding block section.

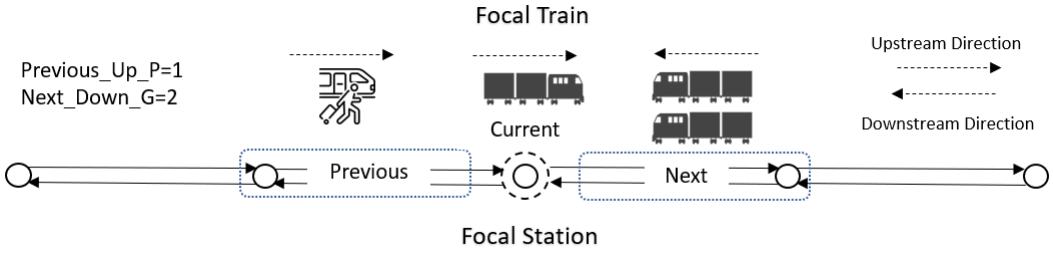

Figure 1.3: A snapshot of the congestion variables in the network. The variables are constructed by counting the passenger and freight trains in the upstream (→) and downstream (←) directions.

larger dwell times than other stations, and dwell times are heterogeneous across different hour slots in a day. All these pieces of evidence help us motivate our model in Section 1.4.

1. Network Congestion Variables: To test whether network-level congestion variables affect dwell times, we run an ordinary least squares regression with total dwell time as the dependent variable and all network congestion variables as explanatory variables.[16] Figure 1.4 shows the estimates from the OLS regression with the associated 95% confidence intervals. All the network congestion variables are significant at a 1% significance level. The adjusted R^2 of the model is 0.93. This shows that dwell times at a station are not only correlated with the focal station's congestion variables but also with the previous and next stations in the network.

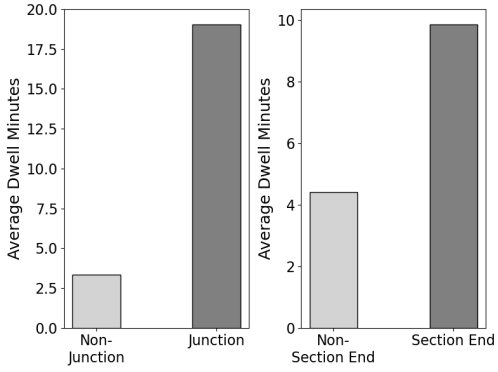

Figure 1.5: Dwell times per train per station for junction and section end stations.

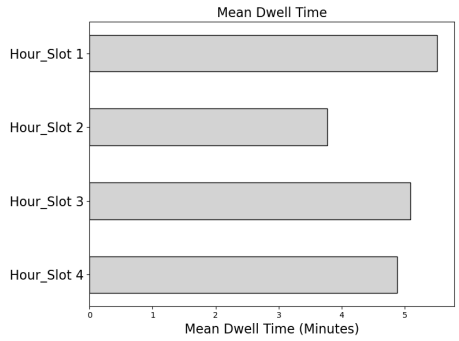

Figure 1.6: Dwell time variation by hour slot.

[16]The dwell time variable is calculated as the total dwell across all trains in 15-minute intervals. The congestion variables are summed up across the 15-minute interval.

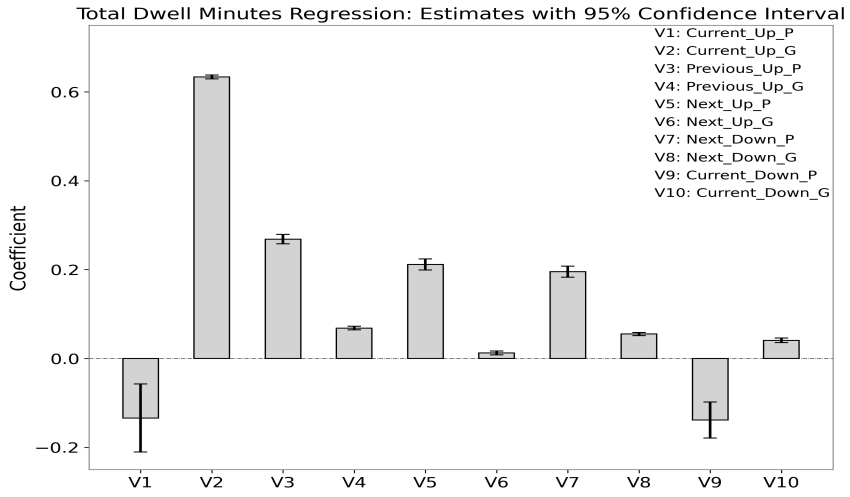

Figure 1.4: OLS estimates and the associated 95% confidence intervals for the total dwell regression (MTGE–PSE block section).

2. Station- and time-level characteristics: We also show that station-level characteristics such as junctions and section ends significantly affect the dwell times. Figure 1.5 shows that the average dwell times for junction stations (19.06 minutes) are 5.75 times that of non-junction stations (3.31 minutes). Similarly, the dwell times at section ends are much higher than at non-section-end stations. We include both these indicator variables as station-level controls in our model. Figure 1.6 depicts the heterogeneity in the dwell times across the six different hour slots in the day. We include fixed effects for these slots in our model to account for this time heterogeneity.

1.4 Model

In this section, we introduce a model to study how section controllers decide to hold and stop freight trains in the network to manage their movement. These decisions ultimately determine the overall freight delay and speed in the network. Hence, modeling them is central to analyzing any downstream policies for delay reduction and speed improvement. In the rest of this section, we introduce the section controllers' decisions and describe why, when, and how they are made.

At any time t, a train is either dwelling at a station or moving, which results in the two sets: (i) dwelling and (ii) moving trains.[17] The section controller makes two types of decisions for these two sets of trains. For a dwelling train, the section controller must decide whether to release the train or continue holding it in the following period. For the moving train, the controller must decide whether or not to stop it at the upcoming station. We call these decisions the *hold/release* and *stop/pass-through* decisions, respectively.

The controller may decide to hold or stop a freight train if the upstream network is congested by many passenger and freight trains. She may also hold a train to let the passenger trains coming from behind pass, since passenger trains get priority over freight trains. On the other hand, she may release a train if freight trains are coming from behind, thus creating a push effect. If the network is congested, a train may be more likely to be stopped/held at a station with a larger capacity (like a junction or a station with multiple tracks). Our model includes the upstream and downstream congestion variables, station and train-level characteristics, and time effects as explanatory variables to capture all these effects.

The hold/release (H/R henceforth) decision is made by the controller *instantaneously* in every period t when the train is dwelling. After consulting with the practitioners, we allowed the length of t to be one minute. Thus, the H/R decision is made frequently. On the other hand, the stop/pass-through (S/P henceforth) decision is made after a freight train leaves the current station and before it reaches the next station (to allow the train to decelerate before stopping). Also, the S/P decision has to be communicated to the station manager before the train reaches the next station. We note that we *do not* observe the exact time of the S/P decision in our data; we only observe whether the train was stopped. Hence, we allow for a hyperparameter $k \in [0, 1]$, representing the fraction of the inter-station distance covered before the S/P decision is made.[18] Figure 1.7 shows the stop and hold decisions in a station-time graph. Next, we describe how these decisions are made.

[17] Trains rarely dwell between stations.

[18] We estimated k in a data-driven manner by minimizing the AIC and include this analysis in Appendix 1.8.1.

1.4.1 Decision 1 (Hold/Release)

For a dwelling train, the section controller chooses one of the two options: $X \in \{H, R\}$. We infer these decisions from a train's arrival and departure time stamps. Specifically, if t^s_{id} is the time of departure of a train i from station s, we infer that the train is released at time t if $t^s_{id} = t$. If $t^s_{id} > t$, the train is held during time period t. We set the binary variable $x_{ist} = 1$ if the train i was held at station s and time t and set it to 0 otherwise. We model this H/R decision with a discrete choice framework. Specifically, the utility obtained from choosing the hold option ($x_{ist} = 1$) is:

$$V_{istx} = \alpha_x + \Theta_1^{'} Z_{st} + \Gamma_1^{'} S_s + \Delta_1^{'} R_i + \Omega_1^{'} T_t + \Upsilon_1^{'} Q_{ist} + \xi^1_{ist} + \epsilon^1_{istx}. \tag{1.1}$$

where Z_{st} is a 11 $\times$ 1 vector of network congestion variables consisting of passenger and freight train counts on the focal station and upstream and downstream links.[19] S_s is a 7 $\times$ 1 vector of station-level characteristics: number of tracks on the station, dummy variables for the focal and next stations being junctions, a dummy variable for the station being at the end of the section, distance of the station from the next, number of lines in the upstream blocksection, and a categorical variable for the railway section. R_i is a 5 $\times$ 1 vector of train-level controls: a dummy variable for the train being loaded or empty, a categorical variable for different wagon types, categorical variables for origin and destination of the train, and a categorical variable for locomotive type used for the train. T_t are time-fixed effects for the day of the week (weekday/weekend only) and the hour of the day slots.[20] Q_{ist} is a 2 $\times$1 vector consisting: a binary variable indicating whether the train is delayed, and the remaining distance to the destination.[21]

ξ^1_{ist} are unobservables that affect the section controllers' decision to hold the trains. These are factors observable to the controllers but not to us as researchers. These unobservables may

[19] We decompose the congestion variable for the current station into two: freight trains that arrived at the station before the focal train (Current_Up_G_Pre) and ones that came after (Current_Up_G_Post). Please see Appendix 1.8.1.

[20] We divide the day into four slots of six hours each starting from midnight and include three fixed effects in the model.

[21] We use log transformation for the remaining distance to the destination.

include crew changes for freight trains, scheduled track maintenance, and any other shocks leading to station congestion. We address the endogeneity issues that arise from these unobservables in Section 1.5. ϵ^1_{istx} are independent and identically distributed idiosyncratic errors that follow type I extreme value distribution. The section controller chooses to hold the train if $V_{istx} > 0$ for $x = 1$. The resultant choice probability of holding for a train conditional on it being dwelling is given by:

$$P_{istx} = \frac{\exp(\alpha_x + \Theta_1' Z_{st} + \Gamma_1' S_s + \Delta_1' R_i + \Omega_1' T_t + \Upsilon_1' Q_{ist} + \xi^1_{ist})}{1 + \exp(\alpha_x + \Theta_1' Z_{st} + \Gamma_1' S_s + \Delta_1' R_i + \Omega_1' T_t + \Upsilon_1' Q_{ist} + \xi^1_{ist})}. \tag{1.2}$$

1.4.2 Decision 2 (Stop/Pass-through)

In addition to making the H/R decision, the section controller chooses one of the two options $Y \in \{S, P\}$ for a train moving between station $s-1$ and s. If $Y = S$, the train is stopped at station s and passed through otherwise. Again, we infer these decisions with data from a train's arrival and departure time stamps. Specifically, if t^s_{ia} and t^s_{id} be the times of arrival and departure of a train i at/from station s, we infer that the train was stopped at s if $t^s_{id} > t^s_{ia}$. We set the binary variable $y_{ist} = 1$ if the train was stopped and set it to 0 otherwise. We also model this binary decision with a discrete choice framework. Specifically, the utility obtained from choosing the stop option $y_{ist} = 1$ for train i and station s at time t is:

$$U_{isty} = \alpha_y + \Theta_2' Z_{st} + \Gamma_2' S_s + \Delta_2' R_i + \Omega_2' T_t + \Upsilon_2' Q_{ist} + \xi^2_{ist} + \epsilon^2_{isty}. \tag{1.3}$$

ξ^2_{ist} are the unobservables that affect the stop decision, and ϵ^2_{isty} are independent and identically distributed idiosyncratic errors that follow type I extreme value distribution. The section controller stops the train if the corresponding utility is positive. The choice probability of stopping the train is given by:

$$P_{isty} = \frac{\exp(\alpha_y + \Theta_2' Z_{st}^{sp} + \Gamma_2' S_s + \Delta_2' R_i + \Omega_2' T_t + \Upsilon_2' Q_{ist} + \xi_{ist}^2)}{1 + \exp(\alpha_y + \Theta_2' Z_{st}^{sp} + \Gamma_2' S_s + \Delta_2' R_i + \Omega_2' T_t + \Upsilon_2' Q_{ist} + \xi_{ist}^2)}. \tag{1.4}$$

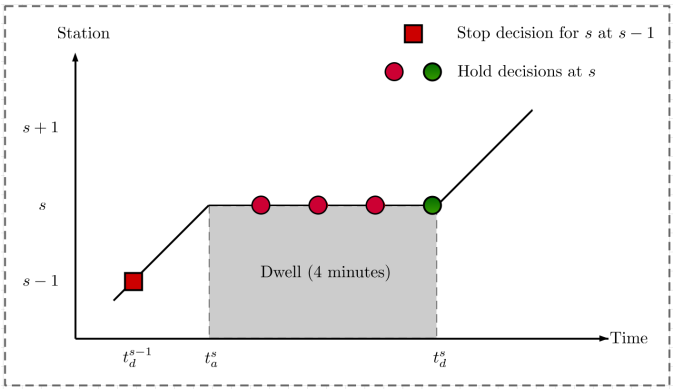

Figure 1.7: Illustration of the H/R and S/P decisions.

1.4.3 Modeling Dependency in Decisions

We first highlight that observations in our data originate from spatially proximate locations (stations). Analyzing such data requires special attention due to the presence of spatial effects. Specifically, our data exhibit *spatial autocorrelation*, where data points located closer together tend to be more (or less) similar. This is a common characteristic of spatial data and has significant implications for estimation. In our context, spatial correlation implies that the section controllers' HR and SP decisions are not independent across stations.

Second, our data also exhibit *spatial heteroscedasticity*, wherein the variance of the error terms varies across different stations (due to local shocks). Ignoring these spatial autocorrelation and heteroskedasticity effects can lead to biased and inconsistent estimates, particularly in non-linear models like discrete choice models [Bhat and Sener, 2009]. Hence, we need a model that can explicitly account for these effects. Our structural model accounts for the correlation between HR and SP decisions across space for trains traveling in *same and opposite directions*.

Copula Approach:

Following Bhat and Sener [2009], we use a *copula approach* to model the spatial correlation between decisions by allowing pairwise correlation across observational units that are close in space. Specifically, we model the correlation between the logistic error terms of different observational units using a multivariate logistic distribution based on the Farlie-Gumbel-Morgenstein (FGM) copula. We next describe the copula approach and how our model handles multiple correlations.

The copula approach: Described simply, a copula is a mathematical function $C : [0,1]^n \rightarrow [0,1]$ that links the marginal distributions of individual random variables (say the spatial logistic errors $\epsilon_1, \epsilon_2 ... \epsilon_n$) to their joint multivariate distribution ($F(\epsilon_1, \epsilon_2, \ldots, \epsilon_n)$). Specifically,

$$F(\epsilon_1, \epsilon_2, \ldots, x_n) = C\left(F_1(\epsilon_1), F_2(\epsilon_2), \ldots, F_n(\epsilon_n)\right) \tag{1.5}$$

where, $F_i(\epsilon_i)$ are the marginal CDFs of the individual random variables and C is the copula function, capturing the dependence structure between the random variables. We use the FGM copula form for the function C to link the spatially dependent logistic error terms.[22] Given the simplest case of two spatially dependent error terms ϵ_1 and ϵ_2, the FGM copula linkage takes the following form:

$$F(\epsilon_1, \epsilon_2) = F(\epsilon_1)F(\epsilon_2)\left[1 + \theta(1 - F(\epsilon_1))(1 - F(\epsilon_2))\right] \tag{1.6}$$

where θ captures the degree of spatial dependence between the two errors. If $\theta = 0$, we are back to the independent error scenario. Arnold [1992], Karunaratne and Elston [1998] propose the following multivariate extension of this bivariate logistics distribution:

$$F(\epsilon_1, \epsilon_2, \ldots, \epsilon_k) = \left[\prod_{i=1}^{k} F(\epsilon_i)\right] \times \left[1 + \sum_{i<j} \delta_{ij}(1 - F(\epsilon_i))(1 - F(\epsilon_j))\right]. \tag{1.7}$$

[22] The choice of the functional form of the copula depends on the application. FGM copulas are analytically simple for both interpretation and computation.

Formulation of Likelihood function: We build on the approach proposed by Bhat and Sener [2009] to formulate the likelihood function for all the decisions in our data. Assume that we observe M Hold/Release (HR) and O Stop/Passthrough (SP) decisions in our data. Then, the probability of observed decisions (d_q and d_k) can be written as the following:

$$P(d_1 = x_1, \ldots, d_M = x_M, d_{M+1} = y_1, \ldots, d_{M+O} = y_O) = \left[\prod_{q=1}^{Q} \frac{e^{w_{iste_q} \cdot d_q}}{1 + e^{w_{iste_q}}}\right]$$

$$\left[1 + \sum_{q=1}^{Q-1} \sum_{k=q+1}^{Q} (-1)^{d_q + d_k} \cdot \theta_{qk} \left\{1 - \frac{e^{w_{iste_q} \cdot d_q}}{1 + e^{w_{iste_q}}}\right\} \left\{1 - \frac{e^{w_{iste_k} \cdot d_k}}{1 + e^{w_{iste_k}}}\right\}\right]$$

where,

$$w_{iste_q} = \begin{cases} \alpha_x + \Theta_1^{'} Z_{st} + \Gamma_1^{'} S_s + \Delta_1^{'} R_i + \Omega_1^{'} T_t + \Upsilon_1^{'} Q_{ist} + \xi_{ist}^1, \text{ if } 1 \leq q \leq M (\text{H/R Decisions}), \\ \\ \alpha_y + \Theta_2^{'} Z_{st} + \Gamma_2^{'} S_s + \Delta_2^{'} R_i + \Omega_2^{'} T_t + \Upsilon_2^{'} Q_{ist} + \xi_{ist}^2, \text{ if } M+1 \leq q \leq M+O (\text{S/P Decisions}). \end{cases}$$

Here θ_{qk} captures the spatial correlation between observational unit q and k. Note that it is not possible to estimate a separate θ_{qk} parameter for each pair of observations. To overcome this, we characterize θ_{qk} as the following:

$$\theta_{qk} = \tanh(\delta' S_{qk}) = \left[\frac{e^{\delta' S_{qk}} - e^{-\delta' S_{qk}}}{e^{\delta' S_{qk}} + e^{-\delta' S_{qk}}}\right], \tag{1.8}$$

where, S_{qk} is a vector of variables that influence the level of spatial correlation between observational units q and k, and $\delta = [\delta_1, \ldots, \delta_6]'$ is a 6×1 vector of correlation parameters. There are a couple of advantages of this characterization. First, our model can account for multiple levels of correlations. We explain this in detail in the next subsection. Second, this characterization

improves on Bhat and Sener [2009] by allowing for both positive and negative correlations.[23]

Incorporating multiple spatial correlation parameters: Our model effectively captures six distinct correlations among decisions by incorporating multiple copula parameters. Specifically, we estimate correlations between spatially adjacent decisions of three types: HR-HR, SP-SP, and HR-SP. Furthermore, our model distinguishes between correlations based on whether the trains involved are moving in the same or opposite directions, resulting in six unique copula parameters (as detailed in Table 1.1). For instance, the parameter δ_1 represents the correlation between two HR decisions for trains at adjacent stations traveling in the same direction.

q and k (next station)	Binary Spatial Variables	Copula Parameter	Observations
HR-HR & Same Direction	s^1_{qk}	δ_1	87859
HR-HR & Opposite Direction	s^2_{qk}	δ_2	155544
HR-SP/ SP-HR & Same Direction	s^3_{qk}	δ_3	47196
HR-HR/ SP-HR & Opposite Direction	s^4_{qk}	δ_4	44664
SP-SP & Same Direction	s^5_{qk}	δ_5	7518
SP-SP & Opposite Direction	s^6_{qk}	δ_6	7406

Table 1.1: Capturing Spatial Dependency through Copulas.

The vector S_{qk} contains elements that are of the binary type. Let S^i_{qk} be the i^{th} element of S_{qk}. If q and k are observations corresponding to HR decisions made at two neighboring stations at the same time for freight trains moving in the same direction, then $S^1_{qk} = 1$ and $S^i_{qk} = 0 \forall i \neq 1$. Similarly if q is an observation corresponding to a SP decision and k corresponds to an HR decision made at two neighbouring stations at the same time for freight trains moving in the opposite directions then $S^4_{qk} = 1$ and $S^i_{qk} = 0 \forall i \neq 4$.[24]

[23] Bhat and Sener [2009] propose the following characterization of θ_{qk}:

$$\theta_{qk} = \pm \left[\frac{(e^{\delta})' s_{qk}}{1 + (e^{\delta})' s_{qk}} \right]. \tag{2}$$

A limitation of this characterization is the need to know the nature of the correlation a priori. In the example study used by Bhat and Sener [2009], a '+' sign is imposed on θ_{qk} based on the domain knowledge. This approach is restrictive, as in many cases, the nature of the dependence is not known in advance.

[24] Note that we populate the S^i_{qk} vectors based on Table 1.1 only when two observations have the exact same time stamp (minute level). As explained earlier, using Copula, our objective is to capture the dependence in decisions made at the same time (contemporaneous) because the lagged effects are already captured through network congestion characterization, as explained in the section.

Incorporating Spatial Heteroscedasticity: A unique feature of this spatial model is that we can directly capture heteroscedasticity and correlation in the error terms ϵ_q across observational units, rather than using a pre-specified spatially autoregressive structure to generate dependence. In our copula approach, the spatial correlation effect is completely delineated from heteroscedasticity effects, allowing a clear capture and testing of each of the heteroscedasticity and correlation effects distinctly. In line with Bhat and Sener [2009], we parameterize σ_q as $exp(\Xi' \varpi_q)$, where ϖ_q represents the variables associated with heteroscedasticity and Ξ is the vector of corresponding coefficients to be estimated.

$$P(d_1 = x_1, \ldots, d_M = x_M, d_{M+1} = y_1, \ldots, d_{M+O} = y_O) = \left[\prod_{q=1}^{Q} \frac{e^{\frac{w_{iste_q}}{\sigma_q} \cdot d_q}}{1 + e^{w_{iste_q}}} \right]$$

$$\left[1 + \sum_{q=1}^{Q-1} \sum_{k=q+1}^{Q} (-1)^{d_q + d_k} \cdot \theta_{qk} \left\{ 1 - \frac{e^{\frac{w_{iste_q}}{\sigma_q}}}{1 + e^{\frac{w_{iste_q}}{\sigma_q}}} \right\} \left\{ 1 - \frac{e^{\frac{w_{iste_k}}{\sigma_k}}}{1 + e^{\frac{w_{iste_k}}{\sigma_k}}} \right\} \right].$$

1.5 Estimation

We describe the estimation of our model in this section. We first explain the problem of endogeneity in the analysis and show that failing to correct it can lead to biased parameter estimates. Then, we describe the set of instrument variables that we employ to address this problem. Finally, we present our estimation algorithm.

1.5.1 Endogeneity

Endogeneity arises because observed network congestion variables may correlate with unobservable factors known to section controllers, influencing their hold/release or stop/pass decisions. For instance, controllers may have private knowledge of track maintenance schedules, leading to longer train holds. Ignoring such unobservables can bias parameter estimates and distort counterfactual policy simulations. Key unobservables that we do not observe in data include crew change sched-

ules—where shifts typically last six hours and changeovers create congestion shocks—and track maintenance schedules, which allocate daily time slots for repairs, increasing station congestion and train hold probabilities. Omitting these factors may underestimate congestion effects, necessitating an endogenous treatment of freight train congestion variables. We provide more details on these unobservables in Appendix 1.8.2. Mathematically, the problem of endogeneity means that:

$$\mathbb{E}\left[z_{st}.\xi_{ist}\right] \neq 0.$$

1.5.2 Network Instrumental Variables

Recent literature has increasingly emphasized the use of network instruments to address endogeneity challenges in network settings [Bramoullé et al., 2009, Drakopoulos and Zheng, 2017, Kim et al., 2023]. The key idea is to use spatially lagged exogenous characteristics in the network as instrument variables for focal endogenous variables. Specifically, let G be a spatially weighted matrix that encodes the extent to which a station j can influence station i by assigning weights g_{ij} to each connection ij. Proposition 3 of Drakopoulos and Zheng [2017] gives us two conditions under which network-lagged measurements of station characteristics S_s serve as valid instrumental variables for endogenous variables at focal station s: (i) the variables in S_s are exogenous, and (ii) the matrices I, G, G^2 are linearly independent. Under these conditions, G^2S_s serve as valid instrumental variables. Intuitively, G^2S_s capture the indirect effects of station s's neighbor's neighbor's exogenous covariates i.e. S_{s-2} on endogenous congestion variable z_{st}. In other words, it captures the influence of nodes that are two steps away in the network. Below, we explain how G^2S_s satisfies the necessary conditions for instrument validity—relevance and exclusion—in our specific setting.

Relevance: In our setting, relevance is fulfilled because lagged station characteristics are correlated with local congestion. For example, if a second-lagged station has limited tracks or high

historical congestion, trains may be delayed upstream, causing congestion that propagates to the focal station. Mathematically, station-specific characteristics at the second-lagged station S_{s-2} directly influence congestion variables at the second-lagged station i.e. $\mathbb{E}[S_{s-2} \cdot z_{(s-2)t}] \neq 0$. Now, since the decisions at focal station s depend on congestion at station $s-2$, the congestion at two stations is also correlated, i.e. $\mathbb{E}[z_{st} \cdot z_{(s-2)t}] \neq 0$. Combining the two, $\mathbb{E}[z_{st} \cdot S_{(s-2)}] \neq 0$. Finally, since $G^2 S_s$ is just a weighted sum of second-lagged characteristics,

$$\mathbb{E}\left[G^2 S_s . z_{st}\right] \neq 0.$$

which satisfies the relevance condition.

Exclusion: Our instrumental variables also satisfy the exclusion restriction. First, the station characteristics S_{s-2} are purely exogenous and do not change in response to the focal station's congestion. Second, $G^2 S_s$ only affect S_s through the intermediate node, ensuring that their impact on the focal node occurs only via network spillover($G^2 S_s \rightarrow z_{(s-2)t} \rightarrow z_{(s-1)t} \rightarrow z_{(s)t}$). Thus,

$$\mathbb{E}\left[G^2 S_s . \xi_{st}\right] = 0.$$

which satisfies the exclusion restriction condition.

We acknowledge that higher-order spatially lagged station characteristics, such as $G^3 S_s, G^4 S_s, \ldots$, may also serve as valid instruments. However, as the lag order increases, the strength of the instruments diminishes due to the progressive dilution of the network effect. Hence, we prioritize second-order spatially lagged instrumental variables ($G^2 S_s$) to employ strong instruments in our analyses.

Construction of the G matrix: We can use different weights $g_{ij}$to construct the G matrix as long as matrix G is row normalized. Some examples of these weights are binary adjacency, inverse of distance, and flow-based weights. We chose inverse distance-based weights ($1/d_{ij}$) as they

allow nearby stations to exert a stronger influence on the endogenous variables than those that are farther away. We explain the construction of network instruments with an example in Figure 1.8, using a linear network of 5 stations. Here, Z represents the endogenous congestion variable. The G matrix creates spatially lagged measures of exogenous station characteristics. For instance, if S represents station characteristics, then G^2S serves as a valid instrument for Z. Specifically, $G_{32}G_{21}S_1$ works as an instrument for Z_3 because S_1 does not directly impact Z_3 but only through Z_1, thereby meeting the requirements for relevance and exclusion.

Station	**Station Characteristic**	**Endogenous Congestion Variable**	**2-Lagged (G^2S) Instruments**
1	S_1	Z_1	-
2	S_2	Z_2	-
3	S_3	Z_3	-
4	S_4	Z_4	-
5	S_5	Z_5	$\frac{w_{52}}{w_{52}+w_{54}}S_1 + \frac{w_{54}}{w_{52}+w_{54}}S_3$
6	S_5	Z_5	$\frac{w_{52}}{w_{52}+w_{54}}S_2 + \frac{w_{54}}{w_{52}+w_{54}}S_4$
7	S_5	Z_5	S_5

Table 1.2: Table of Instruments using Station Characteristics.

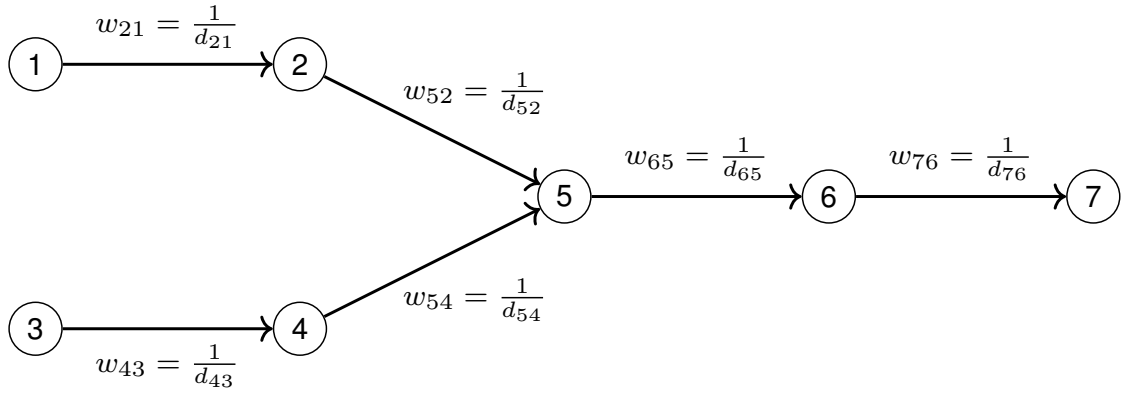

$$G = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & \frac{w_{52}}{w_{52}+w_{54}} & 0 & \frac{w_{54}}{w_{52}+w_{54}} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}, \quad G^2 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{w_{52}}{w_{52}+w_{54}} & 0 & \frac{w_{54}}{w_{52}+w_{54}} & 0 & 0 & 0 & 0 \\ 0 & \frac{w_{52}}{w_{52}+w_{54}} & 0 & \frac{w_{54}}{w_{52}+w_{54}} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}.$$

Figure 1.8: Non-linear Network and Row Normalized G Matrix for Instruments.

1.5.3 Control Function Approach to Estimation

Our estimation proceeds in two steps. First, we estimate control functions for all endogenous variables in the model with the appropriate instruments described in Section 1.5.2. Second, we use these control functions to estimate the parameters of the H/R and the S/P decisions via maximum likelihood estimation.

Step 1 - Estimating Control Functions: The key idea in the control function approach is to control for the part of the error terms in Equations (1.1) and (1.3) that is correlated with the endogenous variables. This endogenous part is obtained using residuals from a first-stage regression of the endogenous variables on the instruments and the control variables. Recall that all the freight train-related congestion variables are endogenous in our setting, which gives us a total of six (five) endogenous variables for the H/R (S/P) model. Next, we show how we estimate these residuals.

We first note that all our endogenous variables are count variables; hence, linear regressions are inappropriate for estimating the corresponding first stages. We model our endogenous variables using negative binomial distributions. We report the distribution of these count variables in Tables 1.6 and 1.7 in the Appendix. We also note from the tables that a large proportion of these variables is zero. This is unsurprising since these variables are *minute-level* train counts in different parts of the railway network. Given the high zero percentage, we model these variables using zero-inflated negative binomial (ZINB) models. Specifically, let Z_{st} be the set of variables corresponding to the ZINB models. Let $z_{ast} \in Z_{st}$, where $a \in \{1, \ldots, 6\}$. Let Y_a be the random variables corresponding to z_{ast}. The probability distributions for the random variables Y_a are given by:

$$P(Y_a = z_{ast}) = \begin{cases} \pi + (1-\pi)\left(\dfrac{\theta}{\theta+\lambda}\right)^{\theta} & \text{if } z_{ast} = 0, \\ (1-\pi)\dfrac{\Gamma(z^{ast}+\theta)}{\Gamma(\theta)\Gamma(z_{ast}^2+1)}\left(\dfrac{\theta}{\theta+\lambda}\right)^{\theta}\left(\dfrac{\lambda}{\theta+\lambda}\right)^{z_{ast}} & \text{if } z_{ast} > 0. \end{cases} \tag{1.9}$$

where, Equation (1.9) correspond to the ZINB models, π is the zero-inflation probability and λ is the mean of the negative binomial random variables. The parameter θ represents the degree of overdispersion. We model π and λ using probit and exponential functions of the following types:

$$\pi = \Phi(\rho + \kappa_s^I W_s + \Lambda_s^I S_s + \Psi_t^I T_t), \tag{1.10}$$

$$\lambda = \exp(\gamma + \kappa_s W_s + \Lambda_s S_s + \Psi_t T_t). \tag{1.11}$$

Here W_s is the vector of instrumental variables, and S_s, T_t are the station- and time-level controls as in Equations (1.10) and (1.11). The parameters $(\rho, \kappa_{st}^I, \Lambda_s^I, \Psi_t^I)$ and $(\gamma, \kappa_{st}, \Lambda_s, \Psi_t)$ correspond to the zero-inflation and main model of the first stage regressions.

Generalized Residuals: The usual residuals $(z_{st} - \hat{z_{st}})$ may be misleading in count regression models due to the presence of a non-linear link function. Specifically, although zero-inflated Poisson and negative-binomial models provide fitted values $\hat{z_{st}}$ (or estimated probabilities), there are no natural residuals that can be recovered. Following Wooldridge [2014], we use generalized residuals to construct the control functions. For the models estimated above, the generalized residuals are equal to the first derivative of the log-likelihood function with respect to the intercept, evaluated at the maximum likelihood estimates [Gourieroux et al., 1987, Vella, 1993]. We derive expressions for generalized residuals for both ZIP and ZINB models in Appendix 1.8.2. For the variables z_{ast}, the generalized residuals gr_{ast} are given by:

$$gr_{ast} = \begin{cases} \dfrac{(\pi - 1)\lambda(\dfrac{\theta}{\theta + \lambda})^{\theta+1}}{\pi + (1 - \pi)(\dfrac{\theta}{\theta + \lambda})^{\theta}} & \text{if } z_{ast} = 0, \\ z_{ast} - \dfrac{(z_{ast} + \theta)\lambda}{\lambda + \theta} & \text{if } z_{ast} > 0. \end{cases} \tag{1.12}$$

Step 2- Maximum likelihood Estimation: We use the generalized residuals obtained from Step 1 as control functions and substitute them into the utility specifications in Equations (1.2) and

(1.3). Finally, we estimate the model parameters by Maximum Likelihood Estimation using the following likelihood function:

$$P(d_1 = x_1, \ldots, d_M = x_M, d_{M+1} = y_1, \ldots, d_{M+O} = y_O) = \left[\prod_{q=1}^{Q} \frac{e^{\left(\frac{w_{iste_q}}{\sigma_q}\right) d_q}}{1 + e^{\left(\frac{w_{iste_q}}{\sigma_q}\right)}} \right]$$

$$\left[1 + \sum_{q=1}^{Q-1} \sum_{k=q+1}^{Q} (-1)^{d_q + d_k} \cdot \theta_{qk} \left\{ 1 - \frac{e^{\left(\frac{w_{iste_q}}{\sigma_q}\right) d_q}}{1 + e^{\left(\frac{w_{iste_q}}{\sigma_q}\right)}} \right\} \left\{ 1 - \frac{e^{\left(\frac{w_{iste_q}}{\sigma_k}\right) d_k}}{1 + e^{\left(\frac{w_{iste_q}}{\sigma_k}\right)}} \right\} \right],$$

where,

$$w_{iste_q} = \begin{cases} \alpha_x + \Theta_1^{'} Z_{st} + \Gamma_1^{'} S_s + \Delta_1^{'} R_i + \Omega_1^{'} T_t + \Upsilon_1^{'} Q_{ist} + \varsigma_1^{'} G_{st}, \text{ if } 1 \leq q \leq M (\text{H/R Decisions}), \\ \\ \alpha_y + \Theta_2^{'} Z_{st} + \Gamma_2^{'} S_s + \Delta_2^{'} R_i + \Omega_2^{'} T_t + \Upsilon_2^{'} Q_{ist} + \varsigma_2^{'} G_{st}, \text{ if } M+1 \leq q \leq M+O (\text{S/P Decisions}). \end{cases}$$

Here G_{st} is a vector of the control functions for the endogenous variables and ς is the vector of the associated coefficients.

1.6 Estimation Results and Model Validation

1.6.1 Estimation Results

This section presents our estimation results. From a statistical point of view, we find that all network-level congestion variables are statistically significant for both the H/R and the S/P decision. This suggests that congestion on a given link influences not only the hold probability for trains on that link but also those on neighboring links, indicating that section controllers account for network-wide congestion. Additionally, for most endogenous congestion variables, the endogeneity-corrected estimates are significantly different in magnitude than the naive estimates, reinforcing the importance of the control function approach in addressing endogeneity bias.

We now highlight key insights from our estimation, focusing on the H/R model. Qualitatively, our findings remain consistent for S/P decisions, hence we relegate that discussion of S/P estimates to Appendix 1.8.3. In the tables, a positive coefficient indicates a higher hold probability as the variable count increases. Throughout, we assume the focal train travels upstream ($\rightarrow$).

1. Passenger trains get priority for clearance and cause freight hold: Our analysis and results confirm that passenger trains receive high movement priority, thereby influencing freight trains' H/R decisions. Our estimates in Table 1.3 show that when a passenger train is present at the focal station, the probability of a freight train being held increases significantly, which is indicated by the magnitude and positive sign of the coefficients for *Current_Up_P* (θ=0.8044, p <0.01) variable. Similarly, suppose a passenger train is present upstream in the *Next* link. In that case, freight trains experience a higher hold probability, indicated by the estimate of *Next_Up_P* (θ=1.3611, p <0.01). Both these observations underscore the clear preference given to passenger trains for movement.

0 (base outcome-Release)	Naive Estimates	Endogeneity Corrected Estimates
Passenger Congestion Variables	Coefficient (Standard Error)	Coefficient (Standard Error)
Current_Up_P	0.6604***(0.0062)	0.8044***(0.0057)
Next_Up_P	1.2542***(0.0026)	1.3611***(0.0050)
Previous_Up_P	1.2501***(0.0087)	1.4299***(0.0099)
Next_Down_P	0.0165***(0.0004)	-0.0066***(0.0004)
Current_Down_P	-0.1299***(0.0002)	-0.1323***(0.0014)

Table 1.3: Passenger congestion & H/R decision : Naive and endogeneity corrected estimates. (*,**,*** indicates statistical significance at 10%,5%,1% level)

Another interesting insight is derived from the *Prev_Up_P* variable's coefficient estimate (θ=1.4299, p <0.01). The positive coefficient indicates *strategic idling* within the non-preemptive system (see Baron et al. [2014]).[25] It implies that section controllers are more inclined to hold a freight train if there is a passenger train in the *Previous* link of the network and is moving upstream, even though the *Next* link is free. This shows that the controllers cannot afford to delay the passenger train,

[25]In a non-preemptive priority queue, once a customer begins service, they can not be interrupted, regardless of the priority of customers arriving after them.

which might happen if the freight train gets released and occupies the *Next* link of the network, causing the passenger train to wait.

2. FCFS nature of freight train queue and the push effect: From the estimates in Table 1.4, we find that when a focal train is trailed by more freight trains, its hold probability increases. This is due to the 'First Come, First Serve'(FCFS) nature of the freight train queue, where trains queued up behind nudge those in front to clear the occupied capacity. We call it the *Push Effect*. The negative coefficients for the variables *Previous_Up_G* (θ=-0.2561, p <0.01), *Current_Up_G_Post* (θ=-0.4798, p <0.01) confirm the presence of this effect. Likewise, suppose additional freight trains are present ahead of the focal train, either at the focal station or in the following link. In that case, the focal train is held with a higher probability to let the preceding trains clear first, as indicated by variables *Current_Up_G_Pre* (θ=0.5611, p <0.01) and *Next_Up_G* (θ=0.6508, p <0.01).

0 (base outcome-Release)	Naive Estimates	Endogeneity Corrected Estimates
Freight Congestion Variables	Coefficient (Standard Error)	Coefficient (Standard Error)
Current_Up_G_Pre	0.2115***(0.0039)	0.5611***(0.0062)
Current_Up_G_Post	-0.1377***(0.0009)	-0.4798***(0.0004)
Next_Up_G	0.5975***(0.0024)	0.6508***(0.0016)
Previous_Up_G	-0.0264***(0.0006)	-0.2561***(0.0015)
Next_Down_G	-0.0867***(0.0008)	-0.2028***(0.0021)
Current_Down_G	-0.0001(0.0011)	-0.0747***(0.0010)

Table 1.4: Freight congestion & H/R decision : Naive and endogeneity corrected estimates.

3. Section controllers strategically hold trains at stations with higher capacity: Interestingly, from Table 1.5, we observe that the hold probability increases if the focal station is a junction (θ=0.4514, p <0.01) and decreases if the *Next* upstream station is a junction (θ=-0.5192, p <0.01). Similarly, the hold probability increases with an increasing number of tracks at a station (as indicated by the *Total_Tracks* variable's coefficient (θ=0.3281, p <0.01). These observations suggest that the controllers strategically hold a train at a bigger station with a larger capacity. We also

find that the hold probability is higher at a station if it is the last station in the administrative section (*Section_End variable*).[26] Finally, block sections that have three rail lines exhibit lower hold probability than the ones having two rail lines due to higher capacity.

0 (base outcome-Release)	Naive Estimates	Endogeneity Corrected Estimates
Station Variables	Coefficient (Standard Error)	Coefficient (Standard Error)
Junction	0.2574***(0.0008)	0.4514***(0.0015)
JunctionNS	-0.4351***(0.0019)	-0.5192***(0.0027)
Tracks_Total	0.2829***(0.0033)	0.3281***(0.0032)
Distance	0.0540***(0.0011)	0.0709***(0.0011)
Section_End	0.9143***(0.0007)	1.1156***(0.0037)
Line_Triple	-2.2661***(0.0037)	-2.3900***(0.0056)
Control: COA_Section	Yes	Yes

Table 1.5: Station characteristics & H/R decision: Naive and endogeneity corrected estimates.

1.6.2 Model Validation

Out-of-sample predictive accuracy

To test our model's accuracy in predicting the decisions of section controllers, we conduct out-of-sample testing. We split our data randomly into ten folds, iteratively select one fold as left-out, train our models for both decisions on the remaining folds, and measure the out-of-sample model performance on the left-out fold. We report the average accuracy after repeating this exercise five times. This is a binary classification exercise, and the accuracy of the classification is defined as:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}},$$

where TP (True Positives) are positive instances predicted correctly, TN (True Negatives) are negative instances predicted correctly, FP (False Positives) are negative instances predicted as positive,

[26] We posit that this is due to the operational friction between the section controllers of adjacent sections. This also implies that dividing the network into more administrative sections would slow down the trains more.

and FN (False Negatives) are positive instances predicted as negative.[27]

Class Imbalance: The H/R and S/P decisions data in our data exhibit class imbalance. Conditional on a train stopping, the average dwell time for the train is around 27 minutes. This dwell corresponds to 26 hold decisions and one release decision in the data. Hence, the hold decisions constitute a majority class with around 97.16% incidence.[28] Models trained on imbalanced data show a bias towards the majority class, leading to them predominantly predicting the majority class. Also, if we predict all instances in H/R decision data as hold, we would erroneously get 97.16% prediction accuracy, even though we would classify all release decisions incorrectly.

One way to improve the sensitivity towards the minority class and ensure a more balanced prediction performance is to shift away from the default probability threshold (0.5) for classifying a decision as hold or release. We vary the decision threshold and select the one where the F1 score is maximum.[29] The optimal decision thresholds for H/R and S/P decisions are 0.8 and 0.2, respectively. We find that our model does well in predicting the hold and stop decisions out-of-sample. Specifically, we are able to achieve a 96% accuracy for the H/R decisions and 74% for the S/P decisions, providing direct evidence that our model does well out-of-sample.

1.7 Conclusion

Railways, as a mode of transportation, are pivotal in the global freight industry due to their cost-effectiveness and environmental sustainability. However, in the past decade, there has been a continuous decline in the market share of railways in freight transportation, a trend evident in the Indian Railways as well. Given that India faces one of the highest logistic costs in the world

[27] We consider Hold and Stop decisions as positive instances and Release and Passthrough decisions as negative instances.

[28] Similarly, 83% of data for the S/P decisions correspond to passthrough decisions.

[29] F1 score is the harmonic mean of precision and recall, where precision is defined as the ratio of correctly predicted positive observations to the total predicted positive observations, and recall is defined as the ratio of correctly predicted positive observations to the total actual positive observations. F1 score combines precision and recall into a single metric, balancing false positives and false negatives.

(14% of GDP) and railways are the most cost-effective freight transportation mode, this decline is particularly problematic for India's rapidly growing logistics sector. It is also a concern for Indian Railways, as most of its revenue is generated from freight transport. One of the important reasons for this decline is the extended delays faced by freight trains in Indian Railways, where a controller makes movement and scheduling decisions for freight trains in real-time. To address this, the government is investing significantly in capacity building, and hence, it is crucial to understand ex-ante how controllers make decisions for freight train movements.

Using a detailed high-frequency dataset capturing congestion changes in the Indian Railway network, we estimate two discrete choice models to uncover the key parameters underlying the section controllers' decisions. Using the control function approach, we account for the endogeneity of the network congestion variables. The estimated parameters confirm that passenger trains are prioritized over freight trains, leading to extended delay times. The estimates also highlight a push effect in the freight train queue and the strategic behavior of section controllers in holding trains at larger stations. Our study is one of the first to empirically examine train flows and investigate the decision-making of railway section controllers. Railway organizations worldwide cater to diverse train types with varying speeds and priorities. While our study focuses on the Indian context, the manual control mechanisms to manage such a heterogeneous mix of trains discussed in the chapter are used in railway networks throughout the world. Hence, the methodology can generally be applied to study controllers' decisions in railway networks across the globe.

1.8 Appendix

1.8.1 Data and Model Appendix

We exclude data for the stations located at the divisional boundary because we cannot capture congestion information of the adjacent division. We also exclude data from the NEWC-FOCM COA section, as it contains a high percentage of missing timestamps. We exclude the trains from our analysis that don't have arrival and departure timestamps (<1%). In the remaining data, 5.4% of observations have missing timestamps. Among these, for observations where only one station in a sequence has a missing arrival and/or departure timestamp for a train, we impute these timestamps using the departure time of the previous station, the arrival time of the next station, and travel time information between stations. We drop the remaining 1.3% of the observations from our analysis. As discussed in the Section 1.5, we use zero-inflated models for the congestion variables. In Tables 1.6 and 1.7, we report the distribution of congestion variables, showing a high percentage of zeros, justifying our modeling choice.

Variable	Mean	Std. Dev.	% Zeros
Current_Up_P	0.15	0.46	88.23%
Current_Up_G_Pre	1.09	2.40	68.59%
Current_Up_G_Post	0.43	1.13	79.69%
Next_Up_P	0.08	0.28	91.41%
Next_Up_G	0.50	0.78	63.27%
Previous_Up_P	0.15	0.39	85.46%
Previous_Up_G	0.71	0.78	46.24%
Next_Down_P	0.12	0.34	88.39%
Next_Down_G	0.57	0.77	56.56%
Current_Down_P	0.12	0.39	89.60%
Current_Down_G	0.63	1.13	65.69%

Table 1.6: HR: Congestion Variables

Variable	Mean	Std. Dev.	% Zeros
Current_Up_P	0.02	0.15	97.88%
Current_Up_G	0.28	0.51	74.85%
Next_Up_P	0.05	0.25	92.68%
Next_Up_G	0.68	1.11	52.59%
Previous_Up_P	0.03	0.17	97.12%
Previous_Up_G	0.44	0.76	66.59%
Next_Down_P	0.08	0.29	92.24%
Next_Down_G	0.61	0.82	54.72%
Current_Down_P	0.06	0.25	93.49%
Current_Down_G	0.37	0.57	67.14%

Table 1.7: SP: Congestion Variables.

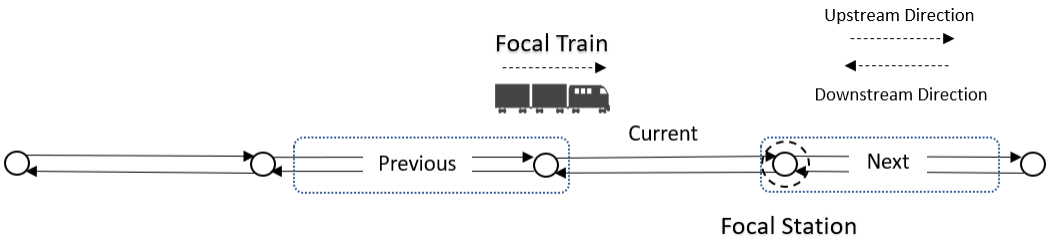

Figure 1.9: Congestion variable: S/P decision.

Congestion Variables

For the H/R model, we decompose the congestion variable Current_Up_G into two components: Current_Up_G_Pre and Current_Up_G_Post. These components represent the count of freight trains traveling upstream that arrived at the focal station before and after the arrival of the focal train. The reason for this decomposition stems from the 'first come, first served' nature of the freight train queue. Specifically, an increase in trains that arrived at the focal station before the focal train (ahead in the queue) is expected to increase the hold probability, whereas an increase in trains that arrived after the focal train (behind in the queue) is expected to decrease the hold probability.

As the S/P decision for a train at the upcoming station is made when the train departs from the current station, there is a slight variation in the definition of the congestion variable for the S/P decision. We use the term 'Current' for the part of the network where the train is when making the decision. The remaining two attributes for congestion characterization, $d \in \{\text{Up, Down}\}$ and $f \in \{\text{P,G}\}$, are defined similarly to the H/R decision (Please see Figure 1.9.) All the results for the S/P model should be interpreted accordingly.

Moran's I: Evidence of Residual Spatial Correlation Among Decisions

We first fit the logit model separately for both HR and SP decisions and recover the residuals for the independent decisions. To capture spatial relationships between decisions, we construct an adjacency matrix W_{ij}, elements of which are determined by the type of decision combination

they represent (refer to decision type combination in Table 1.8). Using this matrix and the model residuals, we compute Moran's I statistic for each decision type combination. This helps assess whether spatial correlation remains in the errors, indicating unaccounted spatial dependence after fitting the independent logit models. The Moran's I statistic is given by:

$$I = \frac{N}{\sum_i \sum_j W_{ij}} \cdot \frac{\sum_i \sum_j W_{ij}(e_i - \bar{e})(e_j - \bar{e})}{\sum_i (e_i - \bar{e})^2}$$

Here$N$ is the number of observations,e_i and e_j are the residuals for observations i and j, $\bar{e}$ is the mean of the residuals, and W_{ij} denotes the spatial weight (adjacency) between observations i and j.

Decision Type Combination	Moran's I	p-Value
HR-HR Same Direction	0.04	0.003
HR-HR Opposite Direction	0.12	0.050
HR-SP Same Direction	0.07	0.047
HR-SP Opposite Direction	-0.09	0.006
SP-SP Same Direction	0.11	0.101
SP-SP Opposite Direction	0.02	0.382

Table 1.8: Moran's I for spatial correlation in model residuals: 5 Days Sample. (Decisions made at neighboring stations)

S/P decision: Value of *k*

As discussed in the Section 1.4, we do not observe the exact time of the S/P decision in our data; we only observe whether the train was stopped. We use a data-driven approach to obtain the optimal $k \in [0, 1]$, representing the fraction of the inter-station distance covered before the S/P decision. We study the S/P decisions for a sample part of the network (GAYA-GRRU) for four different values of k. Though we do not find much difference in the AIC, BIC, and accuracy values, we find that the AIC and BIC are minimized for $k = 0$.

k	0	0.25	0.50	0.75
AIC	22682	22776	22959	22991
BIC	23182	23277	23460	23492
Accuracy	82%	81%	81%	81%

Table 1.9: S/P decision: Value of *k* (GAYA-GRRU subsection: 20,928 observations).

1.8.2 Estimation Appendix

We derive the generalized residuals for the Zero-inflated Negative Binomial model in this section. The probability mass function of the Zero-Inflated Negative Binomial distribution is:

$$P(Y=z)=\begin{cases}\pi+(1-\pi)\left(\dfrac{\theta}{\theta+\lambda}\right)^{\theta} & \text{if } z=0,\\ (1-\pi)\dfrac{\Gamma(z+\theta)}{\Gamma(\theta)\Gamma(z+1)}\left(\dfrac{\theta}{\theta+\lambda}\right)^{\theta}\left(\dfrac{\mu}{\theta+\lambda}\right)^{z} & \text{if } z>0.\end{cases}$$

The log-likelihood for an observation $Y=z$ is given by:

$$\log(L)=[I(z=0)\log\left(\pi+(1-\pi)\left(\frac{\theta}{\theta+\lambda}\right)^{\theta}\right)$$
$$+I(z>0)\log\left((1-\pi)\frac{\Gamma(z+\theta)}{\Gamma(\theta)\Gamma(z+1)}\left(\frac{\theta}{\theta+\lambda}\right)^{\theta}\left(\frac{\lambda}{\theta+\lambda}\right)^{z}\right)]$$

To obtain generalized residuals, we consider two cases separately:

1. The log-likelihood for an observation $z=0$ is:

$$\log(L)=\log\left(\pi+(1-\pi)\left(\frac{\theta}{\theta+\lambda}\right)^{\theta}\right)$$

$$gr=\frac{\partial\log(L)}{\partial\beta_0}=\frac{(\pi-1)\lambda(\frac{\theta}{\theta+\lambda})^{\theta+1}}{\pi+(1-\pi)(\frac{\theta}{\theta+\lambda})^{\theta}}$$

2. The log-likelihood for an observation $z > 0$ is:

$$\log(L) = \log\left((1-\pi)\frac{\Gamma(z+\theta)}{\Gamma(\theta)\Gamma(z+1)}\left(\frac{\theta}{\theta+\lambda}\right)^{\theta}\left(\frac{\lambda_i}{\theta+\lambda}\right)^{z}\right)$$

$$gr = \frac{\partial \log(L)}{\partial \beta_0} = z - \frac{(z+\theta)\lambda}{\lambda+\theta}$$

Unobservables in the Model

One of the factors we do not observe in the data is the crew change schedule. The crew usually works around six hours in a single stretch. The crew changes are necessary for managing crew fatigue and are determined by labor laws and regulations. It typically happens when the drivers/crew reach near the permissible working hours,[30] or if a train is entering a new railway administrative division. More importantly, the crew change points for different freight trains are at different designated stations. These crew change unobservables provide a positive shock to the congestion at the *focal* and nearby stations. The shock also increases the hold probability of a train, and hence, omitting this variable may lead to underestimating the congestion variables' coefficients. These crew changes for freight trains do not impact passenger congestion variables, as passenger trains get prioritized over freight trains. Moreover, since passenger train schedules are predetermined, we only model the freight train congestion variables as endogenous.

Another unobservable is the track maintenance schedule. The railway allocates daily time slots to perform maintenance work on tracks for different routes.[31] Since maintenance provides a positive shock to stations' congestion and increases the hold probability of a train, its exclusion may also lead to an underestimation of the congestion variables' coefficients.[32]

[30] https://tinyurl.com/yc2r5cey

[31] https://tinyurl.com/4mucxdry

[32] This maintenance is usually scheduled during low levels of passenger traffic and hence, does not impact passenger train related congestion variables.

Testing Strength of the Instruments

Weak instrumental variables can pose a challenge in estimation because they can lead to biased estimates of causal relationships [Frazier et al., 2021]. To test the strength of our instrumental variables, we perform the *analysis of deviance* tests for all our endogenous variables. The analysis of deviance tests are used to test the significance of generalized linear models, and the test procedure is analogous to the general linear F test for multiple linear regression.[33]

Construction of the test statistic: let $L_{\mathcal{M}}$ denote the maximum achievable likelihood under a model $\mathcal{M}$ and $L_{\mathcal{S}}$ denote the likelihood under the saturated model (one that achieves a perfect fit). Then the deviance $D_{\mathcal{M}}$ of the model $\mathcal{M}$ is defined as:

$$D_{\mathcal{M}} = -2 \log \left(\frac{L_{\mathcal{M}}}{L_{\mathcal{S}}} \right) = -2 \left(\log L_{\mathcal{M}} - \log L_{\mathcal{S}} \right).$$

While testing the alternative hypothesis that $\mathcal{M}_1$ is a better fit than a model $\mathcal{M}_0$, the test statistic is the difference of deviances:

$$D = D_{\mathcal{M}_0} - D_{\mathcal{M}_1} = -2 \left(\log L_{\mathcal{M}_0} - \log L_{\mathcal{M}_1} \right).$$

The deviance statistics follow the χ^2 distribution with $p_0 - p_1$ degrees of freedom, where p_0 and p_1 are the number of parameters in $\mathcal{M}_0$ and $\mathcal{M}_1$ models. To test the strength of instruments, we consider a model with three constants (one each for the main and the inflation models and one for dispersion) as the $\mathcal{M}_0$ model and a model with all instruments and control variables as the $\mathcal{M}_1$ model. The null and the alternate hypotheses of the test are defined below:

- H_0: A model without IVs and control variables provides a significantly better fit to data.

[33] In an OLS first-stage regression, the F-statistic assesses the joint significance of the instruments in explaining the variation in the endogenous variable. Instead of relying on the F-statistic, hypothesis testing in GLM is performed using likelihood ratio tests or analysis of deviance tests. Please see Fox, John. Applied regression analysis and generalized linear models. Sage Publications, 2015.

- H_1: A model with IVs and control variables provides a significantly better fit to data.

In other words, the null hypothesis states that the instrumental variables in the first stage are not jointly significant. The test statistics and the corresponding p-values for all the endogenous variables in H/R model are given in Tables 1.10 and 1.11. We find that the deviance statistics for all the variables are greater than the critical values. Thus, our instruments are relevant, as the null hypothesis is rejected for all endogenous variables in our model.

Variable	Deviance Statistics
Current_Up_G_Pre	4755974***
Current_Up_G_Post	2598702***
Next_U_G	2202344***
Next_Down_G	625090***
Previous_Up_G	304790***
Current_Down_G	2837752***
df	28
Critical Value ($\alpha = 0.01$)	48.27

Table 1.10: Deviance Statistics for Instrument Strength: H/R Model

Variable	Deviance Statistics
Current_Up_G	55161***
Next_Up_G	252500***
Previous_Up_G	70522***
Next_Down_G	124884***
Current_Down_G	38522***
df	28
Critical Value ($\alpha = 0.01$)	48.27

Table 1.11: Deviance statistics for Instrument Strength: S/P model.

1.8.3 Results Appendix

Estimation results for S/P decision are shown in Table 1.13. In line with the results of H/R estimation, we observe that Passenger trains get priority for clearance and cause freight stops. This is clear from the estimates of Current_Up_P (0.6094***) and Previous_Up_P (0.5211***) variables. The estimates Current_Up_G (0.9713***) and Next_Up_G (0.8532***) show the FCFS nature of

	Naive Estimates	Endogeneity Corrected Estimates
0 (base outcome-Release & Passthrough)	Coefficient (Standard Error)	Coefficient (Standard Error)
HR Congestion Variables		
Current_Up_P	0.6604***(0.0062)	0.8044***(0.0057)
Current_Up_G_Pre	0.2115***(0.0039)	0.5611***(0.0062)
Current_Up_G_Post	-0.1377***(0.0009)	-0.4798***(0.0004)
Next_Up_P	1.2542***(0.0026)	1.3611***(0.005)
Next_Up_G	0.5975***(0.0024)	0.6508***(0.0016)
Previous_Up_P	1.2501***(0.0087)	1.4299***(0.0099)
Previous_Up_G	-0.0264***(0.0006)	-0.2561***(0.0015)
Next_Down_P	0.0165***(0.0004)	-0.0066***(0.0004)
Next_Down_G	-0.0867***(0.0008)	-0.2028***(0.0021)
Current_Down_P	-0.1299***(0.0002)	-0.1323***(0.0014)
Current_Down_G	-0.0001(0.0011)	-0.0747***(0.001)
HR Station Variables		
Junction	0.2574***(0.0008)	0.4514***(0.0015)
JunctionNS	-0.4351***(0.0019)	-0.5192***(0.0027)
Tracks_Total	0.2829***(0.0033)	0.3281***(0.0032)
Distance	0.054***(0.0011)	0.0709***(0.0011)
Section_End	0.9143***(0.0007)	1.1156***(0.0037)
Line_Triple	-2.2661***(0.0037)	-2.3900***(0.0056)
Control: COA_Section	Yes	Yes
HR Train Variables		
Loaded_Flag	0.0753***(0.0016)	0.1198***(0.0012)
Control: Origin	Yes	Yes
Control: Destination	Yes	Yes
Control: Wagon_Type	Yes	Yes
Control: Loco_Type	Yes	Yes
HR Train & Station Variables		
Delayed	0.4577***(0.0003)	0.3653***(0.0018)
Log(1+Dist_To_Dest)	-0.0435***(0.0013)	-0.0627***(0.0017)
HR Control Functions		
Control_Function_Current_Up_G_Pre		-0.4372***(0.0058)
Control_Function_Current_Up_G_Post		0.4035***(0.0023)
Control_Function_Next_Up_G		-0.0191***(0.0019)
Control_Function_Previous_Up_G		0.2380***(0.0014)
Control_Function_Next_Down_G		0.1079***(0.0018)
Control_Function_Current_Down_G		0.0830***(0.0009)
HR Time Variables		
Control: Hour_Slot and Day_Week	Yes	Yes

Table 1.12: Second Stage Results (Part 1).

	Naive Estimates	Endogeneity Corrected Estimates
0 (base outcome-Release & Passthrough)	Coefficient (Standard Error)	Coefficient (Standard Error)
SP Congestion Variables		
Current_Up_P	0.5792***(0.0007)	0.6094***(0.0023)
Current_Up_G	0.3404***(0.0014)	0.9173***(0.0023)
Previous_Up_P	0.4687***(0.0011)	0.5211***(0.0019)
Previous_Up_G	0.2210***(0.0006)	-0.4896***(0.0015)
Next_Up_P	0.4430***(0.0013)	0.4708***(0.0015)
Next_Up_G	0.7408***(0.0023)	0.8532***(0.0041)
Next_Down_P	0.1509***(0.0013)	0.1255***(0.0001)
Next_Down_G	0.0563***(0.0006)	0.6825***(0.0015)
Current_Down_P	-0.0342***(0.0002)	-0.0190***(0.0001)
Current_Down_G	-0.0917***(0.0003)	-1.0274***(0.0015)
SP Station Variables		
Junction	-0.2943***(0.0004)	-0.0606***(0.0014)
JunctionNS	0.4254***(0.0006)	0.4299***(0.0013)
Distance	0.1611***(0.0006)	0.1388***(0.0007)
Tracks_Total	0.1263***(0.0009)	0.0199***(0.0011)
Section_End	1.0932***(0.0028)	1.1534***(0.0009)
Line_Triple	-1.8576***(0.0008)	-1.6812***(0.0006)
Control: COA_Section	Yes	Yes
SP Train Variables		
Loaded_Flag	-0.0586***(0.0002)	0.0244***(0.0007)
Control: Origin	Yes	Yes
Control: Destination	Yes	Yes
Control: Wagon_Type	Yes	Yes
Control: Loco_Type	Yes	Yes
SP Train & Station Variables		
Delayed	-0.0696***(0.0003)	0.0196***(0.0003)
Log(1+Dist_To_Dest)	-0.1534***(0.0005)	-0.1012***(0.0008)
SP Control Functions		
Control_Function_Current_Up_G		-0.5940***(0.0021)
Control_Function_Previous_Up_G		0.7981***(0.0015)
Control_Function_Next_Up_G		-0.1368***(0.0035)
Control_Function_Next_Down_G		-0.6437***(0.0015)
Control_Function_Current_Down_G		0.9601***(0.0014)
SP Time Variables		
Control: Hour_Slot and Day_Week	Yes	Yes

Table 1.13: Second Stage Results (Part 2).

	Naive Estimates	Endogeneity Corrected Estimates
0 (base outcome-Release & Passthrough)	Coefficient (Standard Error)	Coefficient (Standard Error)
Copula Variables		
HR-HR Same Direction	-0.0159***(0.0003)	-0.0261***(0.0004)
HR-HR Opposite Direction	-0.0299***(0.0007)	-0.0490***(0.0009)
HR-SP Same Direction	0.0008***(0.0001)	0.0013***(0.0001)
HR-SP Opposite Direction	0.0007***(0.0001)	0.0014***(0.0001)
SP-SP Same Direction	-0.0072***(0.0001)	-0.0119***(0.0002)
SP-SP Opposite Direction	-0.0067***(0.0001)	-0.0110***(0.0002)
Heteroscedasticity Variables		
Tracks_Total (HR)	0.0277***(0.0006)	0.0403***(0.0005)
Tracks_Total (SP)	0.0160***(0.0005)	0.0164***(0.0004)
Number of Observations	65,35,878	65,35,878

Table 1.14: Second Stage Results (Part 3).

the freight train queue. Similar to the results of HR estimation, we observe that section controllers strategically stop trains at stations with higher capacity, as shown by the coefficient of Tracks_Total variable. Contrary to the results of H/R estimation, we observe that the section controller prefers the freight train to passthrough a junction station. We also observe a higher stop probability at the section end stations.

First Stage Results

We report the results of first-stage zero-inflated models for endogenous freight congestion variables in Tables 1.17 and 1.18. The first part of both tables presents the estimates for the negative binomial models. The latter half of the tables labeled as 'Inflate' provides estimates for zero inflation, modeled using a probit link function.

Copula and Heteroscedasticity Estimates

We find that HR decisions made at neighboring stations are negatively correlated (δ_1=-0.0261, $p < 0.01$, δ_2=-0.0490, $p < 0.01$). This implies that the probability of a hold decision at the focal station decreases when a hold decision is made at the neighboring station and vice versa. We posit that Hold-Hold (1-1) is not preferred because, given everything else remains the same, the

controller will prefer to release at least one train to relieve congestion at nearby stations. Additionally, the controller is unlikely to release both trains simultaneously (Release-Release (0-0)) due to capacity constraints on shared tracks. Thus, the outcomes Release-Hold (0-1) and Hold-Release (1-0) become more probable as these balance congestion relief and capacity limitations across the two neighboring stations. We observe the same pattern for SP decisions as well.

We also accommodate heteroscedasticity in the variance of the error term across decisions caused by variations in the number of tracks at the focal station. Positive and significant coefficients for both HR (Ξ_1=0.0403, p <0.01) and SP (Ξ_2=0.0164, p <0.01) decisions imply that decisions made at the station with a low number of tracks have less variability. We posit that this is because stations with fewer tracks provide limited opportunities for section controllers to hold trains rather than release them, especially in congested conditions, thereby reducing the variability in decision-making.

Section-Level Heterogeneity

Our data do not contain unique identifiers for individual dispatchers (section controllers), so we are unable to directly account for variation at the section-controller level in our estimation. However, we capture some of this heterogeneity through the inclusion of section-level fixed effects in our model. Since dispatchers typically work in fixed sections without rotating across them (albeit in different shifts), the section fixed effects absorb part of the persistent differences in dispatcher decision making across sections. Tables 1.15 and 1.16 present these fixed effects, with the ADSR–MPO section normalized as the baseline.

0 (base-Release)	Naive Estimates	Endogeneity Corrected Estimates
Section Fixed Effect	Coefficient (Standard Error)	Coefficient (Standard Error)
Section_FOCM-ADSR	1.2861***(0.0016)	1.4229***(0.0030)
Section_SEB-GHD	-0.5590***(0.0032)	-0.5764***(0.0024)

Table 1.15: Section Fixed Effects (H/R decision) : Naive and endogeneity corrected estimates.

We observe substantial variation across sections. For example, in Table 1.15, the fixed effect for

0 (base-Passthrough)	Naive Estimates	Endogeneity Corrected Estimates
Section Fixed Effect	Coefficient (Standard Error)	Coefficient (Standard Error)
Section_FOCM-ADSR	1.1363***(0.0009)	0.9875***(0.0012)
Section_SEB-GHD	-0.2752***(0.0009)	-0.2954***(0.0011)

Table 1.16: Section Fixed Effects (S/P decision) : Naive and endogeneity corrected estimates.

the FOCM–ADSR section is significantly positive, indicating a higher baseline propensity to hold trains in this section compared to the baseline section. The coefficient for the SEB–GHD section is significantly negative, suggesting a systematically higher likelihood of releasing trains. These differences suggest that operational decision-making, potentially reflecting dispatcher discretion, traffic load, infrastructure, or local norms, varies meaningfully across sections. While we cannot isolate individual dispatcher decisions, our model provides estimates for the section-level variation that planners can use.

	Current_Up_G_Pre	Current_Up_G_Post	Next_Up_G	Previous_Up_G	Next_Down_G	Current_Down_G
Main						
Lag2_Junction	0.217***	0.802***	-0.027***	-0.185***	-0.261***	-1.247***
	(0.011)	(0.016)	(0.005)	(0.004)	(0.004)	(0.008)
Lag2_Tracks_Total	-0.099***	-0.205***	-0.049***	-0.011***	0.01***	-0.03***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Junction	-0.533***	-3.697***	11.185	-0.964**	-27.641***	-1.829***
	(0.055)	(0.033)	(25.373)	(0.377)	(3.572)	(0.149)
Junction_NS	1.485***	-1.089***	-5.928***	-3.478***	-16.788***	1.112***
	(0.014)	(0.017)	(0.536)	(0.67)	(3.091)	(0.01)
Tracks_Total	-0.714***	0.487***	0.656***	1.289***	2.521***	-0.525***
	(0.009)	(0.005)	(0.004)	(0.078)	(0.286)	(0.003)
Distance	0.131***	0.143***	-0.95***	1.939***	-7.744***	-0.37***
	(0.002)	(0.003)	(0.004)	(0.027)	(0.512)	(0.002)
Section_End	2.067***	1.453***	14.352***	-2.166	-24.697	-1.735***
	(0.029)	(0.017)	(4.621)	(2.847)	(89.363)	(0.017)
Line_Triple	4.74***	-4.774***	-33.127	-7.854***	11.306***	6.483***
	(0.071)	(0.04)	(25.972)	(0.573)	(3.462)	(0.149)
Constant	-3.195***	-3.718***	0.309***	-0.532***	-1.104***	0.742***
	(0.012)	(0.021)	(0.004)	(0.004)	(0.004)	(0.008)
Inflate						
Lag2_Junction	-0.041	-0.571***	-14.693	1.855***	3.556	-4.309***
	(0.056)	(0.022)	(25.373)	(0.048)	(4.92)	(0.171)
Lag2_Tracks_Total	-0.214***	-0.2***	-1.213***	0.304***	-1.881***	-0.005**
	(0.004)	(0.003)	(0.005)	(0.007)	(0.295)	(0.002)
Junction	-0.533***	-3.697***	11.185	-0.964**	-27.641***	-1.829***
	(0.055)	(0.033)	(25.373)	(0.377)	(3.572)	(0.149)
Junction_NS	1.485***	-1.089***	-5.928***	-3.478***	-16.788***	1.112***
	(0.014)	(0.017)	(0.536)	(0.67)	(3.091)	(0.01)
Tracks_Total	-0.714***	0.487***	0.656***	1.289***	2.521***	-0.525***
	(0.009)	(0.005)	(0.004)	(0.078)	(0.286)	(0.003)
Distance	0.131***	0.143***	-0.95***	1.939***	-7.744***	-0.37***
	(0.002)	(0.003)	(0.004)	(0.027)	(0.512)	(0.002)
Section_End	2.067***	1.453***	14.352***	-2.166	-24.697	-1.735***
	(0.029)	(0.017)	(4.621)	(2.847)	(89.363)	(0.017)
Line_Triple	4.74***	-4.774***	-33.127	-7.854***	11.306***	6.483***
	(0.071)	(0.04)	(25.972)	(0.573)	(3.462)	(0.149)
Constant	3.475***	-0.567***	6.341***	-23.045***	25.972***	3.91***
	(0.041)	(0.036)	(0.027)	(0.458)	(2.562)	(0.016)
LN (Dispersion)	-2.189***	-1.92***	-18.51***	-15.445***	-16.496	-13.328***
	(0.004)	(0.009)	(3.065)	(4.249)	(3.791)	(3.394)
Control: Coa_Section	Yes	Yes	Yes	Yes	Yes	Yes
Control: Hour_Slot	Yes	Yes	Yes	Yes	Yes	Yes
Control: Weekday	Yes	Yes	Yes	Yes	Yes	Yes

Table 1.17: First stage results: Hold/Release decision.
(Standard Errors in parentheses and *,**,*** indicates statistical significance at 10%,5%,1% level)

	Current_Up_G	Previous_Up_G	Next_Up_G	Next_Down_G	Current_Down_G
Main					
Lag2_Junction	0.308***	-0.29***	0.213***	-0.129***	-0.719***
	(0.012)	(0.01)	(0.01)	(0.016)	(0.059)
Lag2_Tracks_Total	-0.042***	0.028***	-0.043***	-0.032***	0.047***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)
Junction	0.194***	-0.139***	-0.545***	-0.605***	0.473***
	(0.011)	(0.011)	(0.007)	(0.008)	(0.059)
Junction_NS	0.031***	-0.032***	0.46***	0.2***	-0.15***
	(0.008)	(0.006)	(0.004)	(0.013)	(0.007)
Tracks_Total	0.037***	-0.007***	0.212***	0.13***	-0.015***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)
Distance	0.03***	-0.054***	0.001	0.04***	-0.012***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Section_End	0.183***	0.081***	-0.068***	-0.226***	-0.068***
	(0.009)	(0.009)	(0.005)	(0.006)	(0.009)
Line_Triple	0.151***	1.026***	-0.532***	0.434***	0.634***
	(0.014)	(0.009)	(0.006)	(0.008)	(0.01)
Constant	-2.159***	-0.678***	-1.445***	-1.627***	-1.508***
	(0.016)	(0.009)	(0.009)	(0.009)	(0.01)
Inflate					
Lag2_Junction	-0.413***	-0.657***	4.227***	-0.539	-14.878***
	(0.09)	(0.032)	(0.252)	(1.18)	(0.3)
Lag2_Tracks_Total	0.069***	0.057***	-0.323***	0.511***	1.425***
	(0.009)	(0.003)	(0.019)	(0.014)	(0.03)
Junction	-2.18	-3.467***	-8.385***	1.502	11.96***
	(11.356)	(0.046)	(0.236)	(143.806)	(0.243)
Junction_NS	0.599***	-4.714	1.014***	-2.409**	-2.577***
	(0.048)	(3.292)	(0.046)	(1.214)	(0.06)
Tracks_Total	-0.651***	0.381***	0.654***	-2.479***	-0.496***
	(0.018)	(0.003)	(0.022)	(0.071)	(0.018)
Distance	0.096***	-0.117***	-0.292***	-0.116***	-0.959***
	(0.007)	(0.003)	(0.016)	(0.012)	(0.015)
Section_End	2.177*	-1.221***	-0.522***	-2.616	2.09***
	(1.138)	(0.019)	(0.028)	(4.305)	(0.078)
Line_Triple	-4.315***	5.18***	-0.946	0.057	-9.412***
	(0.541)	(0.429)	(7.377)	(0.303)	(0.19)
Constant	1.537***	-5.438***	-4.55***	3.593***	0.539***
	(0.079)	(0.429)	(0.147)	(0.123)	(0.045)
LN (Dispersion)	-15.941	-2.541***	-10.581***	-16.317	-15.684
	(27.329)	(0.042)	(1.883)	(7.864)	(26.422)
Control: Coa_Section	Yes	Yes	Yes	Yes	Yes
Control: Hour_Slot	Yes	Yes	Yes	Yes	Yes
Control: Day_Week	Yes	Yes	Yes	Yes	Yes

Table 1.18: First stage results: Stop/Passthrough decision.
(Standard Errors in parentheses and *,**,*** indicates statistical significance at 10%,5%,1% level)

Chapter 2

Improving Freight Train Speed in Indian Railways: Counterfactual Analysis

2.1 Introduction

Besides understanding how section controllers make decisions for freight trains, one important objective behind formulating and estimating the structural model is to evaluate policies or interventions for reducing freight delays. In this chapter, we use the estimated parameters from the S/P and H/R models from Chapter 1 to conduct counterfactual analyses and provide prescriptive guidance to Indian Railways for designing such interventions. Our counterfactuals simulate the impact of these interventions on section controllers' S/P and H/R decisions for all trains in the network and, therefore, on their delays and resultant speeds.

We first evaluate the impact of operating Freight Only Corridors (FOCs), which are railway corridors exclusively reserved for freight transportation. The Indian government plans to construct similar dedicated corridors to increase freight speeds and improve throughput. Our study analyzes the FOC's delay reduction and speed improvement benefits (and their sources), much before the entire corridor is constructed. Since such capacity intervention requires a substantial upfront investment (the government plans to spend over 10 billion USD on dedicated corridors), we also simulate alternative policies that do not require such huge monetary investments. Specifically,

we simulate two policies that rely on existing network capacity: (a) threshold-based releases for freight trains dwelling longer than a specified time limit and (b) freight capacity consolidation using *vertically stacked* trains.

For the FOC counterfactual, we simulate the policy by removing the passenger trains in the network, computing the resulting stop and hold probabilities across decisions, and simulating the new freight train flows. For the Threshold Release counterfactual, we set a train's hold/release decision to 'release' once it crosses a dwell time limit and simulate the counterfactual train flow. In the consolidation counterfactual, we study consolidating existing freight capacity into fewer vertically stacked trains. For example, instead of operating five trains with 100 wagons each with origin station o and destination station d, railways could operate four trains with 125 wagons each by vertically stacking 25 wagons on these trains.

The results show that freight trains wait for much less (37% reduction), stop less frequently (31% reduction), and travel faster through the network (16% improvement) due to the FOCs. The speed increase comes from lesser dwell and lesser acceleration and deceleration losses. We also find that dwell thresholds of 60 and 45 minutes result in substantial speed improvements ranging from 13% to 17%. These releases account for less than 2% of the controllers' decisions. Also, the railways can flexibly release freight trains in a suitable time window so that the releases do not interfere with passenger train schedules. For the consolidation counterfactual, we find that just a low (20%) to moderate (30%) level of consolidation leads to noticeable speed improvements in the range of 4% to 8%. Overall, our non-capacity-based interventions are substantially cheaper, can be implemented easily, and can achieve similar speed improvements as those compared to capacity interventions like the FOC.

The remainder of this chapter is organized as follows: In Section 2.3, we discuss the FOC counterfactual; in Section 2.3, we discuss the Threshold Release counterfactual; and in Section 2.4, we discuss the Consolidation counterfactual. We provide a cost-benefit analysis of these three policies in Section 2.4.1 and conclude in Section 2.6.

Measures:

For a given counterfactual policy $\mathcal{P}$, we are interested in Average Dwell Time (Avg_Dwell), Average Stoppages (Avg_Stoppages), and the Average Speed (Avg_Speed) for the freight trains. Let us denote the set of all trains by I and the set of stations in the freight train i's routes by S_i. Let the time of departure from the first station in the set S_i be $t_{id}^{s_i^f}$, the time of arrival at the last station be $t_{ia}^{s_i^l}$ and distance in km between the first and last station for train i be D_i. Here s_i^f and s_i^l denote the first and last entries in the set S_i. Note that $t_{ia}^{s_i^l}$ is not fixed and depends on the simulated H/R and S/P decisions under any policy. While examining the policies, we randomly sample 30 days from our data, simulate all decisions for the trains on sampled days, and compute these measures by averaging across all stations for the sample day under consideration. We then compare the distribution of these measures across the policies. Below, we explain these measures:

- Avg_Dwell: The average dwell time, measured in minutes, for freight trains over all stations is given by: $\frac{\sum_i \sum_{s \in S_i} \sum_t x_{ist}}{\sum_i |S_i|}$, where $|S_i|$ represents the cardinality of set S_i.

- Avg_Stoppages: The number of stoppages for the freight train i is $\sum_{s \in S_i} y_{ist}$. Thus, the average number of stoppages, measured in %, for all freight trains is given by : $\frac{\sum_i \sum_{s \in S_i} y_{ist}}{\sum_i |S_i|}$ and the average passthrough percentage is (1-Avg_Stoppages).

- Avg_Speed: The travel time for freight train i from the first station to last station in its route is: $(t_{ia}^{s_i^l} - t_{id}^{s_i^f})$, and hence the average speed, measured in kmph is: $D_i \cdot 60/(t_{ia}^{s_i^l} - t_{id}^{s_i^f})$. The average speed for all trains is given by : $\frac{\sum_i (D_i \cdot 60/(t_{ia}^{s_i^l} - t_{id}^{s_i^f}))}{|I|}$.

We consider two scenarios: first, a baseline case where the controller makes decisions in the status-quo setting, and second, where the controller makes decisions under the given counterfactual policy $\mathcal{P}$. We compare the key measures of interest in both scenarios to assess the impact of counterfactual policies. We denote the calculated measures under the given counterfactual policy with a

superscript; for example, Avg_Dwell^{FOC} indicates the average dwell time under the Freight Only Corridor policy.

2.2 Capacity Counterfactual: Freight-Only Corridors

The slow pace of freight trains in India has led to low freight capacity and high logistics costs. The speed of freight is notoriously slower in a network of six routes, which consist of only 16% of the total Indian Railways network, yet carry 58% of total freight traffic and 52% of total passenger traffic.[1] Since freight trains share the network infrastructure with passenger trains and are slowed down, Indian Railways has proposed constructing high-speed and high-capacity railway corridors exclusively reserved for freight transportation. Indian Railways expects these Dedicated Freight Corridors (DFCs) to help increase the speed of freight movement, improve throughput significantly, and lower the cost of logistics in the country. The DFCs are, thus, an attractive solution to the network congestion problem, especially when the demand for freight in the country is anticipated to grow in the coming years.[2]

The construction of these corridors requires a substantial upfront monetary investment. The two major corridors, the Eastern DFC and the Western DFC, that are currently under construction, have an estimated investment of 9.8 billion USD.[3] Given the considerable investment, it is crucial to quantify the anticipated benefits of the corridor before making the entire investment. In this counterfactual, we estimate the impact of such freight-only corridors on dwell reduction and the improvement in freight trains' speed using our estimated model. The freight-only corridor we simulate in this counterfactual is similar to the dedicated freight corridors proposed by the government in spirit, with a few operational differences. Specifically, we simulate an environment with no passenger train traffic, similar to the DFCs. Thus, we aim to quantify the freight train speed that

[1] https://pib.gov.in/PressReleasePage.aspx?PRID=1592433

[2] https://www.business-standard.com/article/economy-policy/railways-eyes-1-700-mn-tonnes-freight-in-fy23-15-beyond-budget-estimates-122041500895_1.html

[3] https://dfccil.com/Home/DynemicPages?MenuId=78

can be gained by eliminating the interference from passenger trains.

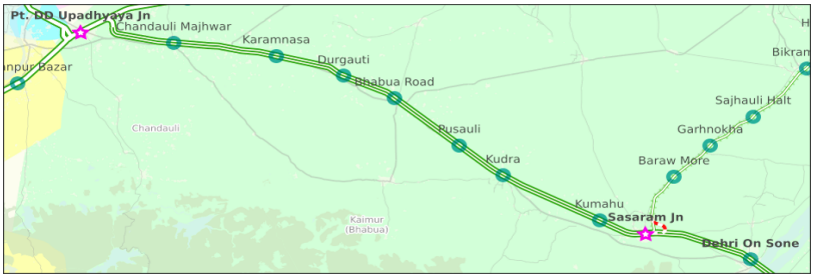

Figure 2.1: The railway section used for the FOC counterfactual analysis.

Simulation: To simulate the FOC counterfactual, we predict the section controllers' S/P and H/R decisions for freight trains in an environment without any passenger trains. The predicted decisions determine the delays and the overall speeds in the network under the "FOC" state of the world. We evaluate the efficacy of implementing the FOC by comparing freight dwell times, stoppages, and speeds with the FOC to the status-quo setting where both freight and passenger trains use the same network capacity.

For this counterfactual, we select the FOCM-ADSR section of the DDU railway division (see Figure 2.1) as it has the highest number of freight trains operating daily. When simulating the freight train flow in the FOC setting, we use the same arrival times for freight trains at the entry point of the section as in the status quo setting. We select 30 random days in our sample period, run our simulation for these 30 days, and report the statistics of our key measures of interest across these days. In every simulation instance, we predict the H/R and S/P decisions and obtain the movement trajectory of the trains. We calculate the control functions for each decision and use them as inputs to predict hold and stop probabilities. Simulating the freight flow on average on a typical day involves predicting around 4125 decisions. The simulations were performed on a system with an Intel i9-12900H CPU clocked at 2.40 GHz and equipped with 16 GB of RAM. Depending on the number of freight trains running daily, the run time for simulating all decisions for a single day ranges from 1.5 to 3 hours (more freight trains $\rightarrow$ more decisions $\rightarrow$ more run time).

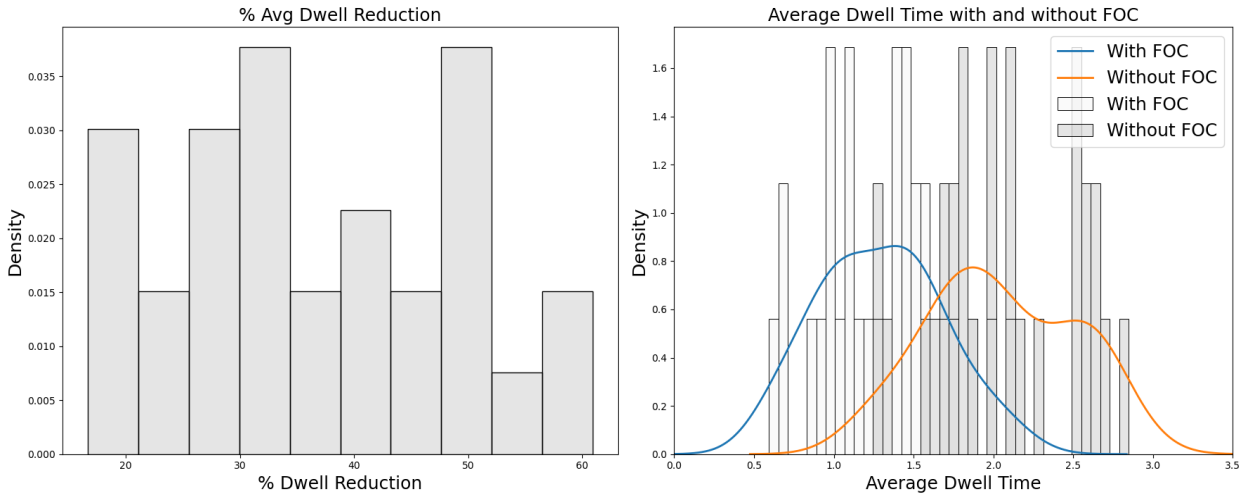

(a) Distribution of % reduction in daily dwell

(b) Dwell distributions: with and without FOC

Figure 2.2: Analysis of freight dwell reduction with FOC

Overall dwell reduction and sources:

Since this policy simulates an environment without passenger trains, we set the passenger train congestion variables to zero in this simulation. We denote the freight-only corridor with superscript FOC, then the reduction in daily average dwell time (ΔAvg_Dwell $=$ Avg_Dwell $-$ Avg_Dwell$^{\text{FOC}}$) is given by:

$$\Delta\text{Avg\_Dwell} = \left(\frac{\sum_i \sum_{s\in S_i} \sum_t x_{ist}}{\sum_i |S_i|}\right) - \left(\frac{\sum_i \sum_{s\in S_i} \sum_t x_{ist}}{\sum_i |S_i|}\right)^{\text{FOC}} \tag{2.1}$$

Figure 2.2 (left panel) shows the distribution of the percentage reduction in daily dwell times with the FOC. This distribution is calculated using the dwell reductions over the 30 sampled days. We note that the average reduction in daily dwell times is around 37%, suggesting that FOC leads to a significant reduction in delays. Figure 2.2 (right panel) shows the two distributions corresponding to the daily dwell times with and without FOC. The dwell time distribution with FOC has a much lower mean than the one without FOC. In addition to reporting the overall reduction, we also report the dwell reduction distribution for each station individually. Figure 2.3 shows the box plots for total dwell time reduction for different stations across the 30 sampled days. We note substantial

heterogeneity in dwell reduction across stations - some stations contribute disproportionately to the average reduction in delays. To understand the overall reduction and heterogeneity better, we investigate various sources of dwell reduction next.

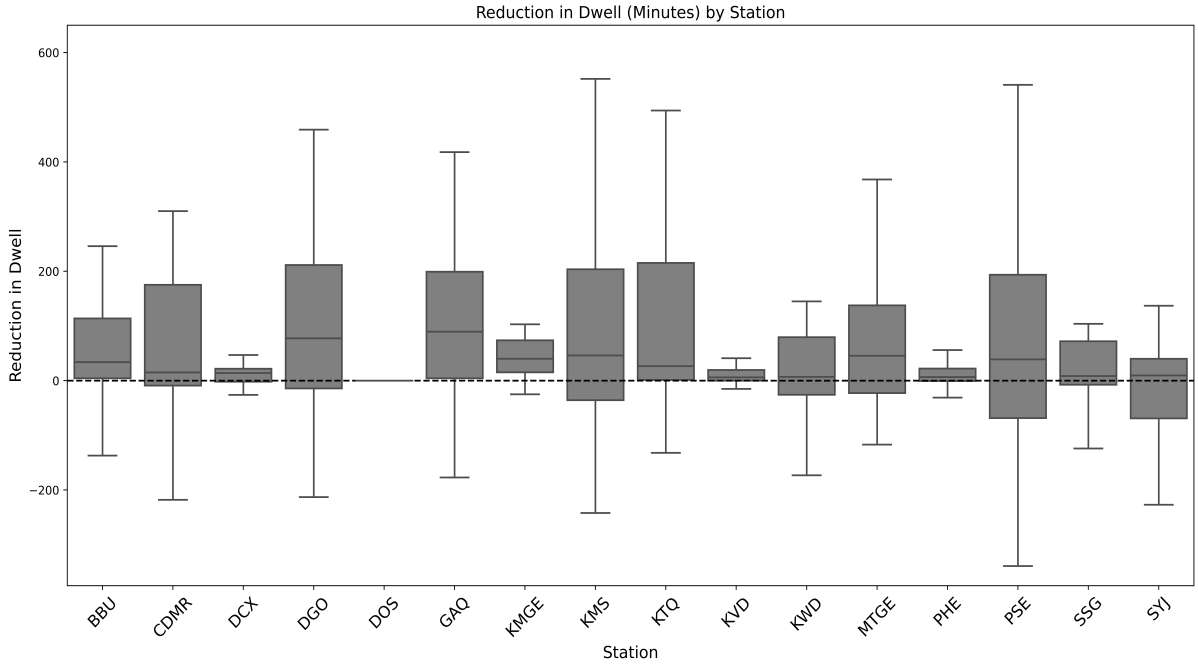

Figure 2.3: Box plots for total dwell reduction (in minutes) across non-junction stations.

Specifically, we analyze the reduction in dwell times by slicing our dwell reduction data by (i) station-level characteristics (junction vs non-junctions and level of track capacity) and (ii) hour slots in the day. Figure 2.4 (left panel) shows a more significant reduction in daily dwell time at junctions (1.68 minutes) than at non-junctions (0.58 minutes). We also categorize the stations based on their track capacity: Low (less than five tracks), Moderate (five to eight tracks), and High (greater than eight tracks). Like junctions, the stations with high track capacity exhibit a more significant reduction in dwell time with FOC. These observations highlight that FOC would reduce congestion at larger stations with a larger dwell capacity.

We also note from Figure 2.4 that the most considerable dwell reductions occur in Hour_Slot 1 (00:00 AM to 06:00 AM) and Hour_Slot 2 (6:00 AM to 12:00 PM). This observation is consistent

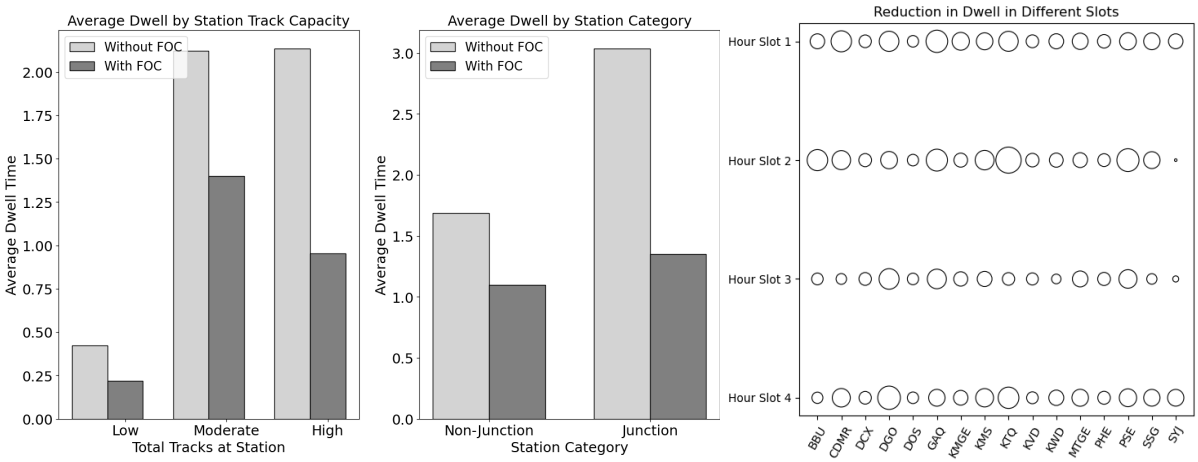

(a) Average dwell reduction : station characteristics (b) Average dwell reduction : hour slot

Figure 2.4: Sources of dwell reduction: station-characteristics (left) and hour slots (right).

with the temporal pattern in the passenger train flow data, as the morning slots have the highest number of passenger trains and the lowest number of freight trains for this railway section. On the other hand, Hour_Slot 3 (12:00 PM to 06:00 PM) has the most freight trains running, leading to a relatively low impact of FOC on dwell reduction. Next, we describe how we estimate the improvement in freight trains' speeds due to the FOC.

Measuring speed improvement due to the FOC

The increase in freight trains' speed due to the FOC can be attributed to two factors: (i) a decrease in dwell times and (ii) a reduction in stoppages. Stoppages lead to travel time losses due to acceleration and deceleration, significantly impacting the travel speed. To measure the speed improvement, we describe the impact of FOC on train stoppages and travel time losses next.

Stoppages Reduction: Since we predict the S/P decisions, we can quantify the change in train stoppages in the FOC setting. The reduction in average stoppages (ΔAvg_Stoppages $=$ Avg_Stops$-$

Avg_Stoppages$^{\text{FOC}}$) is given by:

$$\Delta\text{Avg\_Stoppages} = \frac{\sum_i \sum_{s \in S_i} y_{ist}}{|S_i|} - \left(\frac{\sum_i \sum_{s \in S_i} y_{ist}}{|S_i|} \right)^{\text{FOC}} \tag{2.2}$$

We find that the FOC leads to about 31% reduction in daily train stoppages. The left panel in Figure 2.5 shows the distribution of percentage reduction in daily stoppages, and the right panel shows the two distributions corresponding to the daily stoppages with and without FOC. Unlike dwell times, we find a more significant reduction in stoppages at non-junctions and stations with a moderate number of tracks (see Figure 2.12).

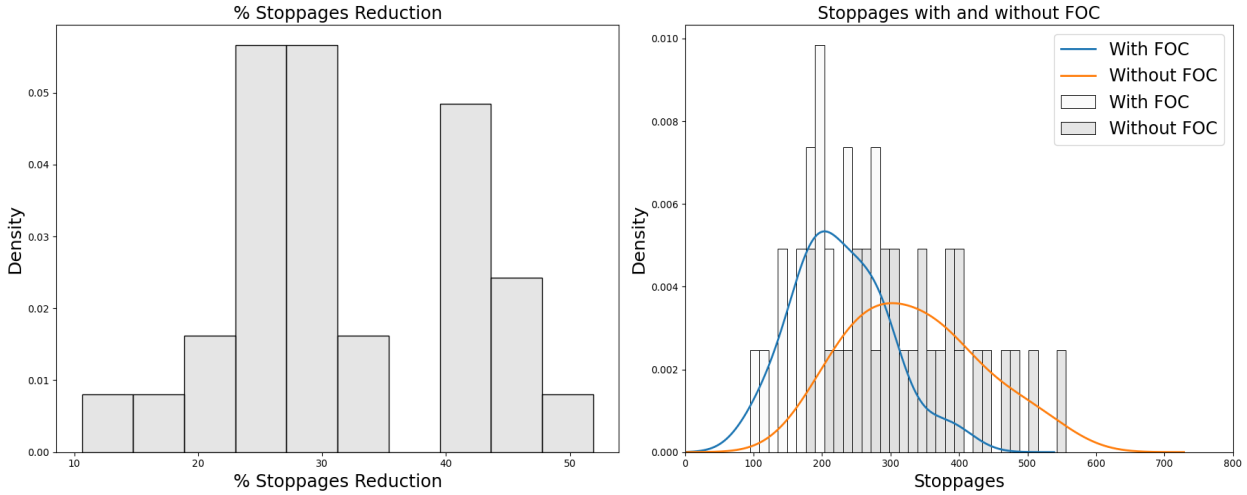

(a) % Stoppages reduction

(b) Stoppages distribution: with and without FOC

Figure 2.5: Stoppages reduction with and without FOC.

Acceleration and deceleration losses: Acceleration and deceleration losses occur along with stoppages because it takes some time for the train to reach its desired speed after starting from rest at a station and slowing down before stopping at the next station. Thus, accelerations and decelerations result in an overall increase in travel time. We provide the summary statistics for these travel time losses in our data in Table 2.1. A train passing between two stations may or may not stop at both stations, giving rise to the four possible stoppage scenarios as highlighted in Table 2.1. For example, when a train passes through two consecutive stations without stopping, the

average travel time is around 8.85 minutes. Similarly, when the train passes through the first station but stops at the subsequent station, the average travel time increases to around 12.75 minutes due to the deceleration losses at the subsequent station. It is worth highlighting that the travel time for the stop-stop scenario is 81% more than the passthrough-passthrough scenario. We also document the respective percentages of these four types of stoppages in the FOC counterfactual in Table 2.1. We calculate the weighted average travel times for the FOC and without FOC cases using these percentages. We find that the FOC leads to about a 5% reduction in average travel time.

Stoppage Scenarios	Average TT	% Loss	% Obs	% Obs (with FOC)	Acc. Loss	Decc. Loss
Passthrough-Passthrough	8.85	0	65.93	75.97	✗	✗
Passthrough-Stop	12.75	44	12.39	9.36	✗	✓
Stop-Passthrough	13.11	48	12.32	9.28	✓	✗
Stop-Stop	16.06	81	9.35	5.38	✓	✓

Table 2.1: Average Travel Time (TT) and Losses.

Increase in speed: There are two sources of a train's speed improvement: dwell reduction at the station and travel time reduction (due to fewer stoppages) on the blocksection. The average dwell time of a train in our sample is about 3.78 minutes, and the corresponding average time to travel on a blocksection is approximately 10.53 minutes. Then the improvement in average speed is given by:

$$\Delta\text{Avg\_Speed} = \frac{\left(\sum_i \frac{D_i}{t_{ia}^{s_i^l} - t_{id}^{s_i^f}}\right) - \left(\sum_i \frac{D_i}{t_{ia}^{s_i^l} - t_{id}^{s_i^f}}\right)^{\text{FOC}}}{|I|} \tag{2.3}$$

Note that $t_{ia}^{s_i^l}$ is a function of dwell time and travel time during the journey of train i, which are governed by the H/R and S/P decisions. We use the previously reported figures of a 37% reduction in dwell time and a 5% reduction in travel time to estimate the speed improvement achievable with the implementation of FOC.

The average length of a block section in this railway section is about 6.14 kilometers. The total time taken by a train to cross the blocksection without FOC is 14.31 minutes (3.78 minutes of dwell and 10.53 minutes of running time on the blocksection). The corresponding total time

with FOC is 12.37 minutes (2.38 minutes of dwell time and 9.98 minutes of travel time). Thus, we find that the average speed of a freight train increases from 25.74 kmph to 29.79 kmph (15.73% increase) with FOC.

	Travel Time	Dwell Time	Flow Time	Speed (KM per Hour)
Without FOC	10.53	3.78	14.31	25.74
With FOC	9.98	2.38	12.37	29.79
% Change	-5.14%	-37.14%	-13.59%	15.73%

Table 2.2: Travel Time, Dwell Time, Flow Time, and Speed.

2.2.1 Partial Sharing of FOC Track

Even with the introduction of DFCs, there are circumstances where passenger and freight trains can share tracks. For instance, the Railway Board has permitted passenger trains to operate on the DFC network during exigencies, such as accidents or maintenance activities on the regular passenger routes. This ensures continuity of passenger services when the primary tracks are unavailable.[4] Conversely, freight trains may continue to use the original shared network in areas where DFC infrastructure is incomplete or not yet extended. To this end, we now simulate a set of partial DFC scenarios where a fraction (0%, 25%, 50%, or 75%) of the current passenger traffic is occupying the DFC corridor. These scenarios recognize that even with DFCs in place, some level of mixed traffic may persist under exigent circumstances. Table 2.3 shows the resulting improvements. As expected, we find that the speed improvement for 0% level (no passenger traffic) is highest, mimicking full FOC.

[4] https://economictimes.indiatimes.com/industry/transportation/railways/passenger-trains-can-run-on-dfc-network-during-exigencies/articleshow/90959882.cms?from=mdr

Measure/ Passenger Traffic on FOC	0%	25%	50%	75%
Δ Average Dwell Time	37.14%	33.57%	24.51%	17.01%
Δ Passthrough Percentage	31.83%	27.90%	20.31%	14.42%
Δ Speed	15.73%	13.63%	9.59%	6.62%

Table 2.3: Partial Offloading of Passenger Traffic.

2.2.2 DFC-like Network Enhancements

Our FOC counterfactual is designed to isolate the core benefit of removing passenger train interference from shared infrastructure. For this reason, we referred to it as a Freight-Only Corridor (FOC) rather than a full-fledged Dedicated Freight Corridor (DFC). DFC involves additional infrastructural enhancements such as fewer stations, gentler track curvature, support for high-powered locomotives, and automatic signaling systems.[5][6] We take into account some of these additional features in this section. Specifically, we reduce the number of stations by 50% and assume a 30% higher average inter-station speed (65 kmph) to account for higher-powered locomotives and reduced track curvature. By combining these changes with different levels of passenger interference, we calculate resultant speed improvement, shown in table 2.4 and figure 2.6.

Measure/ Passenger Traffic on FOC	0%	25%	50%	75%
Δ Speed	46.56%	45.14%	43.10%	42.21%

Table 2.4: Partial Offloading of Passenger Traffic with Network Enhancement.
Baseline Case: No FOC, No Network Enhancement, Speed: 25.74 kmph.

Next, we describe two counterfactuals that use the existing network capacity. In the second counterfactual, we study threshold-based rules as potential levers to improve speed. In the third counterfactual, we study capacity consolidation as a lever to reduce congestion losses and improve the average freight speed through the section.

[5] DFC Corporate Plan (2020)
[6] India Today, 2020

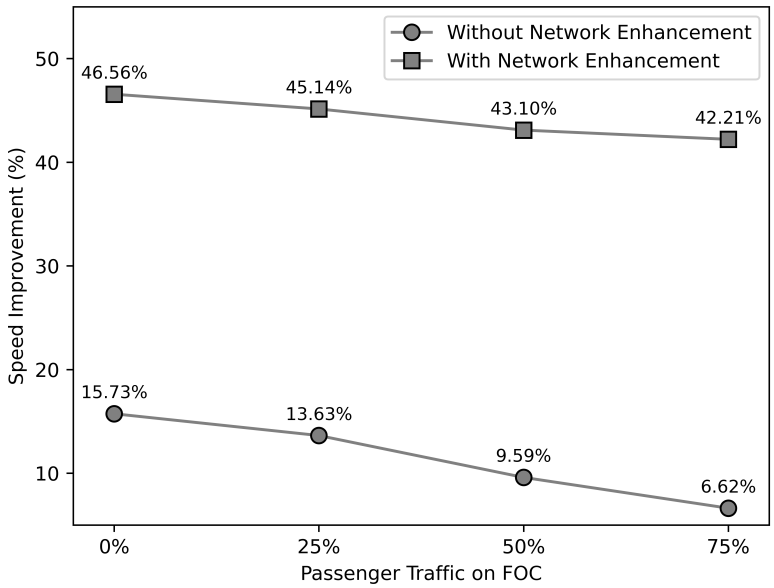

Figure 2.6: Speed Improvement: With and Without DFC-like Network Enhancement.

2.3 Threshold Based Releases

The higher movement priority given to passenger trains results in an extended dwell time for freight trains. In this section, we propose a prioritization rule for freight trains that exceed a predefined threshold for dwell time. Specifically, we evaluate a threshold-based release rule for freight trains dwelling longer than a specified time limit.

For this analysis, we use the freight and passenger movement data for the same section as in 2.2. We simulate freight train trajectories using our estimated H/R and S/P models with a slight difference: we introduce a threshold for dwell times in H/R decisions. Specifically, when a freight train crosses a dwell time threshold at a station, it is subjected to a forced release. Specifically, a freight train i is released from station s at time t' when $\sum_{t=t_{ia}^{s}}^{t'} x_{ist} > \tau$, where t_{ia}^{s} is the arrival time of train i at station s, and τ is the threshold dwell time. We use seven different thresholds and calculate the resultant reduction in dwell times and stoppages, the increase in freight speed, and the

number of such forced releases (please see Table 2.5). We find that the number of forced releases at the release thresholds of 60 and 75 minutes is about 11 and 9, respectively. These releases account for less than 2% of the total section controllers' decisions. We also report the speed improvement obtained from these releases in Table 2.5. It is worth noting that even with these low numbers of forced releases, we get a considerable speed improvement ranging from 6% to 23%.[7]

Measure/τ	30	45	60	75	90	105	120
Δ Average Dwell	-52.37%	-41.66%	-34.25%	-31.98%	-26.91%	-22.13%	-17.59%
Δ Passthrough Percentage	40.52%	32.91%	26.88%	26.83%	21.74%	17.70%	14.52%
Δ Speed	22.73%	17.44%	13.94%	13.18%	10.71%	8.72%	6.71%
Avg. Interference	6.03	3.53	2.83	2.50	1.61	1.40	0.80
Avg. Threshold Releases (%)	1.23%	0.93%	0.74%	0.59%	0.53%	0.46%	0.44%

Table 2.5: Threshold based releases (τ: Threshold in Minutes).

Overall, these threshold-based releases can effectively reduce dwell times and improve trains' speed. For example, the dwell time reduction and speed improvement achieved with a 60-minute threshold level are 34% and 14%, respectively. These improvements are comparable to those achieved with the FOC. What is interesting is that a very low percentage of these forced freight train releases (around 1%) can achieve these impacts. Moreover, implementing the threshold release policy is costless and does not involve new infrastructure. The only additional cost incurred is the cost of actively monitoring the dwelling status of the freight train, which is minimal.

Remark: Implementing the policy in practice: To implement the policy in practice, we recommend releasing the freight trains expeditiously whenever any suitable opportunity arises based on real-time passenger train traffic. Specifically, instead of prescribing a fixed threshold value, we recommend a flexible approach where releases can occur within a specified time range whenever feasible. The flexible approach also provides a cushion to ensure threshold-based releases do not interfere with passenger train schedules.

[7]These results are conditional on the passenger-train schedule for the section under study and may vary under different scheduling scenarios.

2.3.1 Potential Interference to Passenger Trains

One may question why we do not incorporate the potential spillover caused by this threshold policy to passenger trains in our simulation. As we described in Chapter 1, the passenger trains are scheduled, and currently, their scheduling does not consider freight trains. Hence, it is not possible to estimate a model that accounts for the interaction between freight train delays and passenger train traffic (freight delays $\rightarrow$ change in passenger train schedule $\rightarrow$ new freight delays $\rightarrow \ldots$ till equilibrium).

We highlight that the percentage of such threshold-based releases is relatively low (less than 2% of controllers' decisions), limiting potential spillovers to passenger trains. Also, we recommend releasing the freight trains judiciously in a time window that minimizes the impact on the passenger trains in Remark 2. In Figure 2.7 we plot the *performance frontier* between speed improvement and $(1/\text{Average Interference})$ across different threshold values τ. The frontier illustrates that as we aim for speed improvement for freight trains, interference to passenger trains also rises (i.e., $1/\text{Average Interference}$ decreases). In evaluating the overall performance of threshold-based releases, we consider the section controller's dual objective of maximizing freight train speed and minimizing interference to passenger trains. Assuming that the controller has a Cobb-Douglas-type preference for improvements in freight speed and reducing interference to passenger trains, we model the objective function for the section controller as:

$$\text{CB}(\tau;\gamma) = (\Delta\text{Speed}(\tau)_{\text{norm}})^{\gamma} \times \left(\frac{1}{\text{Avg. Interference}(\tau)_{\text{norm}}}\right)^{(1-\gamma)}$$

γ and $1-\gamma$ represent the relative importance given to the two competing objectives.[8] By varying γ, we generate a frontier of solutions. This approach facilitates an assessment of optimal threshold selection, guiding the decision-making for section controllers. The figure 2.8 shows the objective

[8]To make the performance metrics comparable and scale-invariant for computing the Cobb–Douglas score, we apply min-max normalization to both $\Delta\text{Speed}(\%)$ and $1/\text{Avg. Interference}$.

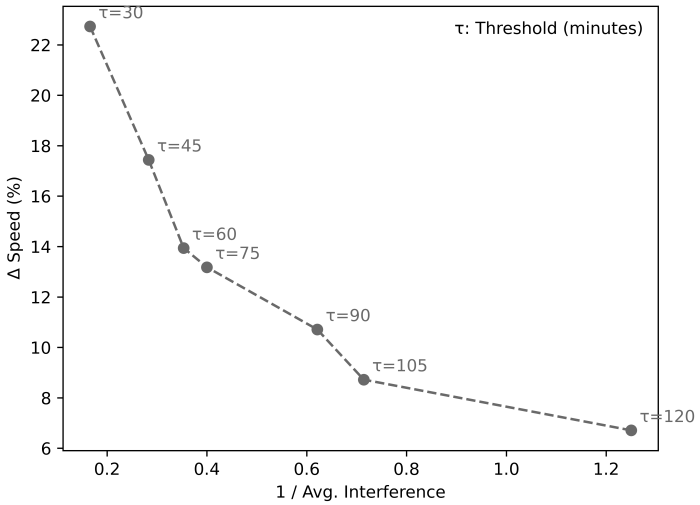

Figure 2.7: Speed Improvement vs Avg. Interference.

function value for different thresholds and γ values.

2.4 Capacity Consolidation: Vertically Stacked Trains

One potential solution to speed up freight that does not require additional capacity investment is consolidating existing freight capacity by using vertically stacked trains. This capacity consolidation would keep the total freight volume transported across the network the same, but result in fewer vertically stacked trains. For example, if the Indian Railways operates five trains with 100 wagons each to transport cargo between origin station o and destination station d, instead of five trains, the railways could operate four trains with 125 wagons each. The extra 25 wagons would be added to the trains at the origin station by vertically stacking them on the existing wagons, resulting in a train depicted in Figure 2.9. This consolidation of train capacity is not uncommon. In the past, Indian Railways has experimented with vertical stacking of freight trains.[9] Using these trains results in a reduction in the number of trains in the network, thereby reducing congestion

[9] https://www.cnbctv18.com/business/companies/gateway-distriparks-starts-double-stack-rail-transport-between-ahmedabad-and-mundra-port-17084151.htm

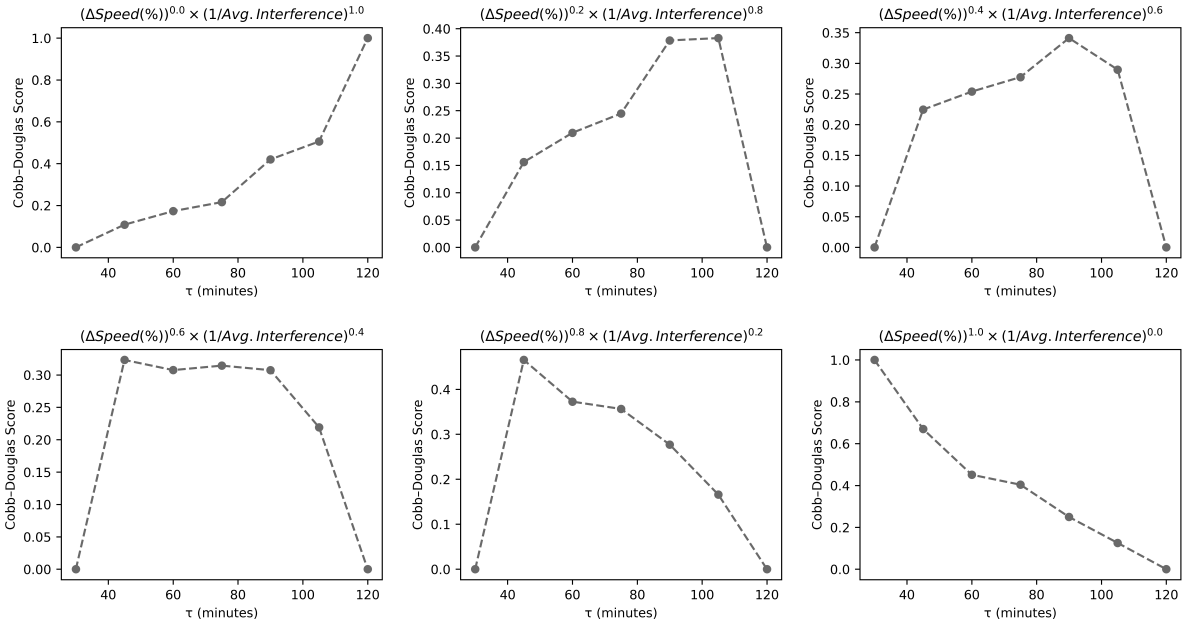

Figure 2.8: Cobb-Douglas Scores for Different Weightages to Competing Objectives.

in busy sections and increasing freight speed. In this counterfactual, we study the impact of using this form of capacity consolidation on trains' dwell times and speeds.

Figure 2.9: A vertically stacked freight train.

A typical IR freight train comprises 58 wagons, each having a capacity of 90 tonnes. A train could be consolidated by vertically stacking a few or all the existing freight train wagons. The relationship between the consolidation level and the number of trains is given by:

$$n_{old} = (1 + m) \cdot n_{new}, \tag{2.4}$$

where m is the consolidation percentage and n_{new} is the number of consolidated trains. For example, a 25% consolidation would reduce the total number of freight trains required to transport the same freight volume by 20%. We apply different consolidation percentages in this counterfactual, which gives us different average numbers of freight trains in the system. To generate the arrival epochs of freight trains, we use discrete event simulation. We sample the inter-arrival times from the empirical distribution of inter-arrival times in our data (see Figure 2.10). Additionally, we use empirical distributions of various train characteristics, such as loading status (loaded or empty) and wagon type of freight trains, while generating the simulation instances. Further, we consider the distribution of route-wise traffic (both in the Up and Down directions) to ensure an accurate representation of real traffic in the simulation.[10]

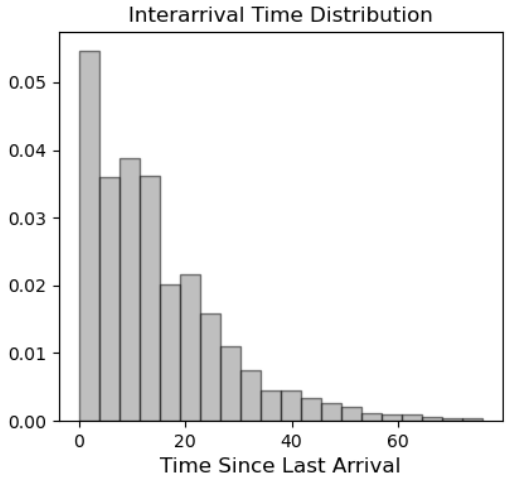

Figure 2.10: Interarrival time distribution for freight trains.

We perform this counterfactual analysis for the same railway section as the one in the previous counterfactuals. To simulate the arrival of freight trains at different consolidation levels, we sample the interarrival time of freight trains from the appropriately scaled empirical distributions. For example, let X be the random variable corresponding to the distribution of inter-arrival times. Then, for simulating a system with 25% consolidation, we sample inter-arrival times from the distribution of 1.25 X. To model the trajectory of freight trains over the network, we use the two

[10]We use histogram-based sampling from the historical freight train data. It involves dividing the historical data into bins, counting the occurrences within each bin, and randomly selecting a data point based on the bin probabilities. This method captures the data's empirical distribution when a suitable parametric distribution is unavailable.

estimated H/R and S/P models like in the previous counterfactuals. We assume that parameters for H/R and S/P decisions remain the same while simulating the trajectories of consolidated freight trains. It is a justified assumption, as we do not analyze consolidation levels that require separate infrastructure to cater to the vertically stacked freight trains. We simulate ten instances for each consolidation level to account for variability in interarrival times and report the average measures across these instances. In Table 2.6, we report the change in average dwell time, passthrough percentage, and speed for different consolidation levels. These values are averaged over ten instances and are relative to the baseline scenario, which represents the current traffic levels for freight trains.

Consolidation Level	10%	15%	20%	25%	30%
Δ Dwell Time	3.72%	6.74%	8.93%	12.05%	14.76%
Δ Passthrough Percentage	6.13%	12.07%	22.07%	34.86%	37.93%
Δ Speed	1.26%	2.49%	4.36%	7.74%	8.43%

Table 2.6: Change in speed: Consolidating freight.

We observe a significant speed improvement due to freight consolidation. Specifically, a 25% consolidation leads to a 12% reduction in average dwell time and a 35% increase in passthrough percentage, which leads to an overall improvement of 8% in speed. We emphasize that, compared to the FOC intervention, capacity consolidation doesn't require capital expenditure.

Remark 3: Practical Implementation: Our recommendation is to run consolidated trains from the origin station to the destination station. Note that we do not propose stopping a train and adding wagons to it when it is en route to the destination. Thus, running consolidated trains does not require additional waiting to assemble the wagons once the train is en route. Also, since we propose stacking the trains vertically instead of making them longer, our approach ensures that no extra track capacity is needed to accommodate the vertically stacked train. This also ensures very limited negative externalities to other trains in the network.

2.4.1 Consolidation Losses

In the preceding analysis, we assumed that the vertical consolidation of freight trains incurs no additional time losses. This simplifying assumption allowed us to isolate the pure congestion-alleviation benefit of consolidation. However, in practice, consolidating trains may involve time delays, particularly at the origin station, which can offset some of the speed benefits. In this section, we account for these losses to realistically evaluate the net impact of consolidation on freight speeds.

We note that the total consolidation time can vary by station, wagon type, train length, and commodity type (e.g. more for coal), etc. Recognizing this heterogeneity, we simulate a range of consolidation scenarios to offer a menu of results that railway managers can use based on station-specific operational realities. To incorporate these losses into our analysis, we assume that the upfront time loss associated with consolidation is distributed proportionally across the train's journey. Specifically, based on the average freight haul distance of 590 km in India[11] and an average inter-station spacing of 5.38 km (from our dataset), we translate the total consolidation time at the origin into a per-station time penalty. For example, a 30-minute delay in consolidation translates to an approximate per-station dwell increase of ($\frac{5.38}{590} \times 30 \approx 0.27$ min). We then re-calculate net speed improvements by adding this per-station delay to the predicted dwell times from our model.

Table 2.7 summarizes the resulting speed improvements across different combinations of consolidation percentages (10% to 30%) and consolidation time losses (15 to 75 minutes). As expected, we find that consolidation yields significant speed gains only when the associated time losses are modest. With higher time penalties, the congestion relief benefits are gradually outweighed, especially at lower levels of consolidation. For instance, a 25% consolidation with a 15-minute penalty yields a 6.74% improvement in speed. However, the same level of consolidation with a 60-minute penalty sees the gain fall to just 3.84%. At the lower end, a 10% consolidation with a 75-minute delay actually leads to a net speed reduction. These results emphasize the im-

[11] https://www.ceicdata.com/en/india/railway-passenger-and-freight-traffic-monthly?

Time Loss / Consolidation	10%	15%	20%	25%	30%
15 Minutes	0.37%	1.58%	3.41%	6.74%	7.41%
30 Minutes	-0.50%	0.68%	2.49%	5.75%	6.42%
45 Minutes	-1.36%	-0.19%	1.58%	4.79%	5.43%
60 Minutes	-2.20%	-1.05%	0.69%	3.84%	4.49%
75 Minutes	-3.03%	-1.90%	-0.19%	2.90%	3.53%

Table 2.7: Change in speed with consolidation losses.

portance of operational feasibility and context-specific implementation. By providing a spectrum of outcomes under varying scenarios, we enable railway managers to assess whether the expected gains from consolidation justify the required operational changes, potentially customizing deployment based on cargo type, station throughput, and resource availability.

2.5 Costs, Benefits and Feasibility Analysis

In closing, we briefly summarize the benefits and the associated costs of the three policies (Table 2.8). While all three policies result in substantial speed improvements, they all have varying implementation costs. The FOC policy leads to about 16% speed improvement but is also the most costly intervention due to the vast capacity investment in building the new corridor. The government plans to spend about 10 billion dollars on constructing the corridor.

Policy $\mathcal{P}$	Δ Speed	Costs
Freight Only Corridor (FOC)	16%	Building new infrastructure (high cost)
Threshold Release (60-45 Minutes)	13-17%	Possible interference to passenger trains (low cost)
Vertical Stacking (25%)	8%	Requirement of higher power locomotives, and OHE Time to stack (low cost)

Table 2.8: Policies: Costs and Benefits.

In comparison, the threshold-based releases offer a 13-17% speed improvement depending on the dwell threshold (60min/45min). One potential but small negative externality of this policy is the potential interference with passenger train movements. However, as we highlighted before, a carefully designed threshold policy can minimize such interferences. Finally, the vertical stacking

policy can provide a speed improvement of up to 8%. A potential cost involved with the policy is the requirement of higher-powered locomotives for older trains to accelerate the heavier vertically consolidated train. However, since the trains' current speeds are low, the locomotives deployed at present can sufficiently power most of the freight trains. Even though we can not quantify the exact monetary costs associated with the non-capacity intervention policies, our conversations with IR suggest that the associated costs are orders of magnitude smaller than the FOC policy. Thus, our proposed non-capacity interventions are cost-effective and yield significant speed improvements compared to the cost-intensive capacity intervention.

Feasibility Analysis

FOC Counterfactual: Dedicated Freight Corridors (DFCs) are already operational in India and represent a major infrastructural advancement. In our analysis, we frame the capacity counterfactual as a Freight-Only Corridor (FOC) to specifically isolate the benefit of removing passenger train interference from shared tracks. While the full DFC includes additional enhancements—such as reduced curvature, high-power locomotives, and automatic signaling—our counterfactual complements the DFC by (i) quantifying the standalone gains from eliminating passenger interference, and (ii) providing performance bounds for scenarios that progressively incorporate DFC-like features. This approach clarifies both the marginal and cumulative benefits of freight prioritization and infrastructure upgrades.

Threshold-based releases: While prioritizing freight over passenger trains is generally not the preferred policy, there are documented instances where Indian Railways (IR) has done so—delaying or canceling passenger services to accommodate critical freight shipments.[12] [13] Moreover, there is growing pressure on IR to improve freight reliability, including government initiatives to compen-

[12] https://timesofindia.indiatimes.com/city/hubballi/many-passenger-trains-delayed-as-swr-gives-priority-to-freight/articleshow/71447387.cms

[13] https://swarajyamag.com/infrastructure/over-16-passenger-trains-cancelled-daily-to-make-way-for-over-400-coal-rakes-on-priority

sate customers for delayed freight deliveries.[14]

Our proposed Threshold-Based release policy builds on these precedents in a more structured and targeted way. Rather than blanket prioritization, we recommend expediting freight trains that have already experienced significant delays—releasing them after a defined threshold wait time while keeping the interference to passenger trains to a minimum. We also explicitly account for the level of interference caused to the passenger trains by releases. Hence, we provide a full menu of solutions to the railway manager to help strike a balance between improving freight performance and maintaining service quality for passengers. The policy can be implemented within the constraints of existing infrastructure and scheduling flexibility.

Double Stacking: Double-stacking is not only feasible but already operational in certain parts of the IR network, like the Western DFC, where successful trials spanned 1,158 km.[15][16] Having said that, there are some challenges with the feasibility of double stacking in other parts of the network because the height of electrified lines in these parts is too low to support double stacking.[17]

To circumvent this, IR has tested double-stack dwarf containers—6 feet 4 inches tall, 662 mm shorter yet 162 mm wider than standard ones. By using these containers, IR can boost volumes by 67% and fit these containers into most existing tracks without structural modifications. Hence, the level of consolidation that we test in our counterfactual (up to 30%) can be easily supported by the existing infrastructure. Moreover, IR is making further infrastructure modifications, such as high-rise overhead equipment and structural adjustments, to enable double stacking with standard containers, which will enhance its feasibility on existing infrastructure.[18][19]

[14] https://economictimes.indiatimes.com/industry/transportation/railways/piyush-goyal-indicates-customers-to-get-compensation-for-delayed-freight-delivery/articleshow/73359509.cms?from=mdr

[15] pib.gov.in/PressReleseDetailm.aspx?PRID=1686785&reg=3&lang=1

[16] www.youtube.com/watch?v=OYgh-I7EX9Y

[17] www.globalrailwayreview.com/article/132297/innovative-logistics-stacking-small-containers-optimised-loading/

[18] https://www.railwaygazette.com/download?ac=32996

[19] https://www.jagranjosh.com/general-knowledge/worlds-first-electrified-double-stack-container-tunnel-1595838008-1

2.6 Conclusion

In this chapter, we conduct counterfactual analyses to provide a set of solutions for addressing the problem of slow freight train speeds. Specifically, we evaluate the impact of Freight Only Corridors (FOC) and other non-capacity-based interventions like threshold-based releases and capacity consolidation. For FOC intervention, we find a considerable dwell reduction (37%), stoppage reduction (35%) and speed improvement (16%). Though the FOC solution effectively reduces freight delays, it is also highly capital-intensive. Our non-capacity interventions offer cost-effective alternatives to improve freight train speed while capturing benefits comparable to those of FOC. Specifically, a 60-minute threshold release policy leads to around 13% reduction in dwell times. The policy can be implemented easily and flexibly, and only leads to very few forced releases. Similarly, consolidating freight capacity by about 25% results in around 8% increase in speed. We provide clear and actionable recommendations for reducing freight delays. By tackling the issue of slow freight train speeds and congestion, the recommendations could enhance the overall efficiency of India's transportation infrastructure, benefiting the country's economic and social development. As Indian Railways is a linchpin in India's logistics, the study's outcomes have far-reaching implications for the broader Indian economy and transportation network.

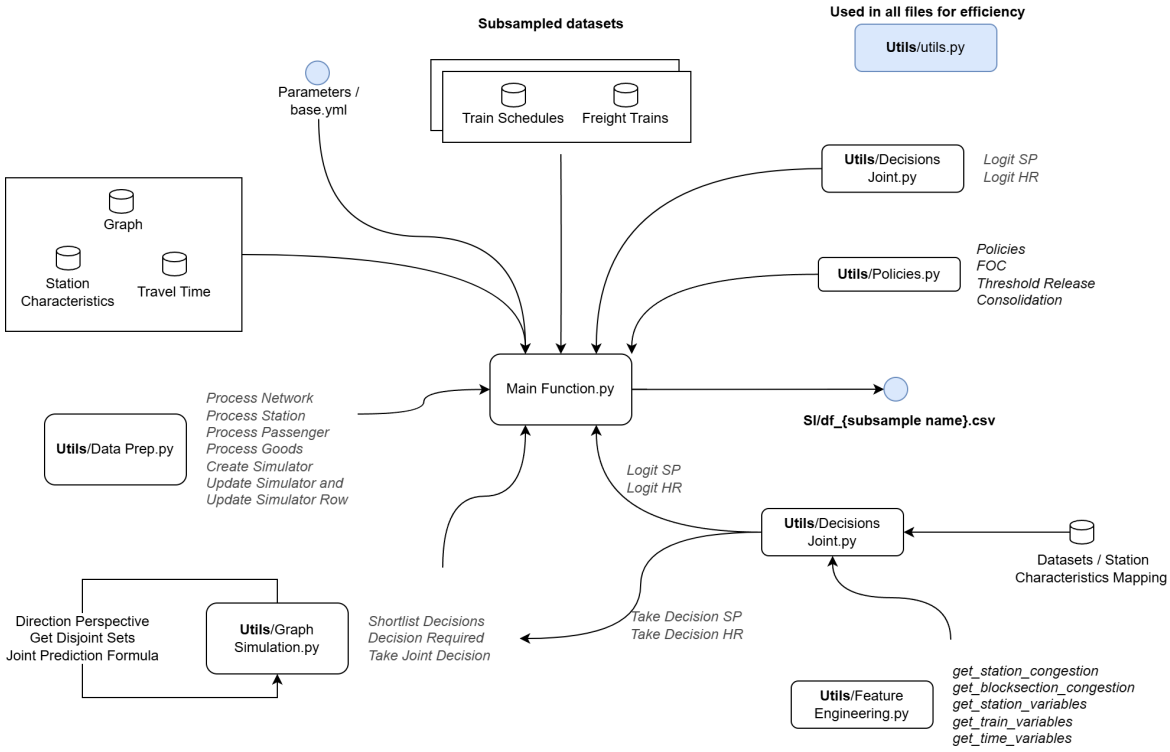

Figure 2.11: Simulator Schema.

2.7 Appendix

2.7.1 Simulator Schema

As we have estimated the H/R and S/P models for only freight trains, we don't predict movement decisions of passenger trains. We treat the passenger train schedule as purely exogenous and use it as an input in our simulator. In Figure 2.11, we present the schema of the simulator, showing the role of passenger train schedules. Thus, in the case of Threshold-Release leading to interference, there is a possibility that the passenger trains will deviate from their schedule. But we restrict ourselves to just quantifying the interference instances.

2.7.2 Results Appendix

Figure 2.12 (top panel) shows the station-wise stoppage reduction. We observe considerable heterogeneity in the extent of stoppage reduction across stations. We find a more significant stoppage reduction at non-junction stations and stations with moderate track capacity (see Figure 2.12 left bottom panel). Similar to the results of dwell reduction, we find that stoppage reduction is higher in Hour_Slot 1 and Hour_Slot 2 (see Figure 2.12 right bottom panel).

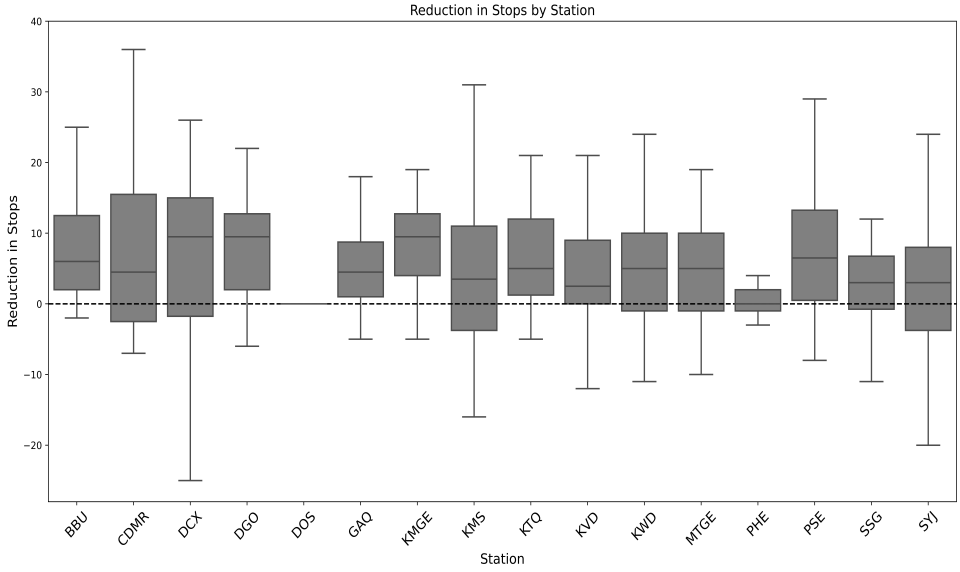

(a) Stop reduction by station

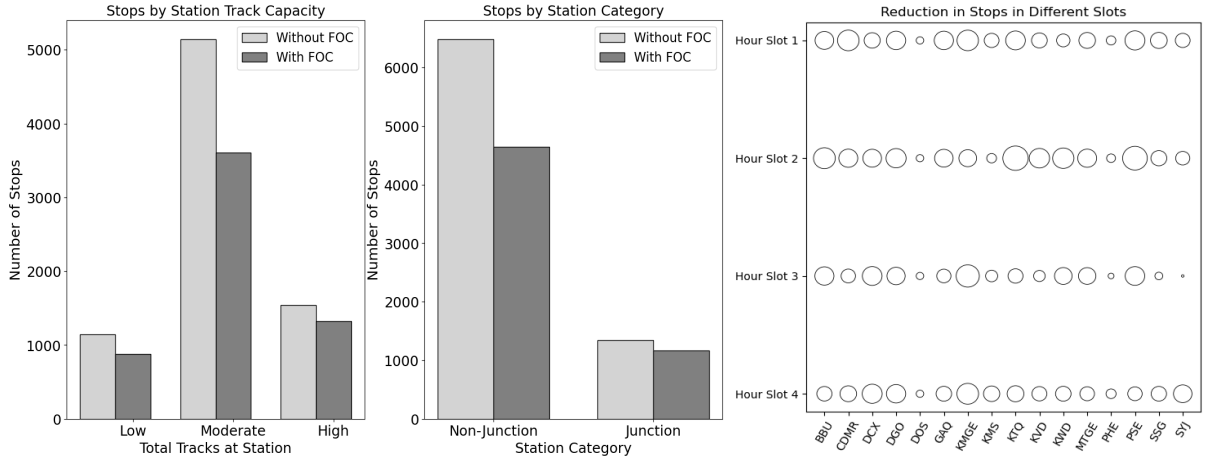

(b) Stop reduction : Station characteristics

(c) Stop reduction : Hour slot wise

Figure 2.12: Sources of stop reduction.

	Travel Time	Dwell Time	Flow Time	Speed (KM per Hour)
Without FOC	9.26	7.53	16.79	21.15
With FOC	8.97	4.81	13.78	25.76
% Change	-3.13%	-36.10%	-17.92%	21.83%

Table 2.9: Travel Time, Dwell Time, Flow Time, and Speed ('ADSR-MPO' Section).

FOC Counterfactual on Alternate Section

In order to test the robustness of the FOC counterfactual results, we run the simulation for an alternate section 'ADSR-MPO'. This section is approximately 88 km long and has 15 stations. In Table 2.9, we provide the speed improvement results for the FOC counterfactual for this section. We observe an average of 36% reduction in dwell time per station and a 35% reduction in stoppages, which translates to a speed improvement of around 22%. These numbers are in line with the results on 'FOCM-ADSR' section and hence support the robustness of our findings.

Economic Impact of Increased Freight Train Speed

Quantifying the exact monetary benefits of speed increases is challenging because it depends on various factors such as network configuration, rolling-stock utilization, and demand conditions. Nevertheless, available reports suggest meaningful savings and revenue gains. In the case of Indian Railways, 1 kmph increase in the average speed is estimated to free approximately 3000 wagons and 70 locomotives, translating into additional annual revenues of about INR 1400 crore from freight and INR 900 crore from passenger train operations.[20] For North American freight operations, where freight trains run at an average speed of 20 miles per hour (mph), 1 mph increase in the average speed is estimated to reduce annual costs by more than USD 200 million for Class-1 railroad.[21] These estimates indicate that even modest improvements in average train speed can deliver significant financial benefits.

[20] https://www.dailypioneer.com/2020/columnists/get-on-the-right-track.html

[21] https://www.informs.org/Impact/O.R.-Analytics-Success-Stories/RailConnect-360-Revolutionizes-Approach-to-Train-Scheduling

Throughput Improvement

Table 2.10 summarises throughput improvement under various policies. Although the numbers indicate potential gains, they should be interpreted with caution. Throughput is jointly determined by demand, passenger schedules, and capacity utilization. Consequently, the reported metric—'Average Additional Freight Trains / Day'—is an indicative number, contingent on sufficient freight demand and on passenger timetables and utilization levels permitting the additional freight trains.

Policy	Flow Time (minutes)	Average Additional Freight Trains / Day
FOC	12.37	15.78
45-minute Threshold	12.18	17.55
60-minute Threshold	12.56	14.03
25% Consolidation	13.28	7.79

Table 2.10: Throughput gains under different policies, relative to a baseline average flow time of 14.31 minutes for a 6.14 km block section.

Chapter 3

Dynamic Control of Priority Queues

3.1 Introduction

Priority queues are commonly found in settings where servers are allocated to customers based on varying levels of urgency or importance. Such queues are prevalent across diverse domains, such as healthcare, call centers, and transportation. While prioritization strategies aim to optimize outcomes for high-priority customers, this may inadvertently increase delays and dissatisfaction among non-priority customers. For instance, in hospitals, such prioritization has significantly increased wait times for elective surgery patients, as observed in recent reports of delays nearly doubling in Australia.[1] Another example is ambulance dispatching, where prioritizing critical cases affects response times for non-critical cases. In Singapore, this approach extends response times for non-critical cases to up to 20 minutes. While this ensures urgent cases are prioritized, it results in extended delays for non-urgent cases.[2] Similarly, in call centers, executives often prioritize premium customers, who hold loyalty memberships, leaving regular customers with prolonged wait times.

We witness a similar scenario in traffic management in railroad systems, where passenger trains

[1] https://www.abc.net.au/news/2024-04-19/elective-surgery-wait-times-almost-double-in-20-years/103742480

[2] https://www.channelnewsasia.com/singapore/scdf-non-emergency-medical-calls-wait-time-ambulance-traffic-signal-myresponder-4456381

are often given priority over freight trains. This prioritization frequently results in significant delays and low speed for freight trains.[3] In some of these cases, such as extended waiting times for ambulances handling non-urgent cases or prolonged freight train delays, resources (ambulances and railway tracks, respectively) are deliberately idled in anticipation of priority customers. Following the nomenclature of Baron et al. [2014], we hereafter refer to this phenomenon as *strategic idling*. Furthermore, the availability and accuracy of customer arrival information play a critical role in the controller's decision-making process when determining whether to idle the resource.

In this work, we study the service control policies governing such priority queues. Specifically, we study a finite-buffer priority queue with two customer classes—a priority class and a non-priority class—both requiring class-specific deterministic service times. The objective is to characterize the optimal service control policies for the non-priority customers.

Modelling Challenges Modelling the above-mentioned priority queue setup involves several challenges. First, while finite buffer systems are essential for capturing real-world capacity limitations, most studies either assume infinite buffer capacity or model a loss system, where customers are turned away when the buffer is full. Second, a shared finite buffer requires tracking class-wise queue lengths in the state. This increases the state-space dimension and makes deriving structural properties harder. Moreover, in settings such as ambulance services and inter-station train runs, service times show low variability. Therefore, modelling the service time using an exponential (memoryless) distribution is inappropriate, as service times are deterministic or near-deterministic.

Though modelling deterministic service times is central to our setting, it makes the problem substantially harder. Under the exponential service time assumption, the memoryless property allows a Markovian specification in terms of the queue length alone. With deterministic service time, however, the *service age* or *residual time* must be additionally tracked to preserve the Markov property. For example, consider a single-server queue with Poisson arrivals and deterministic

[3] https://www.eca.europa.eu/lists/ecadocuments/sr16_08/sr_rail_freight_en.pdf

service time. If we set the decision epochs at arrival and use only the queue length y as a state variable, two states with identical y but different residual service time will have different transition probabilities, because the probability that service completes before the next arrival also depends on residual time, not on y alone. Hence, under deterministic service, the queue-length process alone is not Markovian. To address these challenges, we reformulate the problem as a semi-Markov decision process by redefining the decision epochs. This reformulation enables structural analysis and the derivation of state-dependent control policies for a finite-buffer priority queue with deterministic service times.

Incorporating Visibility Recent advances in real-time tracking technologies have enabled queue controllers to access partial or full information regarding future arrivals. Controllers benefit from real-time information on customers, allowing them to estimate both the time until the next arrival and the type of the next arrival (priority or non-priority). This becomes particularly important, as the controller's ability to anticipate future priority arrival can directly influence their decision to hold or serve a non-priority customer. In order to address this, we extend our analytical model to incorporate visibility in terms of arrival lookahead, wherein the controller receives information about the time and type of the next arrival. While the characterization of optimal policies under this setting is challenging, we propose a modelling framework and present simulation results to show the impact of such visibility.

The remainder of this chapter is arranged as follows: we provide a brief review on the extant literature in Section 3.2, define the problem setting in Section 3.3, analyze and provide the structural results for the no-visibility case in Section 3.4, formulate the visibility case in Section 3.5, present the numerical results and empirical evidence in Sections 3.6 and 3.7, and conclude in Section 3.8.

3.2 Literature Review

In this section, we review the extant literature on (i) Embedded-MDP Approaches to Queue Control, (ii) Queues with Customer Heterogeneity, and (iii) Queues with Lookahead Visibility.

Embedded-MDP Approaches to Queue Control Embedded MDP formulations are helpful when service times are not exponentially distributed, e.g., deterministic or general distributed service times. In these cases, the residual processing time breaks the Markov property. By restricting attention to decision epochs (arrivals, departures, or inspection points), the residual clock resets, allowing the dynamic-programming formulation to be applied without increasing the state space. Early contributions show that simple state-dependent thresholds are optimal in single-server settings. In a seminal study, Balachandran [1973] formulates on/off control for an M/G/1 queue and shows the average-cost optimality of a queue-length threshold. Heyman [1977] studies the T-policy for the $M/G/1$ queue, where the server idles until the workload exceeds a time threshold T. The author shows the existence of a cost-optimal T and proposes thresholds as a control mechanism.

Koole [1998] proposes event-based dynamic programming, reformulating the Bellman equations so that monotonicity and convexity results for individual arrival and service events can be used to get optimal control policies. The author shows that simple threshold policies are optimal in a wide range of service control problems. Lee et al. [2006] examine an $M/G/1$ queue that combines D-policies (delay-until-threshold) with multiple vacation models. They propose new transforms for queue length and waiting time, showing how vacation parameters shift the optimal delay threshold. Liu and Whitt [2011] focus on an overloaded $G/D/S{+}GI$ system with deterministic service times. The authors show how deterministic service time amplifies system congestion under heavy load.

Pacheco et al. [2017] study fixed-cycle traffic lights as an $M/G/1$ queue with server vacations. The authors derive the average vehicle delay in order to tune signal timing and verify it using a detailed simulation study. Lei and Jasin [2020] examine revenue management for reusable resources with deterministic service times (e.g., car-sharing) and, via a fluid approximation, show that a simple state-dependent threshold pricing policy is near-optimal.

Queues with Customer Heterogeneity Queueing systems that account for customer heterogeneity and priority structures address the challenges of resource allocation in environments with diverse service needs. Customer heterogeneity can be of multiple types, such as differences in arrival rates, service times, abandonment costs, and priorities, etc. Bell [1973] studies removable-server queues, showing how server activation policies based on customer priorities can minimize operational costs. Similarly, Tijms [1974] introduces control policies for priority queues with removable servers, balancing idleness and service prioritization under stochastic demand. Harrison [1975a] and Harrison [1975b] explore priority queue models with linear costs and dynamic scheduling, identifying conditions where static priority rules optimize outcomes across customer classes with varying service times and holding costs. Lerzan Örmeci et al. [2001] study two-class loss systems in telecommunications and show that dynamic threshold-based policies effectively maximize revenue under congestion. Green et al. [2006] propose dynamic priority rules for emergency and outpatient service, balancing real-time decisions with revenue and clinical outcomes. Patrick et al. [2008] model the assignment of CT-scanner slots among several urgency classes of patients as an infinite-horizon Markov decision process, with an objective of balancing patient wait-time targets against overtime or diversions. Fatnes [2010], in the railway scheduling context, proposes combining preemptive and non-preemptive priority disciplines to reduce train waiting times on single-track lines.

Kim and Ward [2013] address multiclass queues with abandonment, and formulate dynamic control rules that balance abandonment costs with system efficiency. Wolvers [2020] studies

triage inefficiencies during the COVID-19 pandemic, proposing state-based scheduling policies that adapt to patient deterioration. Hu et al. [2022] analyze a queue in which customers begin as regular but can escalate to urgent status the longer they wait. The authors prove that prioritizing the class with the greatest marginal cost-saving per minute is optimal. Collectively, these studies show that combining customer heterogeneity with adaptive priority policies improves system-wide performance.

Queues with Lookahead Visibility Visibility allows a queue controller to have a lookahead of upcoming arrivals, enabling better scheduling and more efficient use of resources. Spencer et al. [2014] study a system where a controller can reject jobs to keep the queue from getting too long. The authors show that if the controller relies only on the current state, the queue can grow without bound as the system gets congested. In contrast, a controller who has visibility to all future arrivals can keep waiting time low by rejecting those jobs that would never be served. Most importantly, the authors prove that the controller does not need full lookahead: a short visibility is almost as helpful as knowing the whole future arrival stream in order to keep the queue length under control. Building on this study, Xu [2015] quantifies the smallest lookahead required to keep the rejection queue under control. The author shows that if the look-ahead window is shorter than a given factor, the queue will still blow up as traffic nears capacity.

Xu and Chan [2016] examine an overcrowded emergency-department queue where the controller may divert incoming patients to other facilities. The authors propose a proactive diversion rule that uses a rolling forecast of arrivals: as soon as the predicted inflow exceeds capacity, the controller begins diverting, instead of waiting for the queue to get full. They prove that with a finite lookahead window, the rule outperforms the best static threshold policy. Ata and Peng [2020] study a call-center with a single server that can offer callers a call-back option when the line is busy. The authors show that a simple lookahead policy that looks into the future arrivals and service completion times for the next p/h time units (waiting cost of h per unit time and one-time penalty of p) and

uses the information on the current number of customers in the system who previously rejected a callback offer is pathwise optimal. This simple 'p/h-look-ahead' policy matches the performance of a planner with perfect foresight and relies only on the total queue length and a brief forecast. Chen et al. [2025] study a two-class parallel-server system that faces temporary demand surges and analyze how short-term forecasts of arrival rates affect routing decisions between primary and secondary servers. The authors propose a two-stage index-based look-ahead policy that weighs holding versus overflow penalties by estimating how long each queue would take to empty under the forecasted load.

3.3 Queue Control Problem

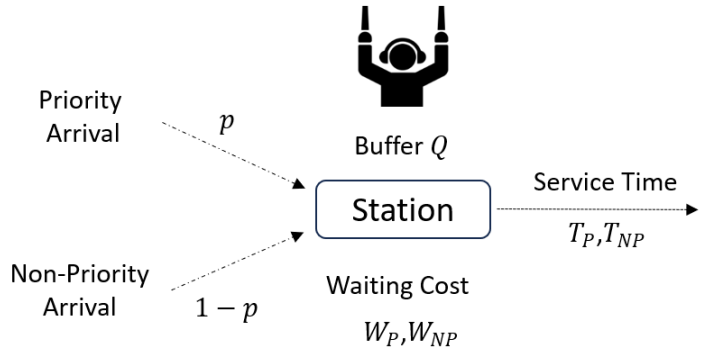

Figure 3.1: Finite Buffer Queue.

We model a finite-buffer priority queueing system with two customer classes: priority and non-priority. We consider deterministic service times, which are customer-type dependent. We assume that (i) the arrival process is a right-1 truncated Poisson process, and thus, during any service process, at most one new customer may arrive, and (ii) waiting costs accrue linearly with the number of queued customers and differ by customer type.[4] The objective of the controller is to minimize the waiting cost. The actions available for the section controller are to either hold a

[4]For this study, we define priority purely based on waiting cost. Hence, the priority class is the one with the higher per-unit-time waiting cost.

customer or serve it.

We formulate the problem both as a long-run average cost MDP and an infinite horizon discounted cost MDP. Our analysis considers two cases. In the first case, we assume no visibility, where the controller doesn't have information about the type and time of the next arrival. In the no-visibility case, we find that the optimal policy under the long-run average cost formulation is to always *serve* the non-priority customer when it is feasible. Under the infinite-horizon discounted cost formulation, we characterize the optimal policy and specify the conditions under which a threshold policy is optimal. The policy is to hold the non-priority customers up to a specified threshold in anticipation of priority customers.

In the second case, we introduce visibility, where the controller has perfect information about the type of the next arrival (priority or non-priority) but imperfect information about the time of the next arrival. In the simulation study, we find that the optimal control policy is defined by a monotone switching curve in the following two variables: the number of non-priority customers waiting at the queue and the expected time to the next priority arrival. The curve determines the transition from hold to serve as the optimal action.

3.4 No Visibility Case

In this section, we study the no-visibility case (no information about time to next arrival and type of next arrival) for both the long-run average cost and the infinite-horizon discounted cost formulation.[5] We assume that arrivals follow a right–1 truncated Poisson process with rate λ and the queue has a finite buffer of size Q. Service times are deterministic: priority customers require T_P time units, whereas non-priority customers require T_{NP} time units. An arriving customer is a priority type with probability p (and a non-priority type with probability $1-p$). Holding costs accumulate at a rate of W_p per unit time for each priority customer and W_{np} per unit time for each

[5]It is important to consider both the long run average cost and the infinite horizon discounted cost formulations because each captures a different operational perspective, as summarized in the Table 3.1.

non-priority customer. In the infinite-horizon formulation, we assume that the costs are discounted exponentially in continuous time at rate γ.

	Long Run Average Cost	Infinite Horizon Discounted Cost
Cost Interpretation	Per unit time (Insensitive to when cost occurs)	Present value of cumulative cost (Sensitive; prioritizes near-term cost reductions)
Managerial Relevance	Capacity planning	Real time control

Table 3.1: Long Run Average Cost and Discounted Cost Formulations.

Decision Epochs & Subcycles

We define the decision epochs as follows: an epoch is an instance when there are no priority customers waiting and the server is free. Let y be the number of non-priority customers waiting to get serviced. We define a cycle as the time between two decision epochs, and the subcycles based on the number of priority customers that are served in a cycle. By defining the cycles in this way, we reduce the state space to a single variable: y. The available actions in any state y are $Hold$ or $Serve$ a non-priority customer. π_p and $\pi_{np} > \pi_p$ are the corresponding probabilities that an arrival happens during the interval T_P, and T_{NP} respectively. $c_h(y) = \alpha_h + \delta_h * y$ is the expected waiting cost for non-priority customers when the $Hold$ decision is made. Similarly, $c_s(y) = \alpha_s + \delta_s * (y-1)$ is the expected waiting cost for non-priority customers when the $Serve$ decision is made.[6] [7] In the cycle initiated by *Hold* decision, the one-step transition probabilities are $p_h^\circ = \Pr\{y \to y\}$ and $p_h^+ = \Pr\{y \to y+1\}$. In the cycle initiated by *Serve* decision, $p_s^- = \Pr\{y \to y-1\}$ and $p_s^\circ = \Pr\{y \to y\}$.

3.4.1 Long Run Average Cost Formulation

We set up the problem as a long-run average cost semi-Markov decision process (SMDP). Let $h(y)$ be the relative cost of being in state y and g be the long-term average cost per unit time. If

[6] α_h and α_s are the expected waiting costs of newly arrived non-priority customers during the cycles starting with $Hold$ and $Serve$ decisions respectively.

[7] $\delta_h = W_{np} * \tau_h$, and $\delta_s = W_{np} * \tau_s$. Here τ_h, and τ_s are the expected durations of the cycle starting with the $Hold$ and the $Serve$ decisions.

Service type	Priority arrival	Non-priority arrival	No arrival
Priority	$\pi_p\, p$	$\pi_p\,(1-p)$	$1-\pi_p$
Non-priority	$\pi_{np}\, p$	$\pi_{np}\,(1-p)$	$1-\pi_{np}$

Table 3.2: Arrival Probabilities During a Service.

(P-Priority, NP-Non Priority)	Hold	Serve
Expected cycle duration	τ_h	τ_s
Transition probabilities		
$\Delta y = +1$	p_h^+	0
$\Delta y = \ \ 0$	p_h°	p_s°
$\Delta y = -1$	0	p_s^-
Cost coefficients		
New NP customer arrived during the cycle	α_h	α_s
New P customer arrived during the cycle	β_h	β_s
Existing NP customers	δ_h	δ_s
One-step cost	$\alpha_h + \delta_h\, y+ \beta_h$	$\alpha_s + \delta_s\,(y-1)+ \beta_s$

Table 3.3: Cost Coefficients and Transition Probabilities for $Hold$ and $Serve$ Decisions.

the optimal decision in state y is to $Hold$ the non-priority customers, then following Theorem 1 of Das et al. [1999], we can write:

$$h(y) + g\tau_h = \alpha_h + \delta_h y + \beta_h + p_h^\circ h(y) + p_h^+ h(y+1) \tag{3.1}$$

If the optimal decision in state y is to $Serve$ a non-priority customer, then:

$$h(y) + g\tau_s = \alpha_s + \delta_s(y-1) + \beta_s + p_s^- h(y-1) + p_s^\circ h(y) \tag{3.2}$$

Equivalently, the optimality equation can be written as:

$$\begin{aligned} h(y) = \min\{&\alpha_h + \delta_h y + \beta_h + p_h^\circ h(y) + p_h^+ h(y+1) - g\tau_h, \\ &\alpha_s + \delta_s(y-1) + \beta_s + p_s^- h(y-1) + p_s^\circ h(y) - g\tau_s\} \end{aligned} \tag{3.3}$$

In the following sections, we derive the expressions for transition probabilities (p_h°, p_h^+, p_s^- and p_s°) and cost coefficients ($\alpha_h, \delta_h, \alpha_s, \delta_s, \beta_h$ and β_s) used in the equation 3.3.

Transition Probabilities

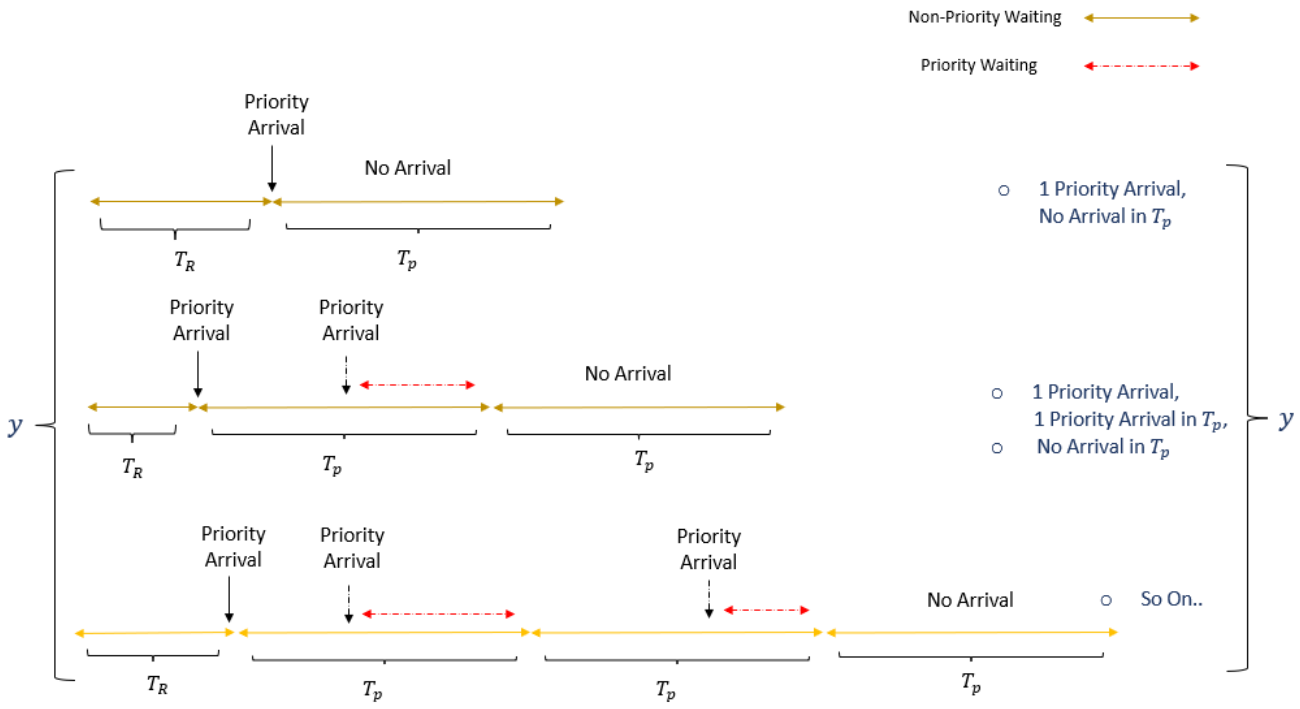

Figure 3.2: Illustration of Possible Sub-cycles in $y \rightarrow y$ Transition During Hold Cycle.

To compute the transition probabilities under a $Hold$ decision, we begin with the case $y \rightarrow y$. As defined earlier, p_h° is the probability of remaining in state y. This occurs when a sequence of consecutive priority arrivals is eventually followed by a period with no arrival (Figure 3.2). The initial subcycles are:

- One priority arrival occurs, followed by a no-arrival interval of length T_p; probability $p(1 - \pi_p)$.
- Two consecutive priority arrivals occur, followed by a no-arrival interval of length T_p; probability $p\pi_p p(1 - \pi_p)$.

In general, the scenario with $n+1$ consecutive priority arrivals followed by no arrival has probability $(p\pi_p)^n\, p(1-\pi_p)$. Thus,

$$P(y \to y, Hold) = p(1-\pi_p) \sum_{n=0}^{\infty} (p\pi_p)^n \tag{3.4}$$

Simplifying, we get:

$$p_h^\circ = P(y \to y, Hold) = \frac{p(1-\pi_p)}{1-p\pi_p} \tag{3.5}$$

p_h^+ represents the probability of moving to state $y+1$ if a $Hold$ decision is made, which occurs if there is a priority customer arrival or continuous arrivals of priority customers followed by a non-priority customer arrival. It involves summing the probabilities of successive priority customer arrivals followed by a non-priority arrival, forming a geometric series.

$$p_h^+ = 1 - p_h^\circ = 1 - \frac{p(1-\pi_p)}{1-p\pi_p} \tag{3.6}$$

Similarly, p_s^- represents the probability of moving to state $y-1$ if a serve decision is made, which occurs if there is no arrival in T_{np} or priority arrival in T_{np} followed by no arrivals. It involves summing up such probabilities. The initial subcycles are:

- No arrival in the interval of length T_{np} ; probability $(1-\pi_{np})$.
- One priority customer arrivals occur, followed by a no arrival interval of length T_p; probability $\pi_{np} p(1-\pi_p)$.

In general, the scenario with one priority arrival in T_{np}, followed by n additional priority arrivals in successive intervals of length T_p, and then a no-arrival window, has probability $p\pi_{np}(1-\pi_p)(p\pi_p)^n$. The series can be summed as:

$$P(y \to y-1, Serve) = (1-\pi_{np}) + p\pi_{np}(1-\pi_p) \sum_{n=0}^{\infty} (p\pi_p)^n \tag{3.7}$$

Simplifying, we get:

$$p_s^- = P(y \to y-1, Serve) = (1-\pi_{np}) + \frac{\pi_{np} p (1-\pi_p)}{1-p\pi_p} \tag{3.8}$$

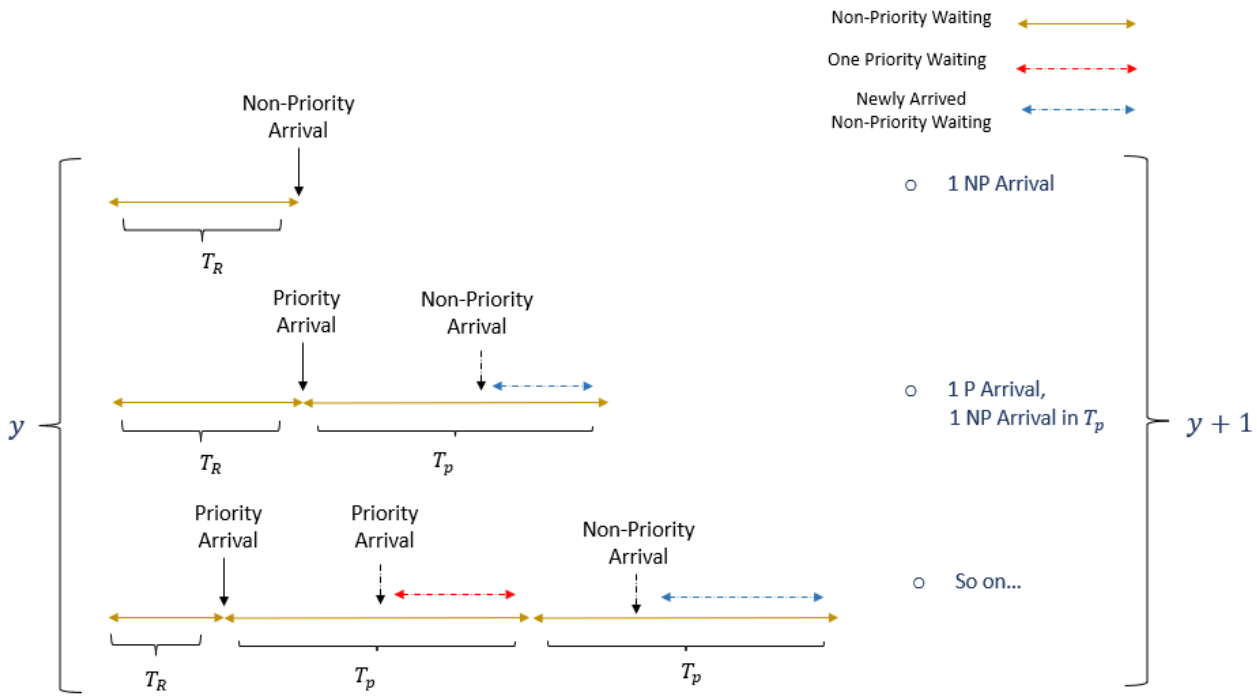

Figure 3.3: Illustration of Possible Sub-cycles in $y \to y+1$ Transition During Hold Cycle .

Similarly, let p_s° represent the probability of moving to state y if a $Serve$ decision is made, which occurs if there is one non-priority arrival in the interval of length T_{np} or priority arrival in the interval of length T_{np} followed by a non-priority arrival. It involves summing up such probabilities, forming a geometric series.

$$p_s^\circ = 1 - p_s^- = 1 - ((1-\pi_{np}) + \frac{\pi_{np} p (1-\pi_p)}{1-p\pi_p}) \tag{3.9}$$

Cost Coefficients

The cost coefficients depend on wait-time expressions that are functions of the problem parameters. For clarity of exposition, we introduce the following notation: $WT_u = 1/\lambda$ denotes the uncondi-

tional expected time to the first arrival in a $Hold$ cycle. WT_p denotes the expected wait time for an arrival occurring within a P-serve cycle of length T_p to the completion of that cycle, and WT_{np} is defined analogously for an NP-serve cycle of length T_{np}. Using these quantities, and aggregating over the subcycles described in Section 3.4.1, we get the expressions for $\alpha_h, \delta_h, \alpha_s, \delta_s, \beta_h$and β_s as follows:

$$\alpha_h = \frac{p\pi_p(1-p)W_{np}WT_p}{(1-p\pi_p)} \tag{3.10}$$

$$\delta_h = W_{np}\left(WT_u + T_p\frac{p}{1-p\pi_p}\right) \tag{3.11}$$

$$\alpha_s = \frac{W_{np}WT_{np}\pi_{np}(1-p)}{1-p\pi_p} \tag{3.12}$$

$$\delta_s = W_{np}\left[T_{np} + \frac{p\pi_{np}T_p}{1-p\pi_p}\right] \tag{3.13}$$

$$\beta_h = \frac{p^2\pi_p W_p WT_p}{1-p\pi_p} \tag{3.14}$$

$$\beta_s = W_p(p\pi_{np}WT_{np} + \frac{p^2\pi_p\pi_{np}WT_p}{1-p\pi_p}) \tag{3.15}$$

For the remainder of this work, we use these terms directly, rather than their explicit expressions in terms of problem parameters. We consider the following boundary conditions:

$$h(0) = \alpha_h + \beta_h + p_h^{\circ}h(0) + p_h^{+}h(1) - g\tau_h \tag{3.16}$$

$$h(Q) = \alpha_s + \delta_s(Q-1) + \beta_s + p_s^{-}h(Q-1) + p_s^{\circ}h(Q) - g\tau_s \tag{3.17}$$

For analyzing the structural properties, we use the general structure given in equation 3.3. Let us define the second difference operator as follows:

$$\Delta^2 h(y) \;:=\; h(y+1) \;+\; h(y-1) \;-\; 2\,h(y) \tag{3.18}$$

We introduce the following two functions to compare $h(y)$ in $Hold$ vs. $Serve$ scenarios:

$$\begin{aligned} l_0(y,g) &:= \frac{c_h(y)+\beta_h-g\tau_h}{p_h^+} = \frac{\alpha_h+\delta_h\, y+\beta_h-g\tau_h}{p_h^+}, \\ l_1(y,g) &:= \frac{c_s(y+1)+\beta_s-g\tau_s}{p_s^-} = \frac{\alpha_s+\delta_s\, y+\beta_s-g\tau_r}{p_s^-}. \end{aligned}$$

Proposition 1. The relative value function $h(.)$ satisfies:

$$h(y) \;\le\; l_0(y,g) \;+\; h(y+1), \quad h(y) \;\le\; l_1(y-1,g) \;+\; h(y-1), \forall y.$$

Proof.

We consider the two cases separately, when $Hold$ or $Serve$ is optimal in state y, respectively.

Case 1: If $Hold$ is optimal in state y, then

$$h(y) = \alpha_h + \delta_h y + \beta_h + p_h^\circ h(y) + p_h^+ h(y+1) - g\tau_h \tag{3.19}$$

As $p_h^\circ + p_h^+ = 1$,

$$h(y) = \frac{\alpha_h + \delta_h y + \beta_h - g\tau_h}{p_h^+} + h(y+1) = l_0(y,g) + h(y+1) \tag{3.20}$$

Case 2: If $Serve$ is optimal in state y, then

$$h(y) = \alpha_s + \delta_s(y-1) + \beta_s + p_s^- h(y-1) + p_s^\circ h(y) - g\tau_s \tag{3.21}$$

As $p_s^- + p_s^\circ = 1$,

$$h(y) = \frac{\alpha_s + \delta_s(y-1) + \beta_s - g\tau_s}{p_s^-} + h(y-1) = l_1(y-1,g) + h(y-1) \tag{3.22}$$

From equations 3.3, 3.20 and 3.22, we get

$$h(y) \leq l_0(y, g) + h(y+1), \quad h(y) \leq l_1(y-1, g) + h(y-1) \forall y.$$

Furthermore, equality holds in the first (respectively second) inequality when y is a $Hold$ (respectively $Serve$) state in the optimal policy.

□

Lemma 1. Let h be a solution to 3.3. For any state i with $0 < i < Q$:

(a) If i is a $Hold$ state, then

$$\Delta^2 h(i) \ \leq \ -\frac{\delta_h}{p_h^+},$$

and if $i-1$ is a $Hold$ state, then

$$\Delta^2 h(i) \ \geq \ -\frac{\delta_h}{p_h^+}.$$

Consequently, if both $i-1$ and i are $Hold$ states, then

$$\Delta^2 h(i) \ = \ -\frac{\delta_h}{p_h^+}.$$

(b) If i is a $Serve$ state, then

$$\Delta^2 h(i) \ \leq \ \frac{\delta_s}{p_s^-},$$

and if $i+1$ is a $Serve$ state, then

$$\Delta^2 h(i) \ \geq \ \frac{\delta_s}{p_s^-}.$$

Consequently, if both i and $i+1$ are $Serve$ states, then

$$\Delta^2 h(i) \;=\; \frac{\delta_s}{p_s^-}.$$

(c) If $i \geq 0$ is $Hold$ and $i+1 \leq Q$ is $Serve$ state, then

$$l_0\Big(i, g\Big) \;=\; -\,l_1\Big(i, g\Big).$$

Furthermore, from (a) and (b)

$$i > 0 \implies -\frac{\delta_h}{p_h^+} \;\geq\; \frac{\delta_s}{p_s^-}, \quad i+1 < Q \implies -\frac{\delta_h}{p_h^+} \;\leq\; \frac{\delta_s}{p_s^-}.$$

Proof.

(a) By applying Proposition 1 at $y = i$, if i is a $Hold$ state, then

$$h(i) = l_0\Big(i, g\Big) + h(i+1).$$

Similarly, applying Proposition 1 at $y = i-1$ yields

$$h(i-1) \;\leq\; l_0\Big(i-1, g\Big) + h(i).$$

Subtracting the second from the first and rearranging gives us:

$$\Delta^2 h(i) \;=\; h(i+1) + h(i-1) - 2\,h(i) \;\leq\; l_0(i-1, g) - l_0(i, g) \;=\; -\frac{\delta_h}{p_h^+}.$$

If $i-1$ is also a $Hold$ state, the roles of the equality and inequality are reversed, resulting in the opposite inequality.

(b) When i is a $Serve$ state, from Proposition 1 we get

$$h(i) = l_1(i-1, g) + h(i-1).$$

Also for $i+1$,

$$h(i+1) \;\leq\; l_1\Big(i, g\Big) + h(i).$$

Subtracting these and simplifying yields

$$\Delta^2 h(i) \;\leq\; l_1(i, g) - l_1(i-1, g) \;=\; \frac{\delta_s}{p_s^-}.$$

As before, if $i+1$ is a $Serve$ state, equality holds, giving $\Delta^2 h(i) = \dfrac{\delta_s}{p_s^-}$.

(c) If i is a $Hold$ state and $i+1$ is a $Serve$ state, then Proposition 1 gives

$$l_0(i, g) + h(i+1) \;=\; h(i), \quad l_1(i, g) + h(i) \;=\; h(i+1).$$

Combining these, we get $l_0(i, g) = -\,l_1(i, g)$. For $i > 0$, using (a) and (b), we get

$$\frac{\delta_s}{p_s^-} \leq \Delta^2 h(i) \;\leq\; -\frac{\delta_h}{p_h^+}.$$

Similarly, if $i + 1 < Q$

$$\frac{\delta_s}{p_s^-} \geq \Delta^2 h(i+1) \;\geq\; -\frac{\delta_h}{p_h^+}.$$

□

Theorem 1 (Optimal Policy for Long Run Average Cost Formulation)**.** Let $\alpha_h, \delta_h, \alpha_s, \delta_s$ and β_h, β_s be the cost coefficients and $p_h^+ = 1 - p_h^\circ$, $p_s^- > 0$ be the transition parameters as defined previously. Then the optimal long-run average cost g and the corresponding optimal policy are determined by the following cases:

1. **Case 1:** $-\dfrac{\delta_h}{p_h^+} = \dfrac{\delta_s}{p_s^-}$.

   All policies are equivalent/optimal and

$$g^* \;=\; \frac{p_s^-\Big(\alpha_h + \beta_h\Big) \;+\; p_h^+\Big(\alpha_s + \beta_s\Big)}{p_h^+\tau_s + p_s^-\tau_h}.$$

2. **Case 2:** $-\dfrac{\delta_h}{p_h^+} < \dfrac{\delta_s}{p_s^-}$.

   It is optimal to switch to $Serve$ at $y = 1$, giving

$$g^* \;=\; \frac{p_s^-\Big(\alpha_h + \beta_h\Big) \;+\; p_h^+\Big(\alpha_s + \beta_s\Big)}{p_h^+\tau_s + p_s^-\tau_h}.$$

3. **Case 3:** $-\dfrac{\delta_h}{p_h^+} > \dfrac{\delta_s}{p_s^-}$.

   It is optimal to $Hold$ until $y = Q - 1$ and switch to $Serve$ at $y = Q$, giving

$$g^* \;=\; \frac{p_s^-\Big(\alpha_h + \beta_h + \delta_h\,(Q-1)\Big) \;+\; p_h^+\Big(\alpha_s + \beta_s + \delta_s\,(Q-1)\Big)}{p_h^+\tau_s + p_s^-\tau_h}.$$

Proof.

Recall from Lemma 1 (c) that a $Hold$ to $Serve$ switch at state i must satisfy

$$\ell_0(i,g) + \ell_1(i,g) \;=\; 0.$$

We analyze this equation by comparing $-\delta_h/p_h^+$ and δ_s/p_s^-.

Case 1: $-\dfrac{\delta_h}{p_h^+} = \dfrac{\delta_s}{p_s^-}$. In this case, $\ell_0 + \ell_1$ depends linearly on y with a zero slope. Since the boundary conditions guarantee that a switch between actions must occur, the switching condition $\ell_0(y,g) + \ell_1(y,g) = 0$ must be satisfied at some state. However, when $\ell_0(y,g) + \ell_1(y,g)$ is independent of y, this condition will hold for all y. As a result, both the $Hold$ and $Serve$ actions

yield the same cost at every state, making every policy optimal. Thus, all policies are equivalent and optimal in this case. That gives

$$g \;=\; \frac{p_s^-(\alpha_h+\beta_h)+p_h^+(\alpha_s+\beta_s)}{p_h^+\tau_s+p_s^-\tau_h}.$$

Case 2: $-\dfrac{\delta_h}{p_h^+} \;<\; \dfrac{\delta_s}{p_s^-}$. Lemma 1(c) forces the $Hold$ to $Serve$ switch to occur at $i=0$. Specifically, if the switch were anywhere else, Lemma 1(c) would imply $-\delta_h/p_h^+ \;\geq\; \delta_s/p_s^-$, contradicting our assumption. Hence, we set $i=0$ in $\ell_0(i,g)+\ell_1(i,g)=0$ and solve for g. That gives

$$g \;=\; \frac{p_s^-\left(\alpha_h+\beta_h\right)+p_h^+\left(\alpha_s+\beta_s\right)}{p_h^+\tau_s+p_s^-\tau_h}.$$

Thus we $Hold$ only at $y=0$ and switch to $Serve$ immediately for $y\geq 1$.

Case 3: $-\dfrac{\delta_h}{p_h^+} \;>\; \dfrac{\delta_s}{p_s^-}$. By a symmetric argument, Lemma 1(c) forces the unique switch to be at $i=Q-1$, as having the switch earlier will result in a contradiction. We solve $\ell_0(Q-1,g)+\ell_1(Q-1,g) \;=\; 0$ to get g. That gives

$$g \;=\; \frac{p_s^-\Big(\alpha_h+\beta_s+\delta_h\left(Q-1\right)\Big) \;+\; p_h^+\Big(\alpha_s+\beta_s+\delta_s\left(Q-1\right)\Big)}{p_h^+\tau_s+p_s^-\tau_h}.$$

Hence, it is optimal to $Hold$ up to $y=Q-1$ and then $Serve$ at $y=Q$. □

Corollary 1. As in our setting, $\delta_h, \delta_s, p_h^+$ and p_s^- are all positive, the optimal policy is to always $Serve$, and the long run average cost is given by:

$$g \;=\; \frac{p_s^-\left(\alpha_h+\beta_h\right)+p_h^+\left(\alpha_s+\beta_s\right)}{p_h^+\tau_s+g\tau_h}.$$

3.4.2 Infinite Horizon Discounted Cost Formulation

In this section, we study the infinite-horizon discounted cost formulation for the problem, assuming continuous-time discounting at rate γ. Let $V(y)$ be the optimal discounted cost-to-go function if the system is currently in state y. Depending on which action is optimal in a given state y, we get:

$$\begin{aligned} V(y) = min\{&c_h(y) + \beta_h + p_h^\circ V(y) + p_h^+ V(y+1), \\ &c_s(y) + \beta_s + p_s^- V(y-1) + p_s^\circ V(y)\} \end{aligned} \tag{3.23}$$

Here p_h°, p_h^+, p_s^- and p_s° are coefficients subsuming both discount and transition probabilities, given by:

$$p_h^\circ = \frac{\lambda}{\lambda+\gamma} \frac{p(1-\pi_p)e^{-\gamma T_p}}{(1-p\pi_p e^{-\gamma T_p})} \tag{3.24}$$

$$p_h^+ = \frac{\lambda}{\lambda+\gamma} \left((1-p) + \frac{p\pi_p(1-p)e^{-\gamma T_p}}{(1-p\pi_p e^{-\gamma T_p})} \right) \tag{3.25}$$

$$p_s^- = (1-\pi_{np})e^{-\gamma T_{np}} + \frac{p\pi_{np}(1-\pi_p)e^{-\gamma(T_{np}+T_p)}}{(1-p\pi_p e^{-\gamma T_p})} \tag{3.26}$$

$$p_s^\circ = (1-p)\pi_{np}e^{-\gamma T_{np}} + \frac{p\pi_p\pi_{np}(1-p)e^{-\gamma(T_{np}+T_p)}}{(1-p\pi_p e^{-\gamma T_p})} \tag{3.27}$$

Refer to the Appendix 3.9.1 for the expressions for $\alpha_h, \delta_h, \alpha_s, \delta_s, \beta_h$and β_s. We first analyze the recurrence (contiguous states with the same optimal action) relations that govern the value function and derive conditions under which the controller switches from $Hold$ decision to $Serve$ decision (and vice versa).

Analysis of the Holding Recurrence

When the $Hold$ decision is optimal at state y, the value function satisfies

$$V(y) \;=\; c_h(y) \;+\beta_h+\; p_h^\circ\, V(y) \;+\; p_h^+\, V(y+1) \tag{3.28}$$

Rearranging,

$$V(y+1) \;=\; \frac{1-p_h^\circ}{p_h^+}\,V(y) \;-\; \frac{1}{p_h^+}\,(c_h(y)+\beta_h) \tag{3.29}$$

Let us denote $\alpha \;:=\; \dfrac{1-p_h^\circ}{p_h^+}$.

Case 1: $p_h^\circ + p_h^+ < 1$**.**

If $p_h^\circ + p_h^+ < 1$, then $\alpha = \frac{1-p_h^\circ}{p_h^+} > 1$. The general solution to this linear recurrence is of the form

$$V(y) \;=\; A_1\,y \;+\; B_1 \;+\; C_1\,\alpha^{y-y_0} \tag{3.30}$$

The constants A_1, B_1 and C_1 are identified by boundary conditions (e.g., $V(0)$ or $V(y_0)$). Specifically:

$$\begin{aligned}
A_1 \;&=\; \frac{\delta_h}{(\,\alpha-1\,)\,p_h^+} \;=\; \frac{\delta_h}{1-p_h^\circ-p_h^+},\\
B_1 \;&=\; \frac{\alpha_h+\beta_h}{(\,\alpha-1\,)\,p_h^+} \;+\; \frac{A_1}{\alpha-1} = \frac{\alpha_h+\beta_h}{1-p_h^\circ-p_h^+} \;+\; \frac{\delta_h p_h^+}{(\,1-p_h^\circ-p_h^+\,)^2},\\
C_1 \;&=\; \frac{V(y_0)-A_1\,y_0-B_1}{\alpha^{\,y_0}}.
\end{aligned}$$

Moreover, since the recurrence relation links $V(y+1)$ to $V(y)$ for every state y where holding is optimal, the above expression remains valid up to the first y_{First_Serve} at which serving becomes optimal. Therefore, we can determine C_1 using:

$$C_1 \;=\; \frac{V(y_{First\_Serve})-A_1 y_{First\_Serve}-B_1}{\alpha^{y_{First\_Serve}-y_0}}$$

Case 2: $p_h^\circ + p_h^+ = 1$**.**

If $p_h^\circ + p_h^+ = 1$, then $\alpha = 1$, making the recurrence a difference equation with no exponential term. In this case, we can show that

$$V(y) \;=\; A_1' \, y^2 \;+\; B_1' \, y \;+\; C_1' \tag{3.31}$$

Where,

$$A_1' \;=\; -\frac{\delta_h}{2\, p_h^+}, \quad B_1' \;=\; \frac{\delta_h}{2\, p_h^+} \;-\; \frac{\alpha_h}{p_h^+}, \quad C_1' \;=\; V(y_0) - A_1' \, y_0^2 - B_1' \, y_0.$$

Analysis of the Serving Recurrence

When the $Serve$ decision is optimal at state y, we have

$$V(y) \;=\; c_s(y) \;+ \beta_s + \; p_s^- \, V(y-1) \;+\; p_s^\circ \, V(y) \tag{3.32}$$

Rearrange to get $V(y)$:

$$V(y) = \frac{p_s^-}{1 - p_s^\circ} \, V(y-1) \;+\; \frac{1}{1 - p_s^\circ} (c_s(y) + \beta_s) \tag{3.33}$$

Let us denote $\beta \;:=\; \dfrac{p_s^-}{1 - p_s^\circ} \leq 1$ as $p_s^- + p_s^\circ \leq 1$.

Case 1: $p_s^- + p_s^\circ < 1$**.**

Then $\beta = \dfrac{p_s^-}{1 - p_s^\circ} < 1$. The solution of the corresponding recurrence takes the following analogous form:

$$V(y) \;=\; A_0 y \;+\; B_0 \;+\; C_0 \, \beta^{\, y} \tag{3.34}$$

We determine A_0, B_0 and C_0 from the recurrence relation and boundary conditions on the serving recurrence. Specifically:

$$\begin{aligned} A_0 &= \frac{\delta_s}{(1-\beta)(1-p_s^\circ)} = \frac{\delta_s}{1-p_s^- - p_s^\circ}, \\ B_0 &= \frac{\alpha_s + \beta_s}{(1-\beta)(1-p_s^\circ)} - \frac{A_0}{1-\beta} = \frac{\alpha_s+\beta_s}{1-p_s^- - p_s^\circ} - \frac{\delta_s(1-p_s^\circ)}{(1-p_s^- - p_s^\circ)^2}. \end{aligned}$$

Since this recurrence relates $V(y)$ to $V(y-1)$, the coefficient C_0 is determined from the last holding state y_{Last_Hold} before serving becomes optimal:

$$C_0 = \frac{V(y_{\text{Last\_Hold}}) - A_0 y_{\text{Last\_Hold}} - B_0}{\beta^{y_{\text{Last\_Hold}}}}.$$

Case 2: $p_s^- + p_s^\circ = 1$**.**

In this special case, $\beta = 1$, and the recurrence again becomes a difference equation that yields a quadratic solution in y.

$$V(y) = A_0' y^2 + B_0' y + C_0' \tag{3.35}$$

Where,

$$\begin{aligned} A_1' &= -\frac{\delta_h}{2\,p_h^+}, \quad B_1' = \frac{\delta_h}{2\,p_h^+} - \frac{\alpha_h}{p_h^+}, \\ C_1' &= V(y_{Last\_Hold}) - A_1' (y_{Last\_Hold})^2 - B_1' y_{Last\_Hold}. \end{aligned}$$

The recurrence relations derived above yield the following structural property of the value function:

Proposition 2. Consider any contiguous range of states $[i, j]$ in which the optimal strategy remains constant (i.e., it is either $Hold$ for all states in $[i, j]$ or $Serve$ for all states). Then, on that interval, the value function $V(\cdot)$ must be either entirely convex or entirely concave.

Switching Analysis

Given that our boundary conditions force holding at $y = 0$ and serving at $y = Q$, there is guaranteed to be at least one switch from $Hold$ to $Serve$. Furthermore, there can only be an odd number of such switches. Specifically:

- A single switch: $Hold$ to $Serve$.
- Three switches: $Hold$ to $Serve$ to $Hold$ to $Serve$ and so on.

However, in our setting, we will see that only one or three switches are possible.

Switching from Holding to Serving

Suppose the optimal decision is to switch from $Hold$ at k to $Serve$ at $k+1$. In that case, the value function must satisfy the following equations:

$$V(k) = c_h(k) + \beta_h \; + \; p_h^{\circ} \, V(k) \; + \; p_h^{+} \, V(k+1) \tag{3.36}$$

$$V(k+1) = c_s(k+1) \; + \beta_s + \; p_s^{-} \, V(k) \; + \; p_s^{\circ} \, V(k+1) \tag{3.37}$$

Rewriting this system and solving for $V(k)$ and $V(k+1)$ yields:

$$V(k) = \frac{(1 - p_s^{\circ}) \, (c_h(k) + \beta_h) + p_h^{+} \, (c_s(k+1) + \beta_s)}{\Big(1 - p_h^{\circ}\Big)\Big(1 - p_s^{\circ}\Big) \; - \; p_h^{+} \, p_s^{-}} \tag{3.38}$$

$$V(k+1) = \frac{p_s^{-} \, (c_h(k) + \beta_h) + \Big(1 - p_h^{\circ}\Big) \, (c_s(k+1) + \beta_s)}{\Big(1 - p_h^{\circ}\Big)\Big(1 - p_s^{\circ}\Big) \; - \; p_h^{+} \, p_s^{-}} \tag{3.39}$$

Thus, the values $V(k)$ and $V(k+1)$ are uniquely determined at the switching point. Notice that both expressions are linear in k, and this will be crucial for identifying optimal strategies. Let us

define two functions encapsulating these expressions:

$$L_0(y) \;:=\; \frac{(1-p_s^\circ)\,(c_h(y)+\beta_h) + p_h^+\,(c_s(y+1)+\beta_s)}{(1-p_h^\circ)(1-p_s^\circ) \;-\; p_h^+\,p_s^-},$$

$$L_1(y) \;:=\; \frac{p_s^-\,(c_h(y-1)+\beta_h) + \left(1-p_h^\circ\right)(c_s(y)+\beta_s)}{(1-p_h^\circ)(1-p_s^\circ) \;-\; p_h^+\,p_s^-}.$$

Thus at the state k where we switch from $Hold$ at k to $serve$ at $k+1$, we must have:

$$V(k) = L_0(k), \quad V(k+1) = L_1(k+1).$$

Since L_0, L_1 were found as the solutions to the above systems, we have the following recurrence relations for the switchover state k:

$$L_0(k) = c_h(k) + \beta_h \;+\; p_h^\circ\,L_0(k) \;+\; p_h^+\,L_1(k+1) \tag{3.40}$$

$$L_1(k) = c_s(k) + \beta_s \;+\; p_s^-\,L_0(k-1) \;+\; p_s^\circ\,L_1(k) \tag{3.41}$$

Lemma 2. Let V be the optimal value function of our problem. Then:

(A) For every $n \in \{0,\ldots,Q-1\}$,

$$V(n) \;\le\; L_0(n), \text{ and } V(n+1) \;\le\; L_1(n+1).$$

The above inequalities become equalities if and only if $Hold$ is optimal at n, and $Serve$ is optimal at $n+1$.

(B) If $k \in \{0,\ldots,Q-1\}$ is a point where the policy is to switch from $Hold$ at k to $Serve$ at $k+1$, then

$$V(k) \;=\; L_0(k), \quad V(k+1) \;=\; L_1(k+1).$$

Proof.

We establish Part (B) by solving for k at which the switch from $Hold$ to $Serve$ happens. For an arbitrary state k:

$$V(k) \;\leq\; c_h(k) \;+\; \beta_h \;+\; p_h^{\circ}\, V(k) \;+\; p_h^{+}\, V(k+1) \tag{3.42}$$

$$V(k+1) \;\leq\; c_s(k+1) \;+\; \beta_s \;+\; p_s^{-}\, V(k) \;+\; p_s^{\circ}\, V(k+1) \tag{3.43}$$

Rearranging (3.42)–(3.43) and solving the resulting linear system gives the component-wise bounds

$$V(k) \;\leq\; L_0(k), \qquad V(k+1) \;\leq\; L_1(k+1). \tag{3.44}$$

Both inequalities in (3.44) are tight *iff* it is optimal to *Hold* in state k and to *Serve* in state $k+1$; otherwise, at least one inequality is strict. This establishes Part (A). □

Lemma 3. Let V be the optimal value function. Then:

(A) If $k \in \{0, \ldots, Q-2\}$ is a point of switching from $Hold$ at k to $Serve$ at $k+1$, then

$$L_1(k+1) \;\leq\; L_0\Big(k+1\Big).$$

Moreover, the inequality is strict unless $Hold$ and $Serve$ are *both* optimal at $k+1$.

(B) Let $k_0 \in \{0, \ldots, Q-1\}$ be such that

$$L_1(k_0) \;\geq\; L_0\Big(k_0\Big) \quad \text{and} \quad L_1(k_0+1) \;\leq\; L_0\Big(k_0+1\Big).$$

Assuming $L_0 \not\equiv L_1$ (that is, these two linear forms are not identical), it follows that k_0 must be a $Hold$ to $Serve$ switching point.

The graph in Figure 3.8 shows an example of the point (B) above, where $k_0 = 7$. We can see that

at the point where $L_1 - L_0$ switches its sign from negative to positive the value function changes from blue (Holding) to red (Serving).

Proof.

Part (A): From Lemma 2 (B) we know that if at state k, the $Hold$ decision is optimal and at state $k+1$, the $Serve$ decision is optimal then

$$V(k+1) \;=\; L_1(k+1) \tag{3.45}$$

On the other hand, by part (A) of the same Lemma and the fact that $k+1 < Q$, we also have

$$V(k+1) \;\leq\; L_0(k+1) \tag{3.46}$$

Combining these two facts gives

$$L_1(k+1) \;=\; V(k+1) \;\leq\; L_0(k+1) \tag{3.47}$$

Part (B). Suppose k_0 is the first index for which $L_1(k+1) \leq L_0(k+1)$. Since for any $k < k_0$ this inequality is not satisfied, by part (A) of the previous Lemma, there cannot be a switch from $Hold$ to $Serve$. Thus, k_0 itself remains a $Hold$ state. Now let $k_1 > k_0$ be the first index after k_0 at which $Serve$ becomes optimal again. By part (A), we must then have $L_1(k_1+1) \leq L_0(k_1+1)$; and by the linearity it follows that $L_1(k_1) < L_0(k_1)$. In this case, we get

$$V(k_1) \;\leq\; L_1(k_1) \;<\; L_0(k_1) \tag{3.48}$$

which contradicts the condition of Lemma 2 (B). Hence, k_0 must be the first switching point.

□

The above Lemma implies the following characterization of the state where the $Hold$ to $Serve$

switch occurs:

Corollary 2. Assume $L_0 \not\equiv L_1$, and let V be the optimal value function of our problem. Then the $Hold$ to $Serve$ can either happen at 0, $Q-1$ or at a $k_0 \in [1, Q-2]$ if and only if

$$L_1(k_0) \;\geq\; L_0\Big(k_0\Big) \quad \text{and} \quad L_1(k_0+1) \;\leq\; L_0\Big(k_0+1\Big).$$

Proof.

The 'if' part is just the Lemma 3-(B). The 'only if' part follows from Lemma 2 (A) and (B). Indeed, since $1 \leq k_0 \leq Q-1$

$$V(k_0) \;\leq\; L_1(k_0), \quad \text{and} \quad V(k_0+1) \;\leq\; L_0(k_0+1).$$

On the other hand, by part (B)

$$V(k_0) \;=\; L_0(k_0), \quad \text{and} \quad V(k_0+1) \;=\; L_1(k_0+1).$$

which yields the result. □

Sufficient Conditions for Threshold Control

We state conditions ensuring that there is exactly *one* hold-to-serve switch, i.e. the optimal policy follows a threshold structure:

Lemma 4. Assume $L_0 \not\equiv L_1$. Then each of the scenarios below guarantees that the optimal control is of threshold type (exactly one switch from $Hold$ to $Serve$):

a) If $L_1(1) < L_0(1)$ and $L_1(Q) \leq L_0(Q)$, then it is optimal to $Hold$ only at $y = 0$ and to $Serve$ at every $y > 0$.

b) If $L_0(1) < L_1(1)$ and $L_0(Q) \leq L_1(Q)$, then it is optimal to $Hold$ at all $y < Q$ and $Serve$

only at $y = Q$.

c) If $L_0(1) \leq L_1(1)$ and $L_0(Q) > L_1(Q)$, then there is a unique threshold $k_0 \in \{1, \ldots, Q-1\}$ such that $Hold$ is optimal for $y \leq k_0$ and $Serve$ is optimal for $y > k_0$, where k_0 is the smallest state satisfying $L_0(k_0+1) \geq L_1(k_0+1)$.

Proof.

Part (a). Since $L_0 \not\equiv L_1$, the given condition implies

$$L_1(k) < L_0(k) \quad \forall\, k \in \{1, \ldots, Q-1\}.$$

By Lemma 2 (A), we then have

$$V(k) \;\leq\; L_1(k) \;<\; L_0(k) \quad \forall\, k \in \{1, \ldots, Q-1\}.$$

Hence, Lemma 2 (B) ensures that in none of the states in $\{1, \ldots, Q-1\}$, we switch from $Hold$ to $Serve$. Therefore, the only possible switch is at $y = 0$.

Part (b). A similar argument applies here, except we have

$$L_0(k+1) < L_1(k+1) \quad \forall\, k \in \{0, \ldots, Q-2\}.$$

By Lemma 2 (A),

$$V(k+1) \;\leq\; L_0(k+1) \;<\; L_1(k+1) \quad \forall\, k \in \{0, \ldots, Q-2\}.$$

Therefore, again by Lemma 2 (B), $k+1$ cannot be the first $Serve$ state, so the only state with $Serve$ as optimal decision is $y = Q$.

Part (c). This case essentially follows from Lemma 3 (B). We need to only show that no additional switch can occur for any $k > k_0$. Since $Serve$ is optimal at $k_0 + 1$, a subsequent switch back

to $Hold$ would necessitate yet another $Hold$ to $Serve$ switch (because we are forced to serve at $y = Q$ eventually). However, by the same reasoning as in part (a), within the region $[k_0 + 1, Q]$ we have

$$L_1(k) \;\leq\; L_0(k).$$

Thus, the same argument excludes the possibility of another switch there. □

Switching from Serving to Holding

So far, we have identified conditions for exactly one switch from $Hold$ to $Serve$ as the optimal decision. However, if these conditions are not satisfied, it is possible that the optimal decision is to switch from $Hold$ to $Serve$ and then back to $Hold$. This scenario is more complex to characterize because we do not have convenient equalities for the states where a $Serve$ to $Hold$ switch can occur. In particular, we have not yet addressed two potential situations:

1. $L_0 \equiv L_1$, i.e. the two linear forms $L_0(\cdot)$ and $L_1(\cdot)$ coincide identically.

2. $L_1(1) \leq L_0(1) \quad \text{and} \quad L_1(Q) \geq L_0(Q)$.

These two cases require additional attention to determine whether further switches (beyond the initial $Hold$ to $Serve$) can arise.

Case 1: $L_0 \equiv L_1$**.**

In this case, we show that every admissible strategy is optimal: either action may be chosen in the intermediate states ($0 < y < Q$) without affecting optimality. Suppose $L_0 \equiv L_1$. Denote the common function as $L := L_0 = L_1$. Substituting $L(n)$ into the recurrences for L_0 and L_1 (see (3.40)–(3.41)), we obtain:

$$L(k) \;=\; c_h(k) + \beta_h \;+\; p_h^{\circ}\, L(k) \;+\; p_h^{+}\, L(k+1) \;=\; c_s(k) + \beta_s \;+\; p_s^{-}\, L(k-1) \;+\; p_s^{\circ}\, L(k) \quad (3.49)$$

Hence, L satisfies the value-iteration equation in such a way that both actions yield the same cost at every state. We show that L is effectively the only solution under this equality condition. Observe that if $k_0 \to k_0 + 1$ is the first $Hold$ to $Serve$ switch, then by Lemma 2**(B)**, we have:

$$V(k_0) = L(k_0), \quad V(k_0 + 1) = L(k_0).$$

For a given state k, consider the value iteration equation:

$$V(k) \;=\; \min\Big\{c_h(k)+\beta_h+p_h^{\circ}\,V(k)+p_h^{+}\,V(k+1),\quad c_s(k)+\beta_s+p_s^{-}\,V(k-1)+p_s^{\circ}\,V(k)\Big\} \tag{3.50}$$

Now, suppose L(k) satisfies the recurrence

$$L(k) \;=\; c_h(k) + \beta_h + p_h^{\circ}\,L(k) + p_h^{+}\,L(k+1) \tag{3.51}$$

Then it follows that

$$L(k) \;=\; \min\Big\{L(k),\; c_s(k) + \beta_s + p_s^{-}\,V(k-1) + p_s^{\circ}\,L(k)\Big\} \tag{3.52}$$

Since $L(k) = c_s(k) + \beta_s + p_s^{-}\,L(k-1) + p_s^{\circ}\,L(k)$, we find that $V(k-1) = L(k-1)$. A similar argument at state $k+1$ shows that $V(k+2) = L(k+2)$. By repeatedly applying both forward and backwards from k, we see $V(n) = L(n)$ for all $n \in \{0, 1, \ldots, Q\}$. In other words, once V and L coincide at two adjacent states, they coincide everywhere. Therefore, in the case $L_0 \equiv L_1$, the function L itself is the only valid solution to the value-iteration equation, and the actual value function V must coincide with it across the entire state space.

Case 2: Non-Trivial Scenario: $L_1(1) \leq L_0(1)$ **and** $L_1(Q) \geq L_0(Q)$.

Here we consider that at $y = 1$, L_1 is below (or equal to) L_0, whereas at $y = Q$, L_1 is above (or equal to) L_0. We determine whether the optimal control requires multiple switches (e.g., $Hold$ to $Serve$ to $Hold$ to $Serve$). Consider the following two strategies:

- **Strategy 1:** $Hold$ at $y = 0$, then switch to $Serve$ for all $y > 0$.
- **Strategy 2:** $Hold$ for all $y < Q$, then switch to $Serve$ at $y = Q$.

We denote the corresponding value functions by $V_1(n)$ and $V_2(n)$, respectively.

Strategy 1: $Hold$ **at** $y = 0$ **then** $Serve$ **for** $y > 0$**.**

By the earlier analysis in Section 3.4.2, we have

$$V_1(0) \;=\; L_0(0), \quad V_1(1) \;=\; L_1(1).$$

We $Serve$ whenever $n \geq 1$. Hence, for $n \geq 1$,

$$V_1(n) \;=\; \alpha_s + \delta_s\,(n-1) + \beta_s \;+\; p_s^-\,V_1(n-1) \;+\; p_s^\circ\,V_1(n) \tag{3.53}$$

Applying the explicit serving-action solution from Section 3.4.2 (with $y_{\text{First\_Hold}} = 0$), we obtain

$$V_1(n) \;=\; A_0\,n \;+\; B_0 \;+\; \Big[L_0(0) - B_0\Big]\,\beta^{\,n}, \quad \beta \;=\; \frac{p_s^-}{1-p_s^\circ}, \quad \text{for } p_s^- + p_s^\circ < 1.$$

If $p_s^- + p_s^\circ = 1$, then $\beta = 1$ and the solution for $V_1(n)$ instead becomes a quadratic in n:

$$V_1(n) \;=\; A_0'\,n^2 \;+\; B_0'\,n \;+\; C_0' \tag{3.54}$$

Strategy 2: $Hold$ **for** $y < Q$ **then** $Serve$ **at** $y = Q$**.**

Under Strategy 2, we do not $Serve$ until the state reaches $y = Q$. Specifically, we switch from $Hold$ to $Serve$ only at $y = Q - 1 \to Q$. Consequently,

$$V_2(Q-1) \;=\; L_0(Q-1), \quad V_2(Q) \;=\; L_1(Q).$$

For $n < Q$, the decision is to $Hold$, so

$$V_2(n) \;=\; \alpha_h \;+\; \delta_h\, n + \beta_h \;+\; p_h^{\circ}\, V_2(n) \;+\; p_h^{+}\, V_2(n+1) \tag{3.55}$$

Using the holding-action formula from Section 3.4.2 (with $y_{\text{Last\_Serve}} = Q$), we get:

$$A_0\, n \;+\; B_0 \;+\; \Big[L_0(0) - B_0\Big]\, \beta^{\,n} \le A_1\, n \;+\; B_1 \;+\; \Big[L_1(Q) - A_1\, Q - B_1\Big]\, \alpha^{\,n-Q},$$

$$V_2(n) \;=\; A_1\, n \;+\; B_1 \;+\; \Big[L_1(Q) - A_1\, Q - B_1\Big]\, \alpha^{\,n-Q}, \quad \alpha \;=\; \frac{1 - p_h^{\circ}}{p_h^{+}}, \quad \text{for } p_h^{\circ} + p_h^{+} < 1.$$

If $p_h^{\circ} + p_h^{+} = 1$, then $\alpha = 1$ and the solution for $V_2(n)$ again takes the following quadratic form:

$$V_2(n) \;=\; A_1'\, n^2 \;+\; B_1'\, n \;+\; C_1' \tag{3.56}$$

Let $V(n)$ be the true value function. Then by construction

$$V(n) \;\le\; V_1(n) \quad \text{and} \quad V(n) \;\le\; V_2(n) \quad \forall n.$$

Lemma 5. Assume $L_1(1) \;>\; L_0(0)$ and $L_1(Q) \;<\; L_0(Q)$. Let V be any solution to the value-iteration equations. Then, for all n,

$$V(n) = \min\left\{V_1(n),\, V_2(n)\right\}.$$

Moreover, exactly one of the following must occur:

(a) $V_1(n) \leq V_2(n)\ \forall n$. In this case, the optimal control effectively uses strategy 1, giving a threshold at $y = 1$.

(b) $V_2(n) \leq V_1(n)\ \forall n$. Then Strategy 2 dominates, giving a threshold at $y = Q$.

(c) There exists a unique n_0 such that

$$V_1(n) < V_2(n) \quad \text{for } n \leq n_0, \quad \text{and} \quad V_1(n) > V_2(n) \quad \text{for } n > n_0.$$

In this case, the optimal policy is: $Hold$ at $y = 0$, $Serve$ for $1 \leq y \leq n_0$, Switch back to $Hold$ for $n_0 + 1 \leq y \leq Q - 1$, and finally $Serve$ at $y = Q$.

Note that by Corollary 2, in this scenario, the only possible switching transitions from $Hold$ to $Serve$ are at: $0 \rightarrow 1$ and $Q-1 \rightarrow Q$, thus the only other possible strategy besides strategy 1 and strategy 2 is a mix of the two, where one switches to $Serve$ at $y = 1$ and continues up to some $n_0 \geq 1$, then switch back to $Hold$ at $n_0 + 1$ and continue holding until $Q - 1$ before serving again at $y = Q$. In this scenario, for $n \leq n_0$, we follow strategy 1, and for $n > n_0$, we follow strategy 2.

$$V(n) = \begin{cases} V_1(n), & n \leq n_0, \\ V_2(n), & n > n_0. \end{cases}$$

We know that $V(n) \leq \min\{V_1(n), V_2(n)\}$, hence the above piecewise solution is valid exactly when $V_1(n) \leq V_2(n)$ for $n \leq n_0$ and $V_1(n) \geq V_2(n)$ for $n > n_0$. This establishes case (c) above. Cases (a) and (b) are simpler scenarios where one strategy's cost remains below the other's for the entire range $n \in \{0, \ldots, Q\}$. We summarize the optimal policy in Theorem 2 below.

Theorem 2 (Optimal Policy Characterization for Discounted Cost Formulation)**.** Let L_0, L_1, V_1, and V_2 be the functions defined previously. The optimal policy is given by the following cases:

- $L_1(1) < L_0(1)$ and $L_1(Q) \leq L_0(Q)$: $Hold$ only at $y = 0$, and $Serve$ for every $y > 0$.
- $L_0(1) \leq L_1(1)$ and $L_0(Q) \leq L_1(Q)$: $Hold$ for all $y < Q$ and $Serve$ only at $y = Q$.
- $L_0(1) \leq L_1(1)$ and $L_0(Q) \geq L_1(Q)$: Define $k_0 := \max\{\, k \geq 0 :\ L_0(k) \ \leq\ L_1(k)\}$.

  Then $Hold$ is optimal for $y \leq k_0$, and $Serve$ is optimal for $y > k_0$.
- $L_1(1) \leq L_0(1)$ and $L_1(Q) \geq L_0(Q)$.:
  - $V_1(n) \ \leq \ V_2(n)$ for all n: $Hold$ only at $y = 0$ and $Serve$ at every $y > 0$.
  - $V_2(n) \ \leq \ V_1(n)$ for all n: $Hold$ for $y \leq Q-1$, and $Serve$ only at $y = Q$.
  - $V_1(n)$ crosses above $V_2(n)$ once: $Hold$ at $y = 0$, $Serve$ at $y = Q$ and

    $Serve$ for $1 \ \leq \ y \ \leq \ n_{\text{cross}}$, and $Hold$ for $n_{\text{cross}} + 1 \ \leq \ y \ \leq \ Q - 1$.

3.4.3 Discussion

Under the long-run average cost formulation, holding non-priority customers adds to the cumulative waiting cost, with no offsetting benefit in the long term. This is because, in this formulation, the future waiting costs are weighted equally to present costs, so postponing service does not yield any advantage. As a result, the optimal policy is simple: the controller should always serve a non-priority customer whenever it is possible to do so.

In the infinite horizon discounted cost formulation, the controller's objective is to minimize the present value of all future waiting costs, with a discount factor making immediate costs more significant than those in the distant future. This discounting alters the optimal policy structure compared to the long-run average cost case. Specifically, it can be optimal to strategically hold a non-priority customer. By holding, the controller can prevent the higher cost that would arise if a non-priority customer is served and subsequently blocks a priority customer. However, as the queue of waiting non-priority customers grows, the marginal cost of holding increases and

eventually outweighs the benefit of waiting, leading to a threshold policy: the controller serves a non-priority customer only when the queue length exceeds a certain threshold.

We find that the fourth ('reverse-threshold') case in the infinite-horizon, discounted-cost formulation arises only under a specific combination of parameters: (i) an arrival is highly likely to occur during the non-priority service—specifically, when parameters λ and T_{np} are large, so that $\lambda T_{\mathrm{np}} \gtrsim 1$, (ii) probability of a customer being priority type is relatively low (small p), and (iii) waiting cost for priority customer is relatively low (small W_p). Under these circumstances, it is possible for $L_1(1) < L_0(1)$ while $L_1(Q) > L_0(Q)$, resulting in the three-switch structure described in case 4 of Theorem 2. *Serve* is initially favorable but becomes less favorable as the number of waiting non-priority customers increases. As the priority share p and the priority waiting cost W_p are small, the immediate cost of blocking a priority arrival when a non-priority customer is served is low. Thus, at small y, we get $L_1(1) < L_0(1)$, and serving is optimal. However, because arrivals are likely within a non-priority service window, each additional NP service *increases* the serve-side marginal cost relative to hold (via $\delta_s > \delta_h$). This produces the observed $Hold \to Serve \to Hold \to Serve$ policy rather than the trivial one-threshold policy.

3.5 Visibility Case

In this section, we present a modelling framework that explicitly incorporates visibility. Prior to each customer's arrival, the controller receives information in the form (θ, τ_p) about the arrival of the next customer, where:

- θ - An estimate of the time to the next arrival (uncertain).
- τ_p - Customer type: Priority or Non-Priority (certain).

A key challenge in modelling this system is to represent and update the belief about the time to the next arrival. After a service period (either T_p for priority or T_{np} for non-priority) elapses without an arrival, θ transitions to a new value.

For this study we assume the initial arrival parameter θ is drawn from the distribution on $\{\theta_1, \ldots, \theta_n\}$ with mass function $q = (q_1, \ldots, q_n)$. Some q_i may be zero, meaning that certain θ_i are not generated in the initial state but can arise only as a result of transitions. To keep the model tractable and suitable for the scenario at hand, we restrict attention to a finite set of possible arrival times, i.e. $\theta_i \in \Theta$. The main reason for this finite parameterization is to ensure that the state space remains finite. We also assume Θ is ordered by expected arrival time.

We model the uncertainty in θ with a two-point distribution, wherein $\theta_{i+1} = \theta_i(1+\nu)$ with probability μ and $\theta_{i+1} = \theta_i(1-\nu)$ with probability $1-\mu$. This characterization helps us ensure that uncertainty reduces as the arrival moves closer. ν is the parameter controlling the spread of the distribution. Thus θ' is first sampled from a given discrete probability distribution as explained earlier, and the actual realization and future trajectory evolve following the two-point distribution. At each decision epoch, the arrival time θ (a realization from the previous iteration) gets updated according to the following rule:

$$\theta_{i+1} = \begin{cases} \left.\begin{array}{ll} \theta'(1+\nu) & \text{with probability } \mu, \\ \theta'(1-\nu) & \text{with probability } 1-\mu. \end{array}\right\} & \text{New Estimate} \\ \left.\begin{array}{ll} (\theta_i - T_{\text{service}})(1+\nu) & \text{with probability } \mu, \\ (\theta_i - T_{\text{service}})(1-\nu) & \text{with probability } 1-\mu. \end{array}\right\} & \text{Updation} \end{cases}$$

The system state is given by $(n_p,\ n_{np},\ \theta_i,\ \tau_p)$, where n_p and n_{np} are the number of priority and non-priority customers waiting at the queue and θ_i and τ_p are the estimates for time to next arrival and type of next-arrival. The action set available in state s is:

$$\mathcal{A}(s) = \begin{cases} \{a_1\}, & \text{if } n_p > 0, \\ \{a_2, a_3\}, & \text{if } 0 < n_{np} < Q, \\ \{a_2\}, & \text{if } n_{np} = Q, \\ \{a_3\}, & \text{otherwise (i.e., } n_1 = n_2 = 0). \end{cases}$$

Here, the three possible actions are: a_1 (Serve a priority customer), a_2 (Serve a non-priority customer), and a_3 (Hold). In our θ updation rule, $T_{service} = T_p$ if action a_1 is taken and $T_{service} = T_{np}$ if action a_2 is taken. Let t denote the time to the next customer arrival, conditional on θ_i. In order to formulate the long-run average cost problem, we define the following additional parameters:

- $m_i = \mathbb{E}[t \mid \theta_i]$ (mean time to the next arrival),
- $f_i^p = \mathbb{E}[(T_p - t)^+ \mid \theta_i]$,
- $f_i^{np} = \mathbb{E}[(T_{np} - t)^+ \mid \theta_i]$,
- $\alpha_i^p = \Pr(t < T_p \mid \theta_i)$,
- $\alpha_i^{np} \equiv \Pr(t < T_{np} \mid \theta_i)$,
- $q_j = \Pr(\theta = \theta_j)$ (Probability that a new customer's time to arrival is θ_j),
- $p'_\epsilon = \Pr(\text{next-customer type} = \epsilon')$, where $\epsilon' \in \{1 \text{ (priority)}, 0 \text{ (non-priority)}\}$,
- W_1, W_2: Waiting costs per unit time for priority and non-priority customers.

  (For notational simplicity, we use W_1, W_2 instead of W_p, W_{np}. This convention helps us write the expressions for the waiting cost of a new arrival in terms of $W_{2-\epsilon}$.)

Let $h(\cdot)$ and g denote the relative value function and the long-run average cost per unit time, respectively. Following Das et al. [1999], we express the relative value equations for the SMDP under the three actions below.

Action: Serve Priority When the controller decides to serve a priority customer, the epoch lasts for T_p, and the immediate cost is given by:

$$C_{\text{P}} = T_p \left((n_1 - 1)W_1 + n_2 W_2\right) + f_i^p W_{2-\epsilon} \tag{3.57}$$

The state transitions to:

- $(n_1 - 1 + \epsilon, n_2 + 1 - \epsilon, \theta_j, \epsilon')$, with probability $\alpha_i^p \cdot q_j \cdot p_{\epsilon'}$, and
- $(n_1 - 1, n_2, \theta_{i+1}, \epsilon)$, with probability $1 - \alpha_i^p$.

The relative value equation is given by:[8]

$$\begin{aligned} Q(s;a_1) + g.T_p =& T_p\left((n_1-1)W_1 + n_2W_2\right) + f_i^p W_{2-\epsilon} \\ &+ \alpha_i^p \overline{\overline{h}}(n_1 - 1 + \epsilon, n_2 + 1 - \epsilon) + (1-\alpha_i^p)h(n_1 - 1, n_2, \theta_{i+1}, \epsilon) \end{aligned} \tag{3.58}$$

Action: Serve Non-Priority When the controller decides to serve a non-priority customer, the epoch lasts for T_{np}, and the immediate cost is given by:

$$C_{\text{NP}} = T_{np}\left(n_1W_1 + (n_2-1)W_2\right) + f_i^{np}W_{2-\epsilon} \tag{3.59}$$

The state transitions to:

- $(n_1 + \epsilon, n_2 - \epsilon, \theta_j, \epsilon')$, with probability $\alpha_i^{np} \cdot q_j \cdot p_{\epsilon'}$, and
- $(n_1, n_2 - 1, \theta_{i+1}, \epsilon)$, with probability $1 - \alpha_i^{np}$.

The relative value equation is given by:

$$\begin{aligned} Q(s;a_2) + g.T_{np} =& T_{np}\left(W_1n_1 + W_2(n_2-1)\right) + f_i^{np}W_{2-\epsilon} \\ &+ \alpha_i^{np} \overline{\overline{h}}(n_1 + \epsilon, n_2 - \epsilon) + (1-\alpha_i^{np})h(n_1, n_2 - 1, \theta_{i+1}, \epsilon) \end{aligned} \tag{3.60}$$

Action: Hold When the controller decides to hold, the epoch lasts for an average of m_i, and the immediate cost is given by:

$$C_{\text{hold}} = m_i\left(n_1W_1 + n_2W_2\right) \tag{3.61}$$

A new customer arrives, and the state transitions to $(n_1 + \epsilon, n_2 + 1 - \epsilon, \theta_j, \epsilon')$.

The relative value equation is given by:

$$Q(s;a_3) + g.m_i = m_i\left(n_1W_1 + n_2W_2\right) + \overline{\overline{h}}(n_1 + \epsilon, n_2 + 1 - \epsilon) \tag{3.62}$$

[8]The averaged relative value function $h(.)$ is given by:

$$\overline{\overline{h}}(n_1, n_2) = \mathbb{E}_{\theta,\epsilon}[h(n_1, n_2, \theta, \epsilon)] = \sum_j q_j\left(p_0 h(n_1, n_2, \theta_j, 0) + p_1 h(n_1, n_2, \theta_j, 1)\right).$$

Value Iteration

Similar to the no-visibility case, the controller always serves the priority customer whenever one is waiting in the queue. Thus, for states with $n_1 > 0$, the relative value function satisfies:

$$\begin{aligned} h(n_1, n_2, \theta_i, \epsilon) + g.T_p =& T_p\left[(n_1 - 1)W_1 + n_2 W_2\right] + f_i^p W_{2-\epsilon} \\ &+ \alpha_i^p \overline{\overline{h}}(n_1 - 1 + \epsilon, n_2 + 1 - \epsilon) + (1 - \alpha_i^p) h(n_1 - 1, n_2, \theta_{i+1}, \epsilon) \end{aligned} \tag{3.63}$$

For states with $n_1 = 0$ and $n_2 > 0$, the relative value function satisfies:

$$\begin{aligned} h(0, n_2, \theta_i, \epsilon) = \min \Big\{ & T_{np}(n_2 - 1)W_2 + f_i^{np} W_{2-\epsilon} + \alpha_i^{np} \overline{\overline{h}}(\epsilon, n_2 - \epsilon) \\ & + (1 - \alpha_i^{np}) h(0, n_2 - 1, \theta_{i+1}, \epsilon) - g.T_{np}, \\ & m_i W_{np} n_2 + \overline{\overline{h}}(\epsilon, n_2 + 1 - \epsilon) - g.m_i \Big\} \end{aligned} \tag{3.64}$$

For the state with $n_1 = 0$ and $n_2 = 0$, i.e. $(0, 0, \theta_i, \epsilon)$, the relative value function satisfies:

$$h(0, 0, \theta_i, \epsilon) + g.m_i = \overline{\overline{h}}(\epsilon, 1 - \epsilon) \tag{3.65}$$

Due to the complexity of the update rule for the time to the next arrival, it is challenging to derive structural properties for this system analytically. However, simulation results indicate that when the next arrival is of priority type, the optimal policy is a switching curve in the (θ_i, n_{np}) space. This curve defines the boundary at which the controller transitions from holding to serving non-priority customers. In Section 3.6, we provide simulation results for this case.

3.6 Numerical Analysis

In this section, we present numerical results showing how key parameters affect the optimal policy under two cases: no visibility and visibility. In Table 3.4, we report long-run average costs for the

λ	T_{np}	p	W_p	g
0.050	6	0.4	100	0.30
0.075	6	0.4	100	0.89
0.100	6	0.4	100	1.81
0.125	6	0.4	100	3.01
0.100	**6**	0.4	100	1.81
0.100	**8**	0.4	100	3.16
0.100	**10**	0.4	100	4.85
0.100	**12**	0.4	100	6.72
0.100	6	**0.2**	100	1.07
0.100	6	**0.4**	100	1.81
0.100	6	**0.6**	100	2.24
0.100	6	**0.8**	100	2.36
0.100	6	0.4	**100**	1.81
0.100	6	0.4	**200**	3.58
0.100	6	0.4	**300**	5.35
0.100	6	0.4	**400**	7.12

Table 3.4: Long Run Average Cost: No Visibility Case ($Q = 4, T_p = 4$ and $W_{np} = 1$).

no-visibility case as we vary the arrival rate λ, non-priority service time T_{np}, priority share p, and priority waiting cost W_p, while keeping $Q = 4$, $W_{np} = 1$, and $T_p = 4$ fixed. We find that the long run average cost increases with λ, T_{np}, and p. In Figures 3.4, 3.5 and 3.6, we show how the optimal policy changes in the infinite-horizon discounted-cost formulation under no visibility, as a function of W_p, T_{np}, p, γ, and λ.

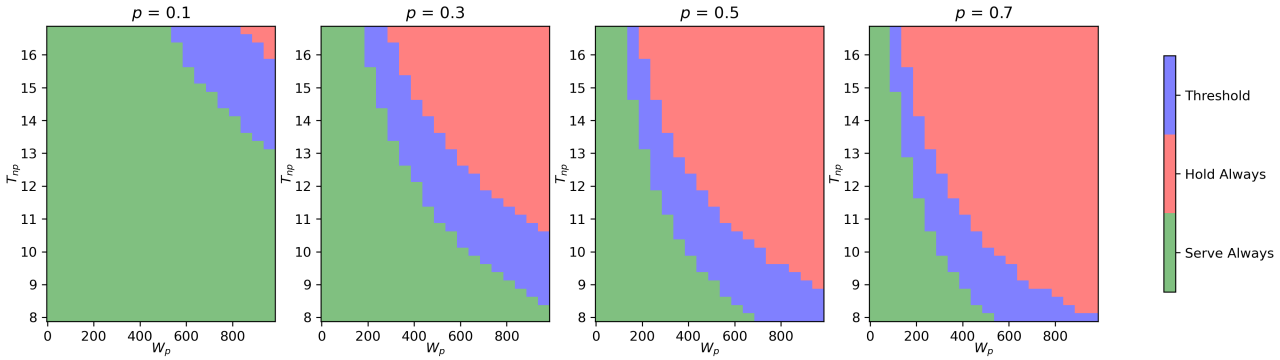

Figure 3.4: Optimal Policy ($Q = 3, T_p = 6, \lambda = 0.02, W_{np} = 1,$ and $\gamma = 0.005$).

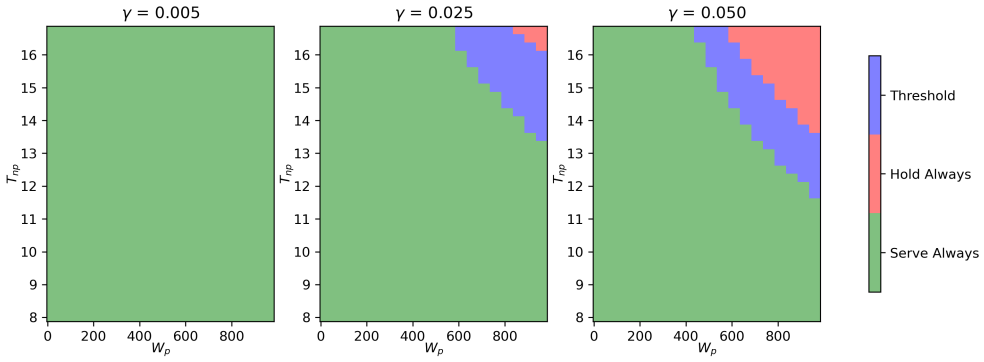

Figure 3.5: Optimal Policy ($Q = 3, T_p = 6, \lambda = 0.01, W_{np} = 1$, and $p = 0.3$).

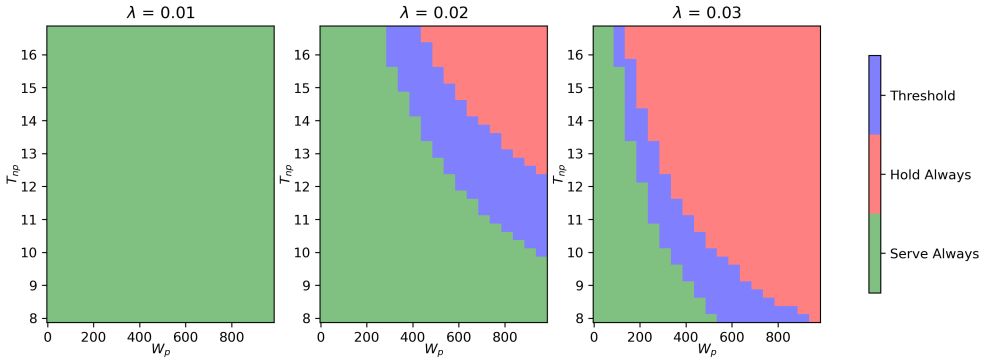

Figure 3.6: Optimal Policy ($Q = 3, T_p = 6, \gamma = 0.005, W_{np} = 1$, and $p = 0.2$).

Our simulation results show that, in the visibility case, when the next arrival is of priority type, the optimal policy is characterized by a two-dimensional switching curve: for each expected arrival time, the controller holds non-priority customers until the queue length reaches a specific threshold. Once this threshold is crossed, the controller serves non-priority customers whenever feasible (Figure 3.7). In Table 3.5, we show the benefit of visibility for the symmetric forecast-noise case ($\mu_1 = \mu_2 = 0.5$, $\nu = 0.5$). For the parameter values considered in this analysis, having exact information about the arrival time of the next priority customer reduces the long-run average cost by approximately 20%.

π	Q	T_{np}	p	W_p	Long Run Average Cost (g)	
					Certain Arrival Time	Uncertain Arrival Time
0.10	4	8	0.4	100	1.83	2.20
0.15	4	8	0.4	100	4.04	4.97
0.02	4	8	0.4	100	8.06	9.42
0.25	4	8	0.4	100	14.32	15.44
0.10	**3**	8	0.4	100	1.91	2.37
0.10	**4**	8	0.4	100	1.83	2.20
0.10	**5**	8	0.4	100	1.82	2.16
0.10	**6**	8	0.4	100	1.82	2.15
0.10	4	**6**	0.4	100	1.78	2.09
0.10	4	**8**	0.4	100	1.83	2.20
0.10	4	**10**	0.4	100	1.92	2.33
0.10	4	**12**	0.4	100	2.07	2.53
0.10	4	8	**0.2**	100	0.68	0.81
0.10	4	8	**0.4**	100	1.83	2.20
0.10	4	8	**0.6**	100	3.77	4.46
0.10	4	8	**0.8**	100	6.48	7.59
0.10	4	8	0.4	**100**	1.83	2.20
0.10	4	8	0.4	**200**	3.29	3.92
0.10	4	8	0.4	**300**	4.74	5.63
0.10	4	8	0.4	**400**	6.19	7.35

Table 3.5: Long Run Average Cost under Visibility: Uncertainty and Certainty in Arrival Time. $T_p = 4, \mu = 0.5,$ and $\nu = 0.5$. To generate the set Θ, we use geometric distribution with success probabilities $\{0.10, 0.15, 0.20, 0.25\}$, truncate the support to include 99% of the probability mass, and round each θ_i to one decimal place so that the resulting state space remains tractable.

3.7 Empirical Evidence

The Indian Railways network provides a setting that matches the setup we modelled in this study. A railway station is equivalent to a finite-buffer priority queue, where passenger trains are treated as priority customers and freight trains as non-priority customers. The travel time between stations is equivalent to service times, which depend on the type of train. For each station, the controller decides whether to hold a freight train or allow it to move, depending on whether a passenger train is expected to arrive soon.

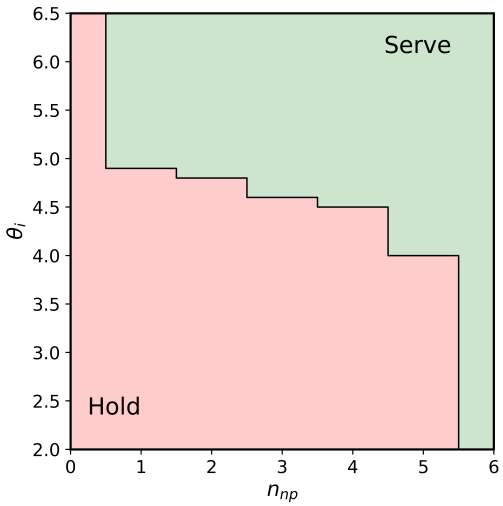

Figure 3.7: Visibility Case: Monotone Switching Curve.
($\lambda = 0.15, Q = 6, T_p = 2, T_{np} = 5, p = 0.8, \mu = 0.5, \nu = 0.5,$ and $W_p = 100$).

To empirically examine hold-versus-release[9] decisions by traffic controllers, we analyze train movement data from 'ADSR–MPO' section of the Deen Dayal Upadhyay division of Indian Railways. The dataset comprises minute-level arrival and departure timestamps for all trains at the 15 stations along this section, spanning the period from September 2021 to August 2022. For the purposes of this study, we restrict the sample to observations that meet the following criteria: (i) at the time of decision, the downstream block section is clear, and (ii) the next arrival is a passenger train.

Let Y_{ist} be a binary indicator equal to 1 if freight train i at station s and decision time t is held by the section controller, and 0 if released. We model the conditional probability of a holding a freight train using logistic regression:

$$\log\left(\frac{\Pr(Y_{ist}=1)}{1-\Pr(Y_{ist}=1)}\right) = \alpha + \sum_{k=2}^{7} \beta_k \, \mathbf{1}_{\{\, cnt\_g_{st}=k \,\}} + \gamma \, \ln(tna_{st}) + \delta' S_s + \theta' T_t + \varepsilon_{ist} \tag{3.66}$$

Here, cnt_g_{st} captures the freight-train count at station s and time t, entered as a seven-level categorical variable (1–7) with level 1 as the reference in the regression. Each β_k captures the

[9] In this context, 'release' is synonymous with 'serve', following the terminology used in our analytical model.

change in log-odds of a hold decision at freight congestion level k relative to level 1. The variable tna_{st} captures the remaining time to the next passenger-train arrival at station s and time t. The station vector S_s contains $Tracks\_Total_s$ (total tracks), $Distance_s$ (next block-section length, km), and $Junction_s$ (indicator for station s being a junction). The vector T_t comprises temporal covariates: indicators for the six-hour interval containing decision time t (four categories) and a weekend indicator.

0 (base -Release)	Coefficient (Standard Error)
Count of Freight Trains (Reference Level - 1)	
2	-0.3730^{***} (0.0561)
3	-1.0140^{***} (0.0597)
4	-1.2114^{***} (0.0771)
5	-1.2097^{***} (0.1212)
6	-1.3783^{***} (0.2958)
7	-1.9775^{***} (0.7227)
ln (Time to Next Passenger Arrival in Minutes)	-0.1040^{***} (0.0204)
Station Characteristics	
Tracks_Total	$0.3653^{***}(0.0115)$
Distance	$0.4023^{***}(0.0189)$
Junction	$-1.9854^{***}(0.0721)$
Control : Hour_Slot and Day_Week	Yes
Observations	3,30,636
Pseudo R^2	0.07

Table 3.6: Effect of queue length and time of arrival of next passenger train on the probability of holding a freight train ($*, **, ***$ indicates statistical significance at 10%,5%,1% level).

The results in Table 3.6 provide two key insights. First, as the number of freight trains waiting in the queue increases, the probability of a release decision increases. Second, a controller is more likely to hold a freight train when a passenger train is due to arrive soon, whereas a longer expected time until the next passenger arrival significantly increases the likelihood of releasing a freight train.

Piecewise Logit Model To assess whether freight congestion shows a threshold effect, specifically, freight train queue length beyond which the release probability increases significantly, we estimate a piecewise logit model [Francesco Ficetola and Denoël, 2009]. For each candidate thresh-

old k in the number of waiting freight trains, we define the following:

$$cnt\_left_{st} = \min(cnt\_g_{st},\, k),$$

$$cnt\_right_{st} = \max(0,\, cnt\_g_{st} - k).$$

The resulting piecewise logit specification is given by:

$$\log\left(\frac{\Pr(Y_{ist}=1)}{1-\Pr(Y_{ist}=1)}\right) = \alpha + \beta_L\, cnt\_left_{st} + \beta_R\, cnt\_right_{st} + \gamma' \ln(tna_{st}) + \delta'\, S_s + \theta'\, T_t + \varepsilon_{ist} \tag{3.67}$$

	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$
AIC	30875.24	30874.22	30847.05	30853.11	30872.65	30876.97
BIC	30982.32	30992.01	30964.85	30970.91	30990.45	30994.77
Log-likelihood	-15427.61	-15426.10	-15412.52	-15415.55	-15425.32	-15427.48

Table 3.7: Model comparison for piecewise logit specification.

We estimate piecewise logit models for different threshold values (k) of cnt_g_{st}. To select the threshold that gives the best fit, we compare models using the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). As shown in Table 3.7, the model with $k = 3$ achieves the lowest AIC and BIC, indicating the best overall fit. This implies that the likelihood of releasing a freight train increases sharply once the count of waiting freight trains reaches three. To test the statistical significance, we compare the piecewise model to the baseline model ($k = 1$) using the likelihood ratio test. The likelihood ratio (LR) statistic is calculated as $LR = -2\,[\ell_{\mathcal{M}_0} - \ell_{\mathcal{M}_1}]$, where $\ell_{\mathcal{M}_0}$ is the log-likelihood of the baseline model and $\ell_{\mathcal{M}_1}$ is the log-likelihood of the piecewise model. A higher LR statistic indicates that the piecewise model provides a significantly better fit. The piecewise logit model with $k = 3$ gives an LR statistic of 30.18, which exceeds the critical value at 1% significance level.[10]

[10]The LR statistic follows a χ^2 distribution with degrees of freedom equal to the difference in the number of parameters between models; the critical value at 1% significance (df = 1) is 6.63.

3.8 Conclusion

In this chapter, we study a finite buffer priority queue with two customer classes and deterministic, class-dependent service times. We first show that, without any visibility, the controller should always serve non-priority customers for the long-run average cost objective. Under a discounted cost criterion, we derive conditions under which a simple threshold policy is optimal. We then introduce the setting where information about the next arrival is available, and formulate the resulting control problem. Through simulation, we show that the optimal policy is a two-dimensional monotone switching curve, defined by the number of waiting non-priority customers and the expected time until the next priority arrival. Using Indian Railways data, we empirically analyze traffic controllers' decisions for freight movement and find that observed behavior aligns with the threshold results of our study. Numerical analysis further shows that key factors such as priority share, holding-cost asymmetry, and service times affect the threshold, and that visibility can yield considerable savings.

The findings of this study have significant managerial implications for Indian Railways and other organizations that manage priority queues of such nature. By identifying optimal service control policies for train release decisions, our research provides actionable insights for section controllers. Specifically, threshold-based policies can serve as practical guidelines for managing freight train movements, ensuring minimal delays while maintaining passenger train priorities. The results also highlight the importance of incorporating visibility into decision-making processes. Providing section controllers with tools to predict the arrival times can significantly enhance their ability to make better decisions, reduce congestion, and improve throughput.

3.9 Appendix

3.9.1 Sub-cycle Analysis for Discounted Cost Formulation

Table 3.8: Cost Coefficient- δ_h (Hold Cycle)

Subcycle	**Probability**	**Discounted Cost**
$\boldsymbol{y \to y}$		
One priority customer arrives, then no arrival during T_p	$p(1-\pi_p)$	$\frac{W_{np}}{\gamma}\left(1-\frac{\lambda e^{\gamma(-T_p)}}{\gamma+\lambda}\right)$
Two consecutive priority customers arrive, then no arrival during T_p	$(p\pi_p)p(1-\pi_p)$	$\frac{W_{np}}{\gamma}\left(1-\frac{\lambda e^{-2\gamma T_p}}{\gamma+\lambda}\right)$
...	...	...
$\boldsymbol{y \to y+1}$		
One non-priority customer arrives	$(1-p)$	$\frac{W_{np}}{\gamma}\left(1-\frac{\lambda}{\gamma+\lambda}\right)$
One priority customer arrives, and then a non-priority customer arrives in T_p	$p\pi_p(1-p)$	$\frac{W_{np}}{\gamma}\left(1-\frac{\lambda e^{\gamma(-T_p)}}{\gamma+\lambda}\right)$
...	...	...
$\delta_h = \frac{W_{np}\left(e^{\gamma T_p}(\gamma+\lambda p)-p(\lambda+\gamma\pi_p)\right)}{\gamma(\gamma+\lambda)\left(e^{\gamma T_p}-p\pi_p\right)}$		

Table 3.9: Cost Coefficient- α_h (Hold Cycle)

Subcycle	**Probability**	**Discounted Cost**
$\boldsymbol{y \to y+1}$		
One non-priority customer arrives	$1-p$	0
One priority customer arrives, and then a non-priority customer arrives in T_p	$p\pi_p(1-p)$	$\frac{\lambda W_{np} e^{-T_p(\gamma+\lambda)}\left(\gamma-e^{\lambda T_p}\left(\gamma+\lambda-\lambda e^{\gamma T_p}\right)\right)}{\gamma(\gamma+\lambda)^2}$
Two priority customers arrive, and then a non-priority customer arrives in T_p	$(p\pi_p)^2(1-p)$	$\frac{\lambda W_{np} e^{-T_p(2\gamma+\lambda)}\left(\gamma-e^{\lambda T_p}\left(\gamma+\lambda-\lambda e^{\gamma T_p}\right)\right)}{\gamma(\gamma+\lambda)^2}$
Three priority customers arrive, and then a non-priority customer arrives in T_p	$(p\pi_p)^3(1-p)$	$\frac{\lambda W_{np} e^{-T_p(3\gamma+\lambda)}\left(\gamma-e^{\lambda T_p}\left(\gamma+\lambda-\lambda e^{\gamma T_p}\right)\right)}{\gamma(\gamma+\lambda)^2}$
...	...	...
$\alpha_h = \frac{\lambda(p-1)p\pi_p W_{np} e^{\lambda(-T_p)}\left(\gamma+\gamma\left(-e^{\lambda T_p}\right)+\lambda\left(e^{\gamma T_p}-1\right)e^{\lambda T_p}\right)}{\gamma(\gamma+\lambda)^2\left(p\pi_p-e^{\gamma T_p}\right)}$		

Table 3.10: Cost Coefficient- δ_s (Release Cycle)

Subcycle	**Probability**	**Discounted Cost**
$\boldsymbol{y \to y-1}$		
No arrival during T_{np}	$(1-\pi_{np})$	$\frac{W_{np}}{\gamma}(1-e^{-\gamma T_{np}})$
One priority customer arrives, and then no arrival in T_p	$\pi_{np}p(1-\pi_p)$	$\frac{W_{np}}{\gamma}(1-e^{-\gamma(T_{np}+T_p)})$
Two priority customers arrive, and then no arrival in T_p	$\pi_{np}p(\pi_p p)(1-\pi_p)$	$\frac{W_{np}}{\gamma}(1-e^{-\gamma(T_{np}+2T_p)})$
Three priority customers arrive, and then no arrival in T_p	$\pi_{np}p(\pi_p p)^2(1-\pi_p)$	$\frac{W_{np}}{\gamma}(1-e^{-\gamma(T_{np}+3T_p)})$
...	...	...
$\boldsymbol{y \to y}$		
A non-priority arrival in T_{np}	$\pi_{np}(1-p)$	$\frac{W_{np}}{\gamma}(1-e^{-\gamma T_{np}})$
One priority customer arrives, followed by a non-priority arrival in T_p	$\pi_{np}p\pi_p(1-p)$	$\frac{W_{np}}{\gamma}(1-e^{-\gamma(T_{np}+T_p)})$
Two priority customers arrives, followed by a non-priority arrival in T_p	$\pi_{np}(p\pi_p)^2(1-p)$	$\frac{W_{np}}{\gamma}(1-e^{-\gamma(T_{np}+2T_p)})$
...	...	...
$\delta_s=\frac{W_{np}e^{\gamma(-\mathrm{Tnp})}\left(p(\pi_p-\pi_{np})+p\pi_p\left(-e^{\gamma\mathrm{Tnp}}\right)+(p\pi_{np}-1)e^{\gamma T_p}+e^{\gamma(T_p+\mathrm{Tnp})}\right)}{\gamma\left(e^{\gamma T_p}-p\pi_p\right)}$		

Table 3.11: Cost Coefficient- α_s (Release Cycle)

Subcycle	**Probability**	**Discounted Cost**
$\boldsymbol{y \to y}$		
A non-priority arrival in T_{np}	$\pi_{np}(1-p)$	$\frac{W_{np}e^{-\mathrm{Tnp}(\gamma+\lambda)}\left(\gamma-e^{\lambda\mathrm{Tnp}}\left(\gamma+\lambda-\lambda e^{\gamma\mathrm{Tnp}}\right)\right)}{\gamma(\gamma+\lambda)}$
One priority customer arrives, followed by a non-priority arrival in T_p	$\pi_{np}p\pi_p(1-p)$	$\frac{e^{-((\gamma+\lambda)(T_p+T_{np}))}\left(\lambda W_{np}\left(e^{\gamma T_p}-1\right)e^{\lambda T_p}-\gamma W_{np}\left(e^{\lambda T}-1\right)\right)}{\gamma(\gamma+\lambda)}$
Two priority customers arrive, followed by a non-priority arrival in T_p	$\pi_{np}(p\pi_p)^2(1-p)$	$\frac{W_{np}\lambda e^{-\lambda(T_{np}+2T_p)}}{\gamma}\left(\frac{e^{-\gamma(T_{np}+T_p)}(1-e^{-\gamma T_p})}{\gamma}-T_p e^{-\gamma(T_{np}+2T_p)}\right)$
...	...	...
$\alpha_s=\frac{(1-p)\pi_{np}W_{np}e^{-T_{np}(\gamma+\lambda)}\left(\gamma+\frac{(p-1)p\pi_p\left(\gamma\left(e^{\lambda T_p}-1\right)-\lambda\left(e^{\gamma T_p}-1\right)e^{\lambda T_p}\right)}{e^{T_p(\gamma+\lambda)}-p\pi_p}-e^{\lambda T_{np}}\left(\gamma+\lambda-\lambda e^{\gamma T_{np}}\right)\right)}{\gamma(\gamma+\lambda)}$		

Table 3.12: Cost Coefficient- β_h (Hold Cycle)

Subcycle	**Probability**	**Discounted Cost**
$\boldsymbol{y \to y}$		
Two consecutive priority arrivals, then no arrival in T_p	$p\pi_p p(1-\pi_p)$	$\frac{\lambda W_p e^{-T_p(\gamma+\lambda)}\left(\gamma - e^{\lambda T_p}\left(\gamma+\lambda-\lambda e^{\gamma T_p}\right)\right)}{\gamma(\gamma+\lambda)^2}$
Three consecutive priority arrivals , then no arrival in T_p	$(p\pi_p)^2 p(1-\pi_p)$	$\frac{\lambda W_p\left(e^{\gamma T_p}+1\right)e^{-T_p(2\gamma+\lambda)}\left(\gamma+\gamma\left(-e^{\lambda T_p}\right)+\lambda\left(e^{\gamma T_p}-1\right)e^{\lambda T_p}\right)}{\gamma(\gamma+\lambda)^2}$
...	...	...
$\boldsymbol{y \to y+1}$		
Two priority arrivals, and then a non-priority in T_p	$(p\pi_p)^2(1-p)$	$\frac{\lambda W_p e^{-T_p(\gamma+\lambda)}\left(\gamma - e^{\lambda T_p}\left(\gamma+\lambda-\lambda e^{\gamma T_p}\right)\right)}{\gamma(\gamma+\lambda)^2}$
Three priority arrivals, and then a non-priority arrival in T_p	$(p\pi_p)^3(1-p)$	$\frac{\lambda W_p\left(e^{\gamma T_p}+1\right)e^{-T_p(2\gamma+\lambda)}\left(\gamma+\gamma\left(-e^{\lambda T_p}\right)+\lambda\left(e^{\gamma T_p}-1\right)e^{\lambda T_p}\right)}{\gamma(\gamma+\lambda)^2}$
...	...	...
$\beta_h = \frac{\lambda p^2 \pi_p W_p e^{\lambda(-T_p)}\left(\gamma+\gamma\left(-e^{\lambda T_p}\right)+\lambda\left(e^{\gamma T_p}-1\right)e^{\lambda T_p}\right)}{\gamma(\gamma+\lambda)^2\left(e^{\gamma T_p}-pP_1\right)}$		

Table 3.13: Cost Parameter- β_s (Release Cycle)

Subcycle	**Probability**	**Discounted Cost**
$\boldsymbol{y \to y-1}$		
One priority arrives, and then no arrival in T_p	$\pi_{np} p(1-\pi_p)$	$\frac{W_p\left(\lambda+\gamma e^{\gamma(-T_{np})-T_p(\gamma+\lambda)}-(\gamma+\lambda)e^{-(\gamma(T_p+T_{np}))}+\gamma e^{-T_{np}(\gamma+\lambda)}+\gamma\left(-e^{\gamma(-T_{np})}\right)\right)}{\gamma(\gamma+\lambda)}$
Two priority arrive, and then no arrival in T_p	$\pi_{np} p\pi_p p(1-\pi_p)$	$\frac{W_p\left(\left(e^{\gamma T_p}+1\right)\left(\gamma+\gamma\left(-e^{\lambda T_p}\right)+\lambda\left(e^{\gamma T_p}-1\right)e^{\lambda T_p}\right)e^{\gamma(-T_{np})-T_p(2\gamma+\lambda)}+e^{-T_{np}(\gamma+\lambda)}\left(\gamma-e^{\lambda T_{np}}\left(\gamma+\lambda-\lambda e^{\gamma T_{np}}\right)\right)\right)}{\gamma(\gamma+\lambda)}$
...	...	...
$\boldsymbol{y \to y}$		
One priority arrives, then a non-priority in T_p	$\pi_{np} p\pi_p(1-p)$	$\frac{W_p\left(\lambda+\gamma e^{\gamma(-T_{np})-T_p(\gamma+\lambda)}-(\gamma+\lambda)e^{-(\gamma(T_p+T_{np}))}+\gamma e^{-T_{np}(\gamma+\lambda)}+\gamma\left(-e^{\gamma(-T_{np})}\right)\right)}{\gamma(\gamma+\lambda)}$
Two priority arrive, then a non-priority in T_p	$\pi_{np}(p\pi_p)^2(1-p)$	$\frac{W_p\left(\left(e^{\gamma T_p}+1\right)\left(\gamma+\gamma\left(-e^{\lambda T_p}\right)+\lambda\left(e^{\gamma T_p}-1\right)e^{\lambda T_p}\right)e^{\gamma(-T_{np})-T_p(2\gamma+\lambda)}+e^{-T_{np}(\gamma+\lambda)}\left(\gamma-e^{\lambda T_{np}}\left(\gamma+\lambda-\lambda e^{\gamma T_{np}}\right)\right)\right)}{\gamma(\gamma+\lambda)}$
...	...	...
$\beta_s = \frac{p\pi_{np} W_p e^{\lambda(-T_p)-T_{np}(\gamma+\lambda)}\left(\gamma p\pi_p e^{\lambda T_p}+\lambda p\pi_p e^{\lambda T_p+T_{np}(\gamma+\lambda)}-(\gamma+\lambda)(p\pi_p-1)e^{\lambda(T_p+T_{np})}-\gamma e^{T_p(\gamma+\lambda)}+\gamma e^{T_p(\gamma+\lambda)+\lambda T_{np}}-\lambda e^{(\gamma+\lambda)(T_p+T_{np})}+\gamma\left(-e^{\lambda T_{np}}\right)\right)}{\gamma(\gamma+\lambda)\left(p\pi_p-e^{\gamma T_p}\right)}$		

Discounted Cost Formulation: Threshold Case

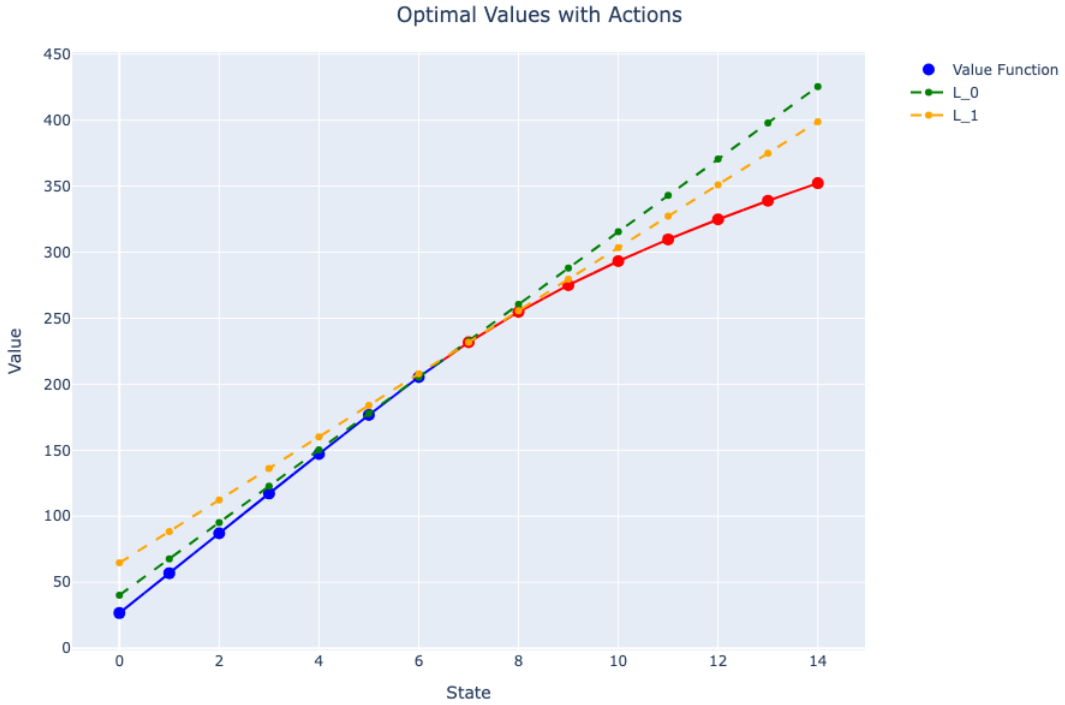

Figure 3.8: Illustration of $L_1 - L_0$ changing sign at $k_0 = 7$ (Threshold Policy).

Bibliography

Mazhar Arıkan, Vinayak Deshpande, and Milind Sohoni. Building reliable air-travel infrastructure using empirical data and stochastic models of airline networks. *Operations Research*, 61(1): 45–64, 2013.

Alessandro Arlotto, Andrew E Frazelle, and Yehua Wei. Strategic open routing in service networks. *Management Science*, 65(2):735–750, 2019.

Barry C Arnold. Multivariate logistic distributions. *Handbook of the logistic distribution*, 237: 244–45, 1992.

Kashish Arora, Fanyin Zheng, and Karan Girotra. Private vs. pooled transportation: Customer preference and design of green transport policy. *Manufacturing & Service Operations Management*, 2023. doi: 10.1287/msom.2022.0569.

Baris Ata and Xiaoshan Peng. An optimal callback policy for general arrival processes: a pathwise analysis. *Operations Research*, 68(2):327–347, 2020.

Baris Ata, Nasser Barjesteh, and Sunil Kumar. Spatial pricing: An empirical analysis of taxi rides in new york city. Technical report, Working Paper, 2019.

KR Balachandran. Control policies for a single server system. *Management Science*, 19(9):1013–1018, 1973.

Andreas Bärmann, Alexander Martin, and Hanno Schülldorf. A decomposition method for multi-period railway network expansion—with a case study for germany. *Transportation Science*, 51 (4):1102–1121, 2017.

Cynthia Barnhart, Hong Jin, and Pamela H Vance. Railroad blocking: A network design application. *Operations research*, 48(4):603–614, 2000.

Cynthia Barnhart, Douglas Fearing, and Vikrant Vaze. Modeling passenger travel and delays in the national air transportation system. *Operations Research*, 62(3):580–601, 2014.

Opher Baron, Oded Berman, Dmitry Krass, and Jianfu Wang. Using strategic idleness to improve customer service experience in service networks. *Operations Research*, 62(1):123–140, 2014.

Colin E Bell. Optimal operation of an m/g/1 priority queue with removable server. *Operations Research*, 21(6):1281–1290, 1973.

Moshe E Ben-Akiva and Steven R Lerman. *Discrete choice analysis: theory and application to travel demand*, volume 9. MIT press, 1985.

Chandra R Bhat and Ipek N Sener. A copula-based closed-form binary logit choice model for accommodating spatial correlation across observational units. *Journal of Geographical Systems*, 11:243–272, 2009.

Ralf Borndörfer, Torsten Klug, Thomas Schlechte, Armin Fügenschuh, Thilo Schang, and Hanno Schülldorf. The freight train routing problem for congested railway networks with mixed traffic. *Transportation Science*, 50(2):408–423, 2016.

Yann Bramoullé, Habiba Djebbari, and Bernard Fortin. Identification of peer effects through social networks. *Journal of econometrics*, 150(1):41–55, 2009.

Alberto Caprara, Matteo Fischetti, and Paolo Toth. Modeling and solving the train timetabling problem. *Operations research*, 50(5):851–861, 2002.

Alberto Caprara, Enrico Malaguti, and Paolo Toth. A freight service design problem for a railway corridor. *Transportation science*, 45(2):147–162, 2011.

Jinsheng Chen, Jing Dong, and Pengyi Shi. Optimal routing under demand surges: The value of future arrival rates. *Operations Research*, 73(1):510–542, 2025.

Veronica Dal Sasso, Leonardo Lamorgese, Carlo Mannino, Antonio Tancredi, and Paolo Ventura. Easy cases of deadlock detection in train scheduling. *Operations Research*, 70(4):2101–2118, 2022.

Tapas K Das, Abhijit Gosavi, Sridhar Mahadevan, and Nicholas Marchalleck. Solving semi-markov decision problems using average reward reinforcement learning. *Management Science*, 45(4):560–574, 1999.

Danial Davarnia, Jean-Philippe P Richard, Ece Icyuz-Ay, and Bijan Taslimi. Network models with unsplittable node flows with application to unit train scheduling. *Operations research*, 67(4): 1053–1068, 2019.

Gerard De Jong and Moshe Ben-Akiva. A micro-simulation model of shipment size and transport chain choice. *Transportation Research Part B: Methodological*, 41(9):950–965, 2007.

Vinayak Deshpande and Mazhar Arıkan. The impact of airline flight schedules on flight delays. *Manufacturing & Service Operations Management*, 14(3):423–440, 2012.

Kimon Drakopoulos and Fanyin Zheng. Network effects in contagion processes: Identification and control. *Columbia Business School Research Paper*, (18-8), 2017.

Johan Narvestad Fatnes. Flow-times in an m/g/1 queue under a combined preemptive/non-preemptive priority discipline.: Scheduled waiting time on single track railway lines. Master's thesis, Institutt for matematiske fag, 2010.

Pnina Feldman, Jun Li, and Hsin-Tien Tsai. Welfare implications of congestion pricing: Evidence from sf park. *Manufacturing & Service Operations Management*, 24(2):1091–1109, 2022.

Gentile Francesco Ficetola and Mathieu Denoël. Ecological thresholds: an assessment of methods to identify abrupt changes in species–habitat relationships. *Ecography*, 32(6):1075–1084, 2009.

David T. Frazier, Eric Renault, Lina Zhang, and Xueyan Zhao. Weak Identification in Discrete Choice Models. Technical report, 2021.

Guillermo Gallego and Robert Phillips. Revenue management of flexible products. *Manufacturing & Service Operations Management*, 6(4):321–337, 2004. doi: 10.1287/msom.1040.0054.

Pin Gao, Yuhang Ma, Ningyuan Chen, Guillermo Gallego, Anran Li, Paat Rusmevichientong, and Huseyin Topaloglu. Assortment optimization and pricing under the multinomial logit model with impatient customers: Sequential recommendation and selection. *Operations Research*, 69 (5):1509–1532, 2021. doi: 10.1287/opre.2021.2127.

Ragavendran Gopalakrishnan, Sherwin Doroudi, Amy R Ward, and Adam Wierman. Routing and staffing when servers are strategic. *Operations Research*, 64(4):1033–1050, 2016.

Christian Gourieroux, Alain Monfort, Eric Renault, and Alain Trognon. Generalised residuals. *Journal of econometrics*, 34(1-2):5–32, 1987.

Linda V Green, Sergei Savin, and Ben Wang. Managing patient service in a diagnostic medical facility. *Operations Research*, 54(1):11–25, 2006.

Jose A. Guajardo, Morris A. Cohen, Sang-Hyun Kim, and Serguei Netessine. Impact of Performance-Based Contracting on Product Reliability: An Empirical Analysis. *Management Science*, 58(5):961–979, May 2012.

Itai Gurvich, Martin A Lariviere, and Can Ozkan. Coverage, coarseness, and classification: Determinants of social efficiency in priority queues. *Management Science*, 65(3):1061–1075, 2019.

J Michael Harrison. Dynamic scheduling of a multiclass queue: Discount optimality. *Operations Research*, 23(2):270–282, 1975a.

J Michael Harrison. A priority queue with discounted linear costs. *Operations Research*, 23(2): 260–269, 1975b.

Brett A. Hathaway, Seyed M. Emadi, and Vinayak Deshpande. Don't call us, we'll call you: An empirical study of caller behavior under a callback option. *Management Science*, 67(3):1508–1526, 2021.

Brett A Hathaway, Evgeny Kagan, and Maqbool Dada. The gatekeeper's dilemma:"when should i transfer this customer?". *Operations Research*, 71(3):843–859, 2023.

Pu He, Fanyin Zheng, Elena Belavina, and Karan Girotra. Customer Preference and Station Network in the London Bike Share System. Working paper, Columbia Business School, 2019.

Daniel P Heyman. The t-policy for the m/g/1 queue. *Management Science*, 23(7):775–778, 1977.

Johan Högdahl and Markus Bohlin. A combined simulation-optimization approach for robust timetabling on main railway lines. *Transportation Science*, 57(1):52–81, 2023.

Yue Hu, Carri W Chan, and Jing Dong. Optimal scheduling of proactive service with customer deterioration and improvement. *Management science*, 68(4):2533–2578, 2022.

Ashish Kabra, Elena Belavina, and Karan Girotra. Bike-share systems: Accessibility and availability. *Management Science*, 66(9):3803–3824, 2020.

P Mahinda Karunaratne and Robert C Elston. A multivariate logistic model (mlm) for analyzing binary family data. *American journal of medical genetics*, 76(5):428–437, 1998.

Jeunghyun Kim and Amy R Ward. Dynamic scheduling of a gi/gi/1+ gi queue with multiple customer classes. *Queueing Systems*, 75:339–384, 2013.

Song-Hee Kim, Fanyin Zheng, and Joan Brown. Identifying the bottleneck unit: Impact of congestion spillover in hospital inpatient unit network. *Management Science*, 2023.

Ger Koole. Structural results for the control of queueing systems using event-based dynamic programming. *Queueing systems*, 30:323–339, 1998.

Leonardo Lamorgese and Carlo Mannino. A noncompact formulation for job-shop scheduling problems in traffic management. *Operations Research*, 67(6):1586–1609, 2019.

Ho Woo Lee, Sahng Hoon Cheon, and Won Joo Seo. Queue length and waiting time of the m/g/1 queue under the d-policy and multiple vacations. *Queueing Systems*, 54(4):261–280, 2006.

Yanzhe Lei and Stefanus Jasin. Real-time dynamic pricing for revenue management with reusable resources, advance reservation, and deterministic service time requirements. *Operations Research*, 68(3):676–685, 2020.

E Lerzan Örmeci, Apostolos Burnetas, and Jan van der Wal. Admission policies for a two class loss system. 2001.

Jun Li, Nelson Granados, and Serguei Netessine. Are consumers strategic? structural estimation from the air-travel industry. *Management Science*, 60(9):2114–2137, 2014.

Na Li, David A Stanford, Peter Taylor, and Ilze Ziedins. Nonlinear accumulating priority queues with equivalent linear proxies. *Operations Research*, 65(6):1712–1721, 2017.

Wenhao Li, Zhankun Sun, and L Jeff Hong. Who is next: Patient prioritization under emergency department blocking. *Operations Research*, 71(3):821–842, 2023.

Yunan Liu and Ward Whitt. Nearly periodic behavior in the overloaded g/d/s+ gi queue. *Stochastic Systems*, 1(2):340–410, 2011.

Alexander H Lovett, C Tyler Dick, and Christopher PL Barkan. Determining freight train delay costs on railroad lines in north america. *Proceedings of Rail Tokyo*, 2015.

Daniel McFadden. The measurement of urban travel demand. *Journal of public economics*, 3(4): 303–328, 1974.

Ken Moon, Kostas Bimpikis, and Haim Mendelson. Randomized markdowns and online monitoring. *Management Science*, 64(3):1271–1290, 2018.

Andrés Musalem, Marcelo Olivares, Eric T Bradlow, Christian Terwiesch, and Daniel Corsten. Structural estimation of the effect of out-of-stocks. *Management Science*, 56(7):1180–1197, 2010.

Marcelo Olivares, Christian Terwiesch, and Lydia Cassorla. Structural estimation of the newsvendor model: an application to reserving operating room time. *Management Science*, 54(1):41–55, 2008.

António Pacheco, Maria Lurdes Simões, and Paula Milheiro-Oliveira. Queues with server vacations as a model for pretimed signalized urban traffic. *Transportation Science*, 51(3):841–851, 2017.

Jonathan Patrick, Martin L Puterman, and Maurice Queyranne. Dynamic multipriority patient scheduling for a diagnostic resource. *Operations research*, 56(6):1507–1525, 2008.

Amil Petrin and Kenneth Train. A control function approach to endogeneity in consumer choice models. *Journal of marketing research*, 47(1):3–13, 2010.

Catherine Ragasa and John Mazunda. The impact of agricultural extension services in the context of a heavily subsidized input system: The case of malawi. *World development*, 105:25–47, 2018.

Nicola Rosaia. Competing platforms and transport equilibrium: Evidence from new york city. Technical report, Working Paper, 2022.

Paat Rusmevichientong and Huseyin Topaloglu. Robust assortment optimization in revenue management under the multinomial logit choice model. *Operations Research*, 60(4):865–882, 2012. ISSN 0030364X, 15265463.

Joel Spencer, Madhu Sudan, and Kuang Xu. Queuing with future information. *The Annals of Applied Probability*, pages 2091–2142, 2014.

Kalyan Talluri and Garrett van Ryzin. Revenue management under a general discrete choice model of consumer behavior. *Management Science*, 50(1):15–33, 2004. ISSN 00251909, 15265501.

Henk Tijms. A control policy for a priority queue with removable server. *Operations Research*, 22 (4):833–837, 1974.

Francis Vella. A simple estimator for simultaneous models with censored endogenous regressors. *International Economic Review*, pages 441–457, 1993.

Jianfu Wang, Opher Baron, and Alan Scheller-Wolf. M/m/c queue with two priority classes. *Operations Research*, 63(3):733–749, 2015.

Zhaodong Wang, Yanfeng Ouyang, and Ruifeng She. On solving a class of continuous traffic equilibrium problems and planning facility location under congestion. *Operations Research*, 70 (3):1465–1484, 2022.

R Wolvers. State based patient scheduling with a scare healthcare server and beds in a finite priority queue. 2020.

Jeffrey M Wooldridge. Quasi-maximum likelihood estimation and testing for nonlinear models with endogenous explanatory variables. *Journal of Econometrics*, 182(1):226–234, 2014.

Kuang Xu. Necessity of future information in admission control. *Operations Research*, 63(5): 1213–1226, 2015.

Kuang Xu and Carri W Chan. Using future information to reduce waiting times in the emergency department via diversion. *Manufacturing & Service Operations Management*, 18(3):314–331, 2016.

Dongyuan Zhan and Amy R Ward. Staffing, routing, and payment to trade off speed and quality in large service systems. *Operations Research*, 67(6):1738–1751, 2019.

Yueyang Zhong, Ragavendran Gopalakrishnan, and Amy R Ward. Behavior-aware queueing: The finite-buffer setting with many strategic servers. *Operations Research*, 2023.

Endong Zhu, Teodor Gabriel Crainic, and Michel Gendreau. Scheduled service network design for freight rail transportation. *Operations research*, 62(2):383–400, 2014.