

THE UNIVERSITY OF CHICAGO

PRICING AND MATCHING IN SERVICE AND MANUFACTURING SYSTEMS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
NASSER BARJESTEH

CHICAGO, ILLINOIS
AUGUST 2020

Copyright © 2020 by Nasser Barjesteh

All Rights Reserved

To my family

TABLE OF CONTENTS

LIST OF FIGURES	vii
LIST OF TABLES	x
ACKNOWLEDGMENTS	xii
ABSTRACT	xiii
1 AN EMPIRICAL ANALYSIS OF TAXI RIDES IN NEW YORK CITY	
1.1 Introduction	1
1.2 Literature Review	5
1.3 Model	9
1.4 Data	18
1.5 Estimation	21
1.5.1 Offline Estimation of the Primitives τ_{ij} , d_{ij} , S_{ij} , c , N_i , M , and m_i^1	21
1.5.2 Estimation of the Standard Deviation of Cost Shocks	23
1.5.3 Estimation of Demand Curve Parameters	25
1.6 Counterfactual Analysis	27
1.6.1 Spatial Pricing	28
1.6.2 Removing Local Search Friction	35
1.7 Concluding Remarks	38
2 DYNAMIC DISPATCH AND CENTRALIZED RELOCATION OF CARS IN RIDE-HAILING PLATFORMS	
2.1 Introduction	40
2.2 Literature Review	46
2.3 Model	53
2.4 Approximating Brownian Control Problem	59
2.5 Reduced Brownian Control Problem	66
2.6 Equivalent Workload Formulation	72
2.7 Solving the Equivalent Workload Formulation	77
2.8 Proposed Policy	87
2.9 Simulation Study and Discussion	89
2.10 Concluding Remarks	96
3 DYNAMIC PRICING OF A MULTICLASS MAKE-TO-STOCK QUEUE	
3.1 Introduction	97
3.2 Literature Review	99
3.3 Model	107
3.4 Approximating Brownian Control Problem	112
3.5 Equivalent Workload Formulation	115
3.6 Solution to the Equivalent Workload Formulation	119
3.6.1 Barrier Policies	119

3.6.2	Bellman Equation	120
3.6.3	Solution to the Bellman Equation	122
3.6.4	Candidate Policy and Its Optimality for the Equivalent Workload Formulation	127
3.7	Proposed Policy and Its Performance	128
3.8	Concluding Remarks	132
A	APPENDIX OF CHAPTER 1	134
A.1	Node Definitions	134
A.2	Further Discussion on the Matching Model	141
A.3	Supplemental Material for Section 1.4	144
A.4	Supplemental Material for Section 1.5	144
A.5	Monte Carlo Experiments	151
A.6	Cross-Validation	152
A.7	Supplemental Material for Section 1.6	154
A.7.1	Mathematical Formulations of the Mechanisms	154
A.7.2	Further Discussion on Origin-Destination Pricing	157
A.7.3	Further Discussion on the Impact of Maximum price multiplier on optimal spatial prices	157
A.8	Maximizing Drivers' Profit	158
A.9	Miscellaneous Proofs and Derivations	161
B	APPENDIX OF CHAPTER 2	171
B.1	Supplemental Material	171
B.1.1	Derivation of the Routing Process	171
B.1.2	Additional Model Primitives for the Motivating Example	171
B.1.3	No Opportunity for Arbitrage	176
B.1.4	A Simplification of the Reduced Brownian Formulation	177
B.1.5	Equivalence of the BCP and the RBCP	177
B.1.6	The Value Function of a Region of Inaction Type Policy.	178
B.1.7	Numerical Solution of the Partial Differential Equation	179
B.1.8	Solving the Equivalent Workload Formulation using the Markov Chain Approximation Method	190
B.1.9	Policy Improvement	202
B.1.10	Detailed Description of the Alternative Policies	204
B.1.11	Auxiliary Simulation Results	206
B.2	Miscellaneous Proofs	212
B.3	Auxiliary Results	222
C	APPENDIX OF CHAPTER 3	226
C.1	Riccati Equation	226
C.2	Derivations	232
C.2.1	Formal Derivation of the Approximating Brownian System	232
C.2.2	Derivation of Equation (3.23)	236
C.3	Miscellaneous Proofs	237

C.4 Auxiliary Results	246
REFERENCES	250

LIST OF FIGURES

1.1	Timing of events and definition of variables in each period.	11
1.2	Average number of pick-ups per minute in NYC. The dashed vertical lines depict the begining and the end of the day shift (6 _{AM} and 4 _{PM} , respectively).	19
1.3	Average number of pick-ups in NYC during the day shift on a weekday of each month. The dashed line highlights the monthly trend.	19
1.4	Average number of pick-ups and drop-offs during the day shift on weekdays in each node.	20
1.5	Estimated potential and satisfied demand (per period) during the day shift. The solid lines correspond to potential demand and the dashed lines correspond to satisfied demand.	27
1.6	Average number of customers and empty taxis (in a five-minute period) during the day shift on weekdays in September 2012 (these are obtained through estimation; see Problem (P1) in Section 1.5).	28
1.7	Optimal origin-only price multipliers for different maximum price variations $\bar{\eta}$ (maximum price multiplier $\hat{\eta} = 5$).	31
1.8	Percentage change in the average number of empty taxis and the fraction of served customers under the optimal origin-only prices compared to the base prices (maximum price variation $\bar{\eta} = 0.5$ and maximum price multiplier $\hat{\eta} = 5$).	33
1.9	Well-served nodes with high (star) and low (diamond) potential revenue. Prices decrease in 88% of the nodes with high potential revenue and increase in 92% of the nodes with low potential revenue.	33
1.10	Average price multipliers for rides originating at each node in the city under origin-destination spatial pricing (the average is calculated with respect to the demand from the node to each destination).	33
1.11	Optimal price multipliers under origin-destination pricing for the rides originating at the node denoted by a star (maximum price variation $\bar{\eta} = 0.5$ and maximum price multiplier $\hat{\eta} = 5$).	35
1.12	Change in consumer surplus from spatial pricing and removing the search friction (maximum price multiplier $\hat{\eta} = 5$).	36
1.13	Optimal price multipliers under citywide origin-only pricing and friction removal (maximum price variation $\bar{\eta} = 0.5$ and maximum price multiplier $\hat{\eta} = 5$).	38
2.1	A map of NYC and the nine areas considered. The arrows represent the non-local dispatch activities. The local dispatch activities are not depicted for visual clarity. An arrow from area i to area k represents the activity of dispatching a car in area i to pick up the next customer in area k	42
2.2	The closed stochastic processing network model of the NYC application and the numbering of the dispatch activities.	43

2.3	A graphical representation of the optimal solution to the static planning problem of the NYC application. The solid arrows from the buffers to the servers depict the basic dispatch activities and the rate at which they are conducted. The solid arrows between the buffers depict the basic relocation activities and the rate at which they are conducted. The dashed arrows depict the nonbasic dispatch activities. The non-basic relocation activities are not shown for visual clarity.	62
2.4	A graphical representation of the optimal solution to the static planning problem and the resulting pools. All local dispatch activities are basic. The basic non-local dispatch activities are depicted with solid black arrows. The basic relocation activities are depicted with dashed gray arrows. Areas with the same shading pattern belong to the same pool. There are three pools in total.	68
2.5	An illustration of the projection mapping L and its inverse L^{-1} for $p = 2$. The dashed gray and solid black lines represents \mathcal{W} and $\hat{\mathcal{W}}$, respectively; L maps $w \in \mathcal{W}$ to $\hat{w} \in \hat{\mathcal{W}}$	73
2.6	An illustration of the relationship between W and \hat{W} for $p = 3$. The gray and black curves depict a sample path of W and \hat{W} , respectively.	73
2.7	The non-dominated pushing controls and the corresponding pushing directions for the NYC application.	77
2.8	An illustration of the region of inaction and the pushing controls used on its boundary in each step of the algorithm for the NYC application.	81
2.9	An example of the discretized region of inaction with a discretization step size of 0.05.	82
2.10	An illustration of the three instances in which the control policy can be improved.	84
2.11	The expected discounted cost under various policies.	93
2.12	Percentage of the demand dropped under various policies.	94
3.1	$v_\gamma^-(0)$ and $v_\gamma^+(0)$ as a function of the average cost γ	125
3.2	Airy functions of the first and second kind.	125
A.1	Boundaries of the zones used in node generation.	141
A.2	Minimum achievable fraction of arcs between adjacent nodes with travel time longer than one period for given \hat{D}	141
A.3	The boundaries of the seventy five nodes used in the analysis and the area under consideration.	142
A.4	An illustration of matching in our imperfect matching model and the model of Buchholz [56].	143
A.5	Percentage change in the number of matches in our imperfect matching model compared to the imperfect matching model of Buchholz [56].	143
A.6	Histograms of the average trip duration and distances across nodes.	144
A.7	Average trip duration and distance for rides originating at each node.	145
A.8	Log likelihood values in the vicinity of the maximum likelihood estimate $\sigma = 1.4048$	147
A.9	Original graph.	148
A.10	Augmented graph.	148
A.11	Optimal origin-only price multipliers for different maximum price multipliers $\hat{\eta}$ (maximum price variation $\bar{\eta} = 0.2$).	158

A.12 Comparison of the optimal origin-only price multipliers for maximizing consumer surplus and drivers' profit (maximum price variation $\bar{\eta} = 0.5$ and maximum price multiplier $\hat{\eta} = 5$).	160
B.1 An example of a hat function in two dimensions. The shaded area depicts the support of the hat function.	188
B.2 An example of a mesh of $\hat{\mathcal{W}}$ in two dimensions, i.e., $p = 3$. Each triangle is an element.	188
B.3 State space of the approximating Markov chain with $\delta = 0.05$ for the NYC application. Black dots represent the grid $\bar{\mathcal{W}}$ and gray circles represents the reflecting boundary $\partial\bar{\mathcal{W}}^+$	192
B.4 Two examples of controls in the approximating Markov chain method. Each control in the approximating Markov chain method results in a random transition to one of two neighboring states.	192
B.5 The discretized control directions for the NYC application.	198
B.6 The control policies derived by our computational method and the Markov chain approximation method for the NYC application ($\delta = 0.01$).	200

LIST OF TABLES

1.1	Improvement in various parameters of interest under spatial pricing and removing friction (prices deviation is limited to 50%)	4
1.2	Consumer surplus under origin-only optimal spatial prices, with maximum price variation $\bar{\eta} = 0.2$	30
1.3	Impact of maximum price variation on various performance metrics (maximum price multiplier $\hat{\eta} = 5$)	31
1.4	Comparison between spatial prices, removing local search friction, and the hybrid mechanism ($\bar{\eta} = 0.5$)	37
2.1	Pushing controls used in the solution (derived by our computational method) to the EWF (2.47)-(2.50) for the NYC application.	89
3.1	Simulation output.	131
3.2	Opt. gap of the static policy.	132
3.3	Opt. gap of the proposed policy.	132
3.4	Optimality gaps.	132
A.1	Parameters used in Problem (P0).	140
A.2	Summary of the primitives.	146
A.3	Cross-validation results.	154
A.4	Impact of maximum price variation when using origin-destination pricing on various performance metrics (maximum price multiplier $\hat{\eta} = 5$)	157
A.5	Impact of origin-only prices on drivers' profit (maximum price variation $\bar{\eta} = 0.2$)	159
A.6	Impact of maximum price variation on drivers' profit (maximum price multiplier $\hat{\eta} = 5$)	160
B.1	The list of the relocation activities of the NYC application.	173
B.2	Pushing controls used under the control policy derived by the Markov chain approximation method for the EWF (2.47)-(2.50) of the NYC application ($\delta = 0.01$)	200
B.3	The expected discounted cost of the control policies derived by our computational method and the Markov chain approximation method for different discretization step sizes.	201
B.4	The expected discounted cost ΛJ_z in millions of dollars under the proposed policy, the static policy, and the original alternative policies (the values denoted in the parentheses indicate the 95% confidence interval)	207
B.5	The expected discounted cost ΛJ_z in millions of dollars under the proposed policy, the static policy, and the modified alternative policies (the values denoted in the parentheses indicate the 95% confidence interval)	208
B.6	The expected discounted cost per driver J_z in thousands of dollars under the proposed policy, the static policy, and the original alternative policies (the values inside the parentheses indicate the 95% confidence interval)	208

B.7 The expected discounted cost per driver J_z in thousands of dollars under the proposed policy, the static policy, and the modified alternative policies (the values inside the parentheses indicate the 95% confidence interval).	209
B.8 The expected discounted cost in millions of dollars under the proposed policy and the percentage difference of the static policy as well as the original alternative policies from the proposed policy (the values inside the parentheses indicate the 95% confidence interval).	209
B.9 The expected discounted cost in millions of dollars under the proposed policy and the percentage difference of the static policy as well as the modified alternative policies from the proposed policy (the values inside the parentheses indicate the 95% confidence interval).	210
B.10 Percentage of the demand dropped under the proposed policy, the static policy, and the original alternative policies (the standard deviation of all estimates is less than or equal to 0.01%).	210
B.11 Percentage of the demand dropped under the proposed policy, the static policy, and the modified alternative policies (the standard deviation of all estimates is less than or equal to 0.01%).	211

ACKNOWLEDGMENTS

I am deeply indebted to my advisors Baris Ata and Sunil Kumar. This dissertation would not have been possible without their guidance and support. They offered encouragement and direction precisely when it was needed the most. Baris provided invaluable technical and expository guidance. He is an excellent mentor, a great teacher, and an outstanding advisor. I learned almost everything I know about research and teaching from him. His influence goes beyond this dissertation. I wish to follow in his footsteps and become a supportive advisor after my doctoral studies. Sunil provided invaluable guidance in defining the research problem and controlling the direction of the work. His broad interests and intellectual curiosity have served as a source of motivation.

I am also grateful to Ozan Candogan, Varun Gupta, Amy Ward, and Yuan Zhong for graciously agreeing to serve on my dissertation committee. Their careful reading of this dissertation and invaluable comments significantly improved the exposition of the dissertation. I also would like to thank Hossein Abouee-Mehrizi, Ravi Mazumdar, and Catherine Rosenberg for introducing me to the field of operations management and for their help and support prior to my doctoral studies.

I also owe a large debt to my family, friends, and colleagues. I wish to thank my loving parents Abdolhossein and Zahra and my dear brother Moein for their constant unconditional love and support. They were a beacon of hope in the darkest and most uncertain moments of my journey. This dissertation is dedicated to them. I am also grateful to my dear friends and classmates Deniz Akturk, Amir Alwan, Kevin Chen, Yifan Feng, Charlie Hannigan, Robert Montgomery, Xiaoshan Peng, Cem Randa, Adam Schultz, Luyi Yang, and Gizem Yilmaz for their support, encouragement, and helpful advice. I feel privileged to have had such friends and classmates. I am also thankful to Shirsho Biswas, David Finer, and Tesary Lin for helpful discussions about my research.

ABSTRACT

This dissertation focuses on the control of service and manufacturing systems through the two levers of pricing and matching. Chapters 1 and 2 focus on the ride-hailing industry. Chapter 3 studies a make-to-stock manufacturing system. This dissertation is the culmination of Ata and Barjesteh [20], Ata et al. [21], and Ata et al. [22].

In Chapter 1, we study how spatial pricing and search friction can impact the taxi market in New York City. We use a mean field model, in which the taxi drivers strategically search for customers in different neighborhoods across the city, taking into account the spatial and temporal distribution of the supply and demand as well as the prices across the city. Our model captures the interplay between spatial pricing, where prices depend on either the origin of the ride alone or both its origin and destination, and search friction, due to empty taxis and customers within the same neighborhood failing to pair efficiently. Spatial pricing can incentivize relocation of empty taxis to a neighborhood while the use of mobile applications can alleviate search friction within that neighborhood. We fit our model to a dataset of New York City taxi rides over four years and conduct a series of counterfactual studies to explore how spatial pricing impacts demand for and supply of rides, consumer welfare, and drivers' profit. Our analysis reveals that spatial prices that only use origin information can increase consumer surplus by 7.0% of the average fare and serve 2.6% more customers without hurting the drivers' profit. Moreover, we find that eliminating the (local) search inefficiency alone can increase consumer surplus by 13.9% of the average fare and serve 4.3% more customers while simultaneously increasing drivers' profit by 2.5% of the average fare. We also observe that improving search efficiency primarily impacts under-served neighborhoods such as upper Manhattan, Brooklyn and Queens, while pricing primarily impacts well-served neighborhoods, for example, the airports, midtown, and downtown Manhattan. This underscores the value of a hybrid mechanism. We propose a mechanism in which (local) search is eliminated in all neighborhoods while spatial pricing is only used in well-served neighborhoods. This mechanism increases consumer surplus by

21.5% of the average fare and serves 8.7% more customers, while avoiding price discrimination in less affluent neighborhoods of the city. The proposed mechanism achieves 96.3% of the benefits of a citywide spatial pricing and friction removal mechanism.

In Chapter 2, we consider a ride-hailing platform that seeks to maximize its profit by dynamically dispatching cars to pick up customers and centrally relocating cars from one area to another. We model the ride-hailing platform as a closed stochastic processing network. Because the problem appears intractable, we resort to an approximate analysis in the heavy-traffic regime and consider the resulting Brownian control problem. This problem is simplified considerably and reduced to a lower-dimensional singular control problem called the workload formulation. We develop a novel algorithm to solve the workload problem numerically. We apply this algorithm to the workload problem derived from the New York City taxi dataset. The solution helps us derive a dynamic control policy for the New York City application. In doing so, we prescribe the ride-hailing platform to first solve an offline linear program, whose optimal solution can be interpreted as the optimal static control policy. This solution helps partition the areas of the city into pools of areas. The platform only uses the information on the fraction of cars in the various pools, which reduces the state space dimension significantly, making the problem computationally tractable. When the distribution of cars among the pools is balanced, the platform follows the optimal static control policy. Otherwise, the platform intervenes to move the system to a more balanced state by either dropping demand or using a dispatch or relocation activity that is not used under the optimal static control policy. We demonstrate the effectiveness of the proposed dynamic control policy for the New York City application using a simulation study.

In Chapter 3, we consider a make-to-stock manufacturing system selling multiple products to price-sensitive customers. The system manager seeks to maximize the long-run average profit by making dynamic pricing, outsourcing, and scheduling decisions: First, she adjusts prices dynamically depending on the system state. Second, when the backlog of work is judged excessive, she may outsource (or reject) new orders thereby incurring outsourcing

costs. Third, she decides dynamically on which product to prioritize in the manufacturing process, i.e., she makes dynamic scheduling decisions. This problem appears analytically intractable. Thus, we resort to an approximate analysis in the heavy-traffic regime and consider the resulting Brownian control problem. We solve this problem explicitly by exploiting the solution to a particular Riccati equation. The optimal solution to the Brownian control problem is a two-sided barrier policy with drift rate control: Outsourcing and idling processes are used to keep the workload process above the lower reflecting barrier and below the upper reflecting barrier, respectively. Between the two barriers, a state-dependent drift rate is used to control the workload process.

CHAPTER 1

AN EMPIRICAL ANALYSIS OF TAXI RIDES IN NEW YORK CITY

1.1 Introduction

Taxi industry is an essential part of the transportation sector.¹ For example, in New York City (NYC), taxis offer over 150 million rides per year (TLC 1). In this market, taxis and customers search for each other. This search friction results in a substantial welfare loss. Prices set by the taxi industry affect the inter-temporal and spatial distribution of the supply and demand and the interaction between them, impacting the consumer surplus and drivers' profit. This chapter explores spatial pricing, a mechanism that prices rides based on their origin and destination. In this context, we study the following questions: How does spatial pricing impact consumer surplus and drivers' profit? What is the pattern of the optimal spatial prices? How does pricing based solely on the origin of the ride compare with pricing based on the origin-destination pair? How does spatial pricing compare with removing the local search friction using mobile applications?²

We consider two inefficiencies: First, the mismatch between supply and demand. Second, the (local) search friction. To address these, we consider spatial pricing and friction removal (through a better matching technology) as levers. Although spatial pricing³ is not widely used in the taxi industry or ride-sharing platforms, it can help match supply and demand. Supply can be redistributed since the profitability of different neighborhoods is closely tied to prices and taxi drivers make relocation decisions to maximize their profit. The distribution of demand can also be adjusted since customers are price sensitive. By adjusting both supply

1. With an annual revenue of \$18.9 billion and annual profit of \$1.5 billion, taxi and limousine industry is one of the major segments of the US economy; see e.g., Sayler [185].

2. Internet-based mobile applications such as Arro and Curb can provide a better match between customers and taxis; see <https://www.ridearro.com> and <https://main.gocurb.com>.

3. Spatial pricing should not be confused with dynamic pricing.

and demand, spatial pricing provides the policy maker with a unique and powerful tool to intervene in the details of the market. Moreover, spatial pricing can be implemented with the existing equipment⁴ and technological solutions such as mobile applications.

Spatial pricing does not resolve all the inefficiencies. The search by customers for taxis is spatially localized. A customer looks for a taxi only within a relatively small geographical area at any given time. Although empty taxis have more freedom and may relocate to different areas to seek customers, within any single area their ability to connect with customers is imperfect. As the spatial densities of customers and taxis increase, a higher fraction of customers are matched with taxis. However, it is not uncommon to find both empty taxis and unfulfilled customers in the same area. On the one hand, spatial pricing can impact the relocation decisions of the empty taxis. This in turn, impacts the spatial density of taxis, and consequently, the efficacy of the search. On the other hand, mobile (search) applications can improve the search efficacy as well, enabling better connections between empty taxis and customers even in locations with low density of customers and empty taxis, as measured by the number of customers and the number of empty taxis per street mile, respectively. Relocation of empty taxis induced by spatial pricing and the reduction of search friction complement each other. Understanding the interplay between spatial pricing and the removal of (local) search friction is the primary focus of this chapter. Although removing the (local) search friction resolves the issue of having unfulfilled customers and unutilized taxis simultaneously at the same location, it is merely a local solution. It has little to no impact on neighborhoods with high density of supply and demand, where the majority of rides initiate (due to the already high efficacy of search in these neighborhoods). In other words, removing the (local) search friction does not address the global mismatch between supply and demand; see e.g., Lagos [149].

Our model captures the interaction of supply and demand both locally and globally. To

4. Smart meters installed on all NYC yellow taxis (since 2009) report all the information (longitude and latitude of the pick-up and drop-off locations) required for the implementation of spatial prices. Mobile applications such as Arro and Curb can inform customers of their fare before they hail a taxi.

capture the global aspect, we use a mean field model that approaches the problem from a macroscopic perspective. In this model, an individual taxi driver is irrelevant and the focus is on the distribution of the taxis. To capture the local aspect, we use an aggregate matching function that captures the spatial and microscopic aspect of the search.

We take the view of a social planner and optimize total consumer surplus. We observe that a spatial pricing scheme that prices rides based on their origin and destination and allows prices to change between 50% and 150% of the current prices⁵ in NYC can increase consumer surplus by \$168,000 in every day shift (\$0.79 per ride or 8.5% of the average fare), serve 3.2% more customers, and increase customer miles (total number of miles traveled by customers) by 7.4%, without hurting the drivers' profit. A similar spatial pricing scheme that only uses origin information can increase consumer surplus by \$135,000 in every day shift (\$0.63 per ride or 7.0% of the average fare), serve 2.6% more customers, and increase customer miles by 3.9%.

Prices in under-served areas (such as upper Manhattan, Brooklyn, and Queens)⁶ increase to attract more taxi drivers. Due to higher prices, each served customer in under-served areas is worse off. However, this effect is offset by the increase in the number of served customers. In well-served areas, prices increase in nodes with low demand or shorter trips and decrease in nodes with high demand and longer trips.

Considerably higher prices in the less affluent neighborhoods of the city (upper Manhattan, Brooklyn, and Queens) is an undesirable outcome of using spatial pricing alone. This effect can be mitigated by removing the (local) search friction. Table 1.1 compares the impact of origin-only spatial pricing (with price deviation smaller than 50%) and removing friction on various metrics of interest. Removing the (local) search friction alone (with no spatial pricing) can increase consumer surplus by \$1.26 per ride, serve 4.3% more customers, and increase drivers' profit by \$0.22 per ride. The majority of the benefits of removing the

5. In this dissertation, we do not change the ratio of the price per mile to the fixed portion of the fare. Instead, we use price multipliers to change the fares; see Section 1.6.

6. See Figure A.1 for the definition of these geographical areas.

Table 1.1: Improvement in various parameters of interest under spatial pricing and removing friction (prices deviation is limited to 50%).

	Origin-Only Pricing	Origin-Destination Pricing	Removing Local Search Friction	Hybrid Mechanism	Proposed Mechanism	Citywide Origin-Only Pricing & Friction Removal
Consumer Surplus						
Total increase	\$135K	\$168K	\$268K	\$322K	\$417K	\$433K
Per ride	\$0.63	\$0.79	\$1.26	\$1.52	\$1.96	\$2.04
In terms of average fare	7.0%	8.5%	13.9%	16.7%	21.5%	22.4%
Number of served customers	2.6%	3.2%	4.3%	5.7%	8.7%	8.9%
Miles traveled by customers	3.9%	7.4%	6.2%	8.1%	11.7%	12.0%
Drivers' Profit	\$0	\$0	\$48K	\$0	\$0	\$0

(local) search friction come from under-served nodes while the majority of the benefits of spatial pricing are from well-served nodes. This highlights the value of a hybrid mechanism. A hybrid mechanism, that uses spatial pricing in well-served neighborhoods and removes friction in under-served neighborhoods (keeping the price pattern), can increase consumer surplus by \$1.52 per ride. Under the hybrid mechanism, we need only a little price variation to achieve the majority of the benefits. Table 1.1 also presents the impact of citywide spatial pricing and friction removal. This mechanism (with only 50% price variation) increases consumer surplus by \$2.04 per ride, which is considerably higher than spatial pricing or removing friction alone. Since policy makers prefer avoiding price discrimination in less affluent neighborhoods of the city, we propose a mechanism in which friction is removed in the entire city while spatial pricing is used only in well-served neighborhoods. The mechanism captures almost all the benefits of citywide spatial pricing and friction removal.

Spatial prices can increase drivers' profits, as well. Spatial prices, that only use origin information and do not deviate from the current prices in NYC by more than 50%, can increase drivers' profits by \$100,000 in every day shift (\$0.47 per ride or 5.3% of the average base-fare), without lowering consumer surplus.

The rest of the chapter is organized as follows. Section 1.2 reviews the literature. Section 1.3 introduces the mean field model. Section 1.4 describes the data. Section 1.5 describes the estimation procedure and results. Section 1.6 describes the counterfactual anal-

ysis and Section 1.7 concludes. Appendix A.1 discusses the procedure used to define the nodes/neighborhoods. Appendix A.2 provides complementary discussions on the matching model. Appendices A.3-A.4 provides supplementary material for the Data and Estimations sections, respectively. Appendix A.5 describes a Monte Carlo simulation study to illustrate the identification of our model. Appendix A.6 uses five-fold cross-validation to examine the ability of our model in predicting the relocation decisions of the drivers. Appendix A.7 provides supplementary material for the counterfactual analysis and Appendix A.8 studies spatial pricing for maximizing drivers' profit. Appendix A.9 provides the proofs and derivations.

1.2 Literature Review

This chapter is related to four streams of literature. The first stream focuses on mechanisms and regulations used for improving the performance of the taxi market. The second stream studies spatial models of search with an emphasis on the ride-hailing industry. The third stream focuses on dynamic discrete choice models. The last stream studies mean field games and their applications.

Mechanisms and regulations used for improving the performance of the taxi market have been studied extensively. Entry restrictions and price controls are the most studied mechanisms in the literature; see e.g., Coffman and Shreiber [75], Foerster and Gilbert [91], Schroeter [187], and Häckner and Nyberg [105]. The common theme in many of these papers is that price regulations and entry restrictions are helpful since they increase the availability of taxis in times and locations with low demand.⁷ Frechette et al. [92] follows this literature by showing that search frictions and entry restrictions are important inefficiencies in the taxi market, and one reason for the success of ride-sharing platforms is the fact that they can address these issues. Similar to Frechette et al. [92], we observe that search friction is

7. This is achieved by reducing the friction of price negotiation and ensuring a suitable minimum profit for the drivers. For example, in exchange for serving all neighborhoods of a city, a firm could be granted a monopoly position.

a major issue and removing it increases consumer surplus and drivers' profit by \$1.26 and \$0.22 per ride, respectively.

Following the success of ride-sharing platforms, taxi industry has received increasing scrutiny. Cramer and Krueger [79] shows that the utilization of Uber drivers is higher than the utilization of taxi drivers. They argue that this is in part due to Uber's efficient matching technology and inefficient regulations in the taxi industry. Buchholz [56] proposes a model of spatial search of taxi rides and computes the gain in consumer surplus from perfect matching of customers and taxis in each neighborhood. Buchholz concludes that although the elimination of the (local) search friction in NYC results in a 7.1% increase in the number of served customers, it has a negative impact on consumer surplus. In contrast, we observe that the elimination of the (local) search friction results in a 4.3% increase in the number of served customers while increasing consumer surplus by \$1.26 per ride. The majority of these benefits come from neighborhoods with low density of supply and demand. Our observation is consistent with the conclusions made in Lam and Liu [150] and Shapiro [189] that a platform that posses a superior matching technology outperforms NYC yellow taxis primarily in neighborhoods with low density of supply and demand.

Another class of mechanisms closely related to this chapter is dynamic and spatial pricing in ride-hailing networks. This literature can be divided into spatial and non-spatial treatments of the problem. Banerjee et al. [41], Bai et al. [36], Özkan and Ward [172], and Cachon et al. [61] use a non-spatial analytical approach. Banerjee et al. [41] models the problem of dynamic pricing of rides in a single region as a queueing system. It shows that the throughput and revenue of no dynamic pricing strategy can exceed that of the optimal static pricing policy. However, dynamic pricing strategies are more robust to fluctuations in the system parameters. Relevant empirical non-spatial studies include Hall et al. [106], Cohen et al. [76], Ming et al. [166], Lam and Liu [150], and Shapiro [189]. Cohen et al. [76] uses data on Uber rides in four major cities in the United States and estimates that for each dollar spent by customers, 1.6 dollars of consumer surplus is generated. Lam and Liu

[150] uses a dynamic choice model in which customers in NYC choose between Uber, Lyft, and Taxi rides. It finds that customers who use ride-sharing platforms gain 0.72 dollars for every dollar spent on rides and 64% of this gain is due to dynamic pricing. Although this chapter focuses on spatial pricing and we observe more modest increases in consumer surplus from pricing and using a better matching technology (e.g., an increase in consumer surplus of \$2.04 per ride, 22.4% of the average ride, from city-wide spatial pricing and friction removal), important qualitative insights of Lam and Liu [150] agree with ours. For example, Lam and Liu [150] observes that pricing is more welfare-enhancing in thick markets (such as midtown Manhattan) while the matching technology is most beneficial in the outer boroughs.

The second stream of literature studies spatial search models. Lagos [149] is one of the first papers that studied the taxi market with an emphasis on its spatial aspect and the strategic behavior of its drivers. Lagos [149] proposed a model in which taxi drivers search for customers on a graph, and highlighted the friction resulting from this search. A number of papers followed this work, by generalizing it, or using it to address other issues in the ride-hailing industry.

Bimpikis et al. [52] builds on Lagos [149] to study spatial pricing in ride-sharing platforms. It shows that if the demand pattern is not balanced, spatial pricing is beneficial and optimal prices can be written in terms of the optimal dual variables corresponding to the flow balance equations. Under the assumption that prices and compensations can be decoupled, the authors are able to solve for the optimal prices by focusing on the mass balance equations only. Since in the taxi market in NYC the entire fare is collected by the drivers and the Taxi and Limousine Commission (TLC) does not collect a fee, prices and compensations cannot be decomposed in our setting. Hence, no such simplification occurs in our analysis. We allow for a non-complete graph with different distances/travel times between nodes, a non-uniform demand pattern on the graph, and inter-temporal variations in the system.

Buchholz [56] builds upon Lagos [149] to study the pricing of taxi rides based on pick-up location, time of day, and the distance of the ride. It concludes that distance-based pricing

outperforms pricing based on the pick-up location and time of day. Buchholz finds that distance-based pricing outperforms origin-only pricing by 766% with respect to the increase in consumer surplus and by 356% with respect to the increase in the number of served customers. Although we do not analyze distance-based pricing, the origin-destination pricing subsumes distance-based pricing. However, we see a more modest increase with respect to consumer surplus (24%) and the number of served customers (23%) in origin-destination compared to origin-only pricing.

Other related papers in the spatial literature are Braverman et al. [55], Banerjee et al. [39], Yang et al. [218], Afèche et al. [10], and Besbes et al. [50]. Braverman et al. [55] studies centralized empty car routing in ride-shairng platforms. It shows that a fluid-based optimization can be used to solve for the optimal network utility, which is an upper bound on the utility of all static and dynamic routing policies in the finite-car system. Banerjee et al. [39] studies pricing in shared vehicle systems where taxis do not relocate when they are empty. It proposes a static spatial pricing algorithm with an approximation ratio that improves as the average number of vehicles per location grows. Afèche et al. [10] takes prices as fixed and studies admission control of customers and centralized re-positioning of drivers and shows that the value of these control policies are largest at moderate capacity and they increase with demand imbalances. Besbes et al. [50] uses a stylized model on a line, where taxis make myopic decisions based on the next ride and can relocate instantaneously. They show that the pricing problem can be spatially decomposed based on the attraction regions. The platform can use prices to create regions in which driver congestion is artificially high in order to motivate drivers to relocate to more profitable regions.

The third stream of literature studies discrete choice theory and structural estimation; see Ben-Akiva et al. [46] for an introduction and Anderson et al. [16] for examples. The most relevant papers in this area are Rust [183], Nair [169], and Su and Judd [194]. Rust [183] studies a structural model in which a manager has to decide in each period whether to replace a bus engine or postpone the decision. Nair [169] studies an equilibrium model in which

consumers decide when to purchase a durable product based on their expectation of future prices and the seller decides on the optimal temporal pricing scheme. In a setting similar to Nair [169], in this chapter, taxi drivers make relocation decisions based on their expectation of future supply and demand, and the central planner sets the prices hoping to impact the relocation decisions of the drivers. Su and Judd [194] shows that structural estimation problems can be viewed as constrained optimization problems. We use an approach similar to Su and Judd [194] in Sections 1.5-1.6. Other relevant papers in this area include Akşin et al. [13, 14]⁸, Ata et al. [23, 24], Li et al. [155], Zheng [222], and Zheng et al. [223].

The last stream of literature studies mean field games. The mean field games theory was developed to study systems with an infinite number of rational agents in competition. The theory was developed independently by the mathematics community in Lasry and Lions [152] and by the engineering community in Huang et al. [131]. This framework has been used to model various economic and engineering systems. Examples of such papers in the operations community include Adlakha and Johari [5], Xu and Hajek [215], Iyer et al. [133], Adlakha et al. [6], Gummadi et al. [103], and Balseiro et al. [37].

1.3 Model

This section introduces the empirical model used to estimate the effect of spatial pricing of taxi rides in NYC. We index the months between January 2010 and December 2013 by $k \in \{1, \dots, K\}$ ($K = 48$). In what follows, we fix k and focus on non-holiday weekdays of month k . We assume that the primitives of the model are the same in all weekdays of the same month. In other words, although there is variation across months, we assume that the non-holiday weekdays in a month are i.i.d. copies of each other. Consider a connected directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with n nodes, i.e. $|\mathcal{V}| = n$. Each node represents an area in Manhattan, Brooklyn, Queens or one of the airports (La Guardia and JFK). Any pair of distinct nodes

8. Ata et al. [25] and Ata and Peng [32] rigorously establish the existence and uniqueness of the equilibria in a similar setting.

are connected with at most two directed edges, denoted by (i, j) and (j, i) , respectively. Furthermore, each node has a loop, which represents travel within the area/node.

Let $t \in \{1, \dots, T\}$ index time. Each t represents a five-minute interval between the hours of 6_{AM} and 4_{PM} (a typical NYC day-time shift), corresponding to $T = 120$. For all $i, j \in \mathcal{V}$, the average time and distance to travel from node i to node j are denoted by τ_{ij} and d_{ij} , respectively.⁹ We allow (τ_{ij}, d_{ij}) to be different from (τ_{ji}, d_{ji}) ¹⁰ and do not impose a relationship between τ_{ij} and d_{ij} . We assume that τ_{ij} is integer-valued for all $i, j \in \mathcal{V}$ and $\tau_{ii} = 1$ for all i . To facilitate the analysis to follow, define

$$S_{ij} = \begin{cases} 1 & \text{if } (i, j) \in \mathcal{E}, \\ 0 & \text{otherwise.} \end{cases}$$

Let P_{ij} and F_{ij} denote the price per mile and the fixed portion of the fare, respectively, for a ride from node i to node j . The fare paid by a customer to his driver for a ride from node i to node j is $F_{ij} + P_{ij}d_{ij}$. The fixed portion is used to model the flag-drop and the JFK flat fare.¹¹ Taxes and tolls are not included in the calculations as they are fixed (over time) and taxis do not keep them. The price per mile and the fixed portion of the fare are allowed to depend on the origin and the destination of the ride in order to capture the flat fare structure of rides between JFK and Manhattan.¹² This also gives us sufficient flexibility to explore the origin-destination pricing in our counterfactual study.

Given prices $F = [F_{ij}]_{i,j=1}^n$ and $P = [P_{ij}]_{i,j=1}^n$, the arrival rate (per period) of potential

9. If $(i, j) \in \mathcal{E}$, τ_{ij} and d_{ij} are the travel time/distance on the edge connecting i to j . Otherwise, τ_{ij} and d_{ij} correspond to the travel time/distance of full taxis, that are calculated from the data.

10. This allows us to capture the spatial variation in the traffic speed as well as the impact of one-way streets and streets with different traffic speeds in each direction.

11. The fare of a ride from JFK to Manhattan (and vice versa) after September 2012 is \$52 plus the \$4.5 rush hour surcharge (4PM to 8PM on weekdays, excluding legal holidays). Since this chapter focuses on the day shift (6AM to 4PM) on weekdays, the rush hour surcharge is not included in the calculations.

12. Consider a customer at JFK airport at a non-rush hour time in September of 2012. If he is headed to Manhattan, his fare is \$52. This corresponds to $F_{ij} = 52$ and $P_{ij} = 0$. However, if he is headed to another location, for example to La Guardia airport for a connecting flight, his fare is calculated based on the distance he intends to travel. In this case, $F_{ij} = 2.5$ and $P_{ij} = 2.5$ per mile.

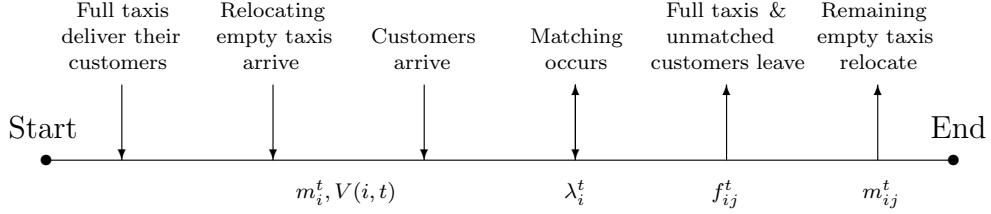


Figure 1.1: Timing of events and definition of variables in each period.

customers who wish to go to node j from node i in period t , denoted by $\Lambda_{ij}^t(F, P)$, is given as follows:

$$\Lambda_{ij}^t(F, P) = A_{ij}^t [F_{ij} + P_{ij}d_{ij}]^\alpha \exp(\beta k) \quad \text{for all } i, j, t. \quad (1.1)$$

The parameter A_{ij}^t is the fixed effect of demand and $\alpha < 0$ is the price-elasticity. As will be shown in Section 1.4, our data exhibits a trend over time (see Figure 1.3) and β captures the monthly trend in demand. Since α does not depend on the fare, the demand model in (1.1) is a constant-elasticity demand model.¹³ The (potential) demand at node i in period t , denoted by $\Lambda_i^t(F, P)$, is given by

$$\Lambda_i^t(F, P) = \sum_{j=1}^n \Lambda_{ij}^t(F, P) \quad \text{for all } i, t.$$

The probability that a customer at node i is headed to node j , denoted by $\pi_{ij}^t(F, P)$, is equal to

$$\pi_{ij}^t(F, P) = \frac{\Lambda_{ij}^t(F, P)}{\Lambda_i^t(F, P)} \quad \text{for all } i, j, t. \quad (1.2)$$

In what follows, we denote satisfied demand at node i in period t by λ_i^t . Note that $\lambda_i^t \leq \Lambda_i^t(F, P)$.

The timing of events in each period is depicted in Figure 1.1. We denote the number of active taxis during the day shift by M and the average of M across the 48 month by \bar{M} .

13. See Van Zandt [199, Page 152] for a discussion on constant-elasticity demand models.

We normalize the number of active taxis to an M/\bar{M} mass (as done usually in mean field models).¹⁴ The mass of empty cars at node i at the begining of period t is denoted by m_i^t . Also, let m_{ij}^t denote the mass of empty cars that decide to relocate from node i to node j at the end of period t , and f_{ij}^t denote the mass of full cars that picked up a customer at node i , who wishes to go to node j , in period t .

An empty taxi at node i in period t that could not pick up a customer relocates to node j with probability q_{ij}^t . In particular, it stays idle at node i until the next period with probability q_{ii}^t . These probabilities arise endogenously as drivers make their relocation decisions. Letting $\mathcal{A}(i) = \{j : S_{ij} = 1\}$ denote the set of nodes that can be reached from node i , an empty taxi at node i can only relocate to nodes in $\mathcal{A}(i)$. Moreover, since $q_i^t = \{q_{ij}^t : j \in \mathcal{A}(i)\}$ is a probability distribution for all i , it follows that

$$\begin{aligned} \sum_{j \in \mathcal{A}(i)} q_{ij}^t &= 1 && \text{for all } i, t, \\ q_{ij}^t &= 0 && \text{for all } j \notin \mathcal{A}(i), \\ q_{ij}^t &\geq 0 && \text{for all } i, j, t. \end{aligned}$$

Let c denote the mean travel cost per mile (incurred by the drivers). We assume that the cost of traveling from node i to node j is equal to $c d_{ij} + \epsilon_{ij}$, where ϵ_{ij} denotes a Gumbel Min¹⁵ (minimum extreme value type I) distributed idiosyncratic shock to the travel cost of the driver. The idiosyncratic cost shocks correspond to unobservable variables in the empirical industrial organization literature (see, e.g., Rust [183]), which are observed by the driver but not recorded in the data. The cost shocks are assumed to be i.i.d. with mean zero and scale parameter σ . Furthermore, we let $V(i, t)$ denote the value function of an empty taxi's driver at node i in period t (before customers arrive); and N_i denotes the number of

14. Using M/\bar{M} as opposed to the unit mass in all months allows us to capture the variation in the number of active taxis while ensuring that all quantities of interest are comparable across months.

15. A Gumbel Min distributed random variable has the same distribution as the negative of a Gumbel Max (commonly referred to as Gumbel) distributed random variable.

regions in node i . What constitutes a region will be discussed below in detail.

For the remainder of this section, we assume that the problem primitives \bar{M} , M , N_i , c , F_{ij} , P_{ij} , τ_{ij} , d_{ij} , S_{ij} , A_{ij}^t , α , β , k , σ , and the initial distribution of empty taxis, m_i^1 , are given for all i, j, t . We then proceed to characterize λ_i^t , m_i^t , m_{ij}^t , f_{ij}^t , q_{ij}^t , and $V(i, t)$ by deriving a set of equations they must satisfy. To characterize the satisfied demand λ_i^t , we focus on the following question: *If there are m taxis and Λ customers at the beginning of a period in an area (node), what is the expected number of served customers (matches between taxis and customers)?* The answer depends on whether the matching of customers and taxis is perfect (frictionless) or imperfect (with friction). Matching can be perfect in nodes where customers and taxis are matched through a mobile application. In most other cases, the matching is imperfect and taxis need to search for customers block by block.

Perfect Matching. When matching is perfect, the number of matches is the minimum of the number of taxis and customers in the node in that period, i.e.

$$\lambda_i^t = \min(m_i^t, \Lambda_i^t(F, P)). \quad (1.3a)$$

When matching is imperfect, there can be unfulfilled customers and unutilized taxis simultaneously. We propose the following matching model in this case. Consider a node partitioned into $N \geq 2$ regions, which can be done so that it takes one period (i.e. five minutes) for a taxi to explore any of the regions.¹⁶ We assume that customers and taxis randomly choose one of N regions (with equal probabilities) to explore and the number of matches in each region is the minimum of the number of customers and taxis in that region. The total number of matches in the node is equal to the sum of the matches made in the N regions.

Since customers and taxis choose regions with equal probabilities, all regions have the

16. Regions are the smallest geographic units in our model and they are obtained by dividing the street miles in a node such that each region takes exactly one period to explore. One way to compute the number of regions N is to use the data on the average speed of taxis and the total street-miles in that node.

same expected number of matches. Consider an arbitrarily chosen region. Let X and Y denote the number of taxis and customers that chose that region. Since there are m taxis and each taxi chooses this region with probability $1/N$, we have $X \sim \text{Binomial}(m, 1/N)$. Similarly, $Y \sim \text{Binomial}(\Lambda, 1/N)$. Therefore, the expected number of matches in the region is equal to $\mathbb{E}[\min(X, Y)]$, and that in the node is equal to $N \times \mathbb{E}[\min(X, Y)]$. We would like to find a tractable approximation for $N \times \mathbb{E}[\min(X, Y)]$ that is monotone in Λ and equals zero when $\Lambda = 0$. To do so, we use a normal approximation to the Binomial distribution that is further refined by a linear approximation for small values of demand to ensure monotonicity.¹⁷

Imperfect Matching. When matching is imperfect, we propose the following model:

$$\lambda_i^t = \begin{cases} G(\Lambda_i^t; m_i^t) & \text{if } \Lambda_i^t \geq \hat{\Lambda}_{m_i^t}, \\ G(\hat{\Lambda}_{m_i^t}; m_i^t) \frac{\Lambda_i^t}{\hat{\Lambda}_{m_i^t}} & \text{otherwise,} \end{cases} \quad (1.3b)$$

where $G(\Lambda; m) = \Lambda \Phi(\nu) + m \Phi(-\nu) - (m - \Lambda) \phi(\nu)/\nu$ with $\Phi(\cdot)$ and $\phi(\cdot)$ denoting the cdf and pdf of the standard normal distribution, respectively,

$$\nu \triangleq \sqrt{\frac{\bar{M}}{N_i - 1}} \frac{m - \Lambda}{\sqrt{m + \Lambda}},$$

and $\hat{\Lambda}_m = \min \{ \Lambda \geq 0 : G(\Lambda; m) \text{ is strictly increasing on } (\Lambda, \infty) \}$; see Appendix A.9 for the derivation of (1.3b).

The term $G(\Lambda; m)$ in (1.3b) is derived using a normal approximation for the Binomial distributions. Although this approximation elegantly captures the key features of search friction that we seek to model, it is not monotone in a neighborhood of the origin and fails to satisfy $\lambda_i^t = 0$ when $\Lambda_i^t = 0$. The linear approximation at $\Lambda_i^t < \hat{\Lambda}_{m_i^t}$ in (1.3b) remedies

17. Our numerical results are robust to the approximation used in small values of demand. This approximation is required at demand values much smaller than one customer per period, which are observed in less than 1% of the (i, t) pairs in the numerical experiments of Section 1.6.

these two issues and retains the other desirable features of the normal approximation.

Buchholz [56] tailors the imperfect matching model proposed in Burdett et al. [59] to the taxi market. In Buchholz [56]'s model, first the number of drivers and their locations are revealed to the customers. Then, each customer chooses a taxi. When more than one customer chooses a taxi, only one of the customers is served and the others leave the system unfulfilled. In our matching model, taxis and customers in each region observe each other (and not the taxis and customers in other regions) and matching in each region is frictionless. Due to this frictionless matching in each region, our matching model results in a higher expected number of matches when the number of taxis and customers are sufficiently higher than the number of regions (nodes with high density of supply and demand). However, when the number of taxis or customers is low (nodes with low density of supply or demand), due to the fact that taxis can be matched only with customers in their region, our matching model results in a lower number of matches. For a detailed comparison of our matching model with the matching model of Buchholz [56], see Appendix A.2.

Next, we describe the flow balance equations and the Bellman equation, governing the system dynamics and the relocation decisions of empty taxis, respectively. These equations hold under both the perfect and imperfect matching.

Flow Balance of Full Cars. The mass of empty cars that picked up a customer at node i headed to node j in period t , f_{ij}^t , must be equal to λ_i^t , the satisfied demand at node i in period t , multiplied by the probability that the customers are headed to node j , $\pi_{ij}^t(F, P)$. Therefore, we must have

$$f_{ij}^t = \lambda_i^t \pi_{ij}^t(F, P) \quad \text{for all } i, j, t. \quad (1.4)$$

Flow Balance of Empty Cars. The mass of empty cars that decided to relocate from i to j at the end of period t (after they could not obtain a customer), m_{ij}^t , must be equal to the mass of empty cars after the customers are served, $(m_i^t - \lambda_i^t)$, multiplied by the probability

that a car will relocate to node j , q_{ij}^t ; see Figure 1.1 for the timing of events in a period. Therefore, we must have

$$m_{ij}^t = (m_i^t - \lambda_i^t) q_{ij}^t \quad \text{for all } i, j, t. \quad (1.5)$$

Flow Balance at Nodes. The mass of empty cars in node i at the begining of period t is equal to the sum of the mass of empty cars that arrived at node i in period t and the mass of full cars that dropped their customer at node i in period t . Therefore, we must have

$$m_i^t = \sum_{j \in \mathcal{A}(i)} m_{ji}^{t-\tau_{ji}} + \sum_{j=1}^n f_{ji}^{t-\tau_{ji}} \quad \text{for all } i, t. \quad (1.6)$$

Mass Balance. Since taxis are either empty or full and the mass of (relocating) empty taxis and full taxis at any period t must sum to M/\bar{M} , the total mass of taxis on the graph, we must have

$$\sum_{i=1}^n \sum_{j \in \mathcal{A}(i)} \sum_{s=1}^{\tau_{ij}} m_{ij}^{t-s} + \sum_{i=1}^n \sum_{j=1}^n \sum_{s=1}^{\tau_{ij}} f_{ij}^{t-s} = M/\bar{M} \quad \text{for all } t. \quad (1.7)$$

Next, we describe how taxi drivers make their relocation decisions.

Bellman Equation. Consider an infinitesimal driver and assume that the mass of the empty and full taxis, $[m_{ij}^t]_{i,j=1}^n$ and $[f_{ij}^t]_{i,j=1}^n$, and satisfied demands, $[\lambda_i^t]_{i=1}^n$, are given for all $t \in \{1, \dots, T\}$. The driver moves around on the graph until $t = T$, picks up customers, delivers them to their destination and makes relocation decisions when he can not pick up a customer. The objective of the driver is to maximize his total profit. Recall that the travel cost of the infinitesimal driver from node i to node j is equal to $cd_{ij} + \epsilon_{ij}$, where ϵ_{ij} denotes a Gumbel Min (minimum extreme value type I) distributed idiosyncratic shock to the travel cost of the driver. The cost shocks are assumed to be i.i.d. with mean zero and scale parameter σ .

The state of the empty infinitesimal taxi is $s = (i, t, \epsilon^{(i)})$, where $i \in \{1, \dots, n\}$ denotes the location (node) of the taxi, $t \in \{1, \dots, T\}$ denotes the period, and $\epsilon^{(i)} = (\epsilon_{ij}; j \in \mathcal{A}(i))$ is the vector of iid cost shocks. Each element of $\epsilon^{(i)}$ has a Gumbel Min (minimum extreme value type I) distribution with mean zero and scale parameter σ . We denote the observable (to the researcher) state of the empty taxi by $x = (i, t)$. Therefore, $s = (i, t, \epsilon^{(i)}) = (x, \epsilon^{(i)})$. The driver of the empty taxi at state $s = (i, t, \epsilon^{(i)})$ can choose to relocate to node $j \in \mathcal{A}(i)$. Therefore, the action set of the driver at state $s = (i, t, \epsilon^{(i)})$ is $\mathcal{A}(i)$. The following proposition characterizes the driver's value function and the relocation probabilities; see Appendix A.9 for its proof.

Proposition 1. *The value function of the infinitesimal empty taxi is given by*

$$V(i, t) = \frac{\lambda_i^t}{m_i^t} \left(\sum_{j=1}^n \pi_{ij}^t(F, P) \left(F_{ij} + [P_{ij} - c] d_{ij} \right) + \sum_{j=1}^n \pi_{ij}^t(F, P) V(j, t + \tau_{ij}) \right) + \sigma \left(1 - \frac{\lambda_i^t}{m_i^t} \right) \log \left[\sum_{j \in \mathcal{A}(i)} \exp \left(\frac{V(j, t + \tau_{ij}) - c d_{ij}}{\sigma} \right) \right] \quad \text{for all } i \text{ and } t \leq T, \quad (1.8)$$

and $V(i, t) = 0$ for all i and $t > T$. The relocation probabilities are given by

$$q_{ij}^t = \begin{cases} \frac{\exp \left([V(j, t + \tau_{ij}) - c d_{ij}] / \sigma \right)}{\sum_{l \in \mathcal{A}(i)} \exp \left([V(l, t + \tau_{il}) - c d_{il}] / \sigma \right)} & \text{for } j \in \mathcal{A}(i), \\ 0 & \text{otherwise.} \end{cases} \quad (1.9)$$

Proposition 1 applies to any matching model. The first term on the right-hand side of Equation (1.8) captures the expected value of picking up a customer and the second term captures the expected value of the best relocation decision, where λ_i^t/m_i^t is the probability of picking up a customer. The value of picking up a customer is the expected fare of the ride plus the expected value function at the destination of the ride. Note that the arguments in the exponential function are the scaled values of the value function at the destination of the ride minus the travel cost. For an analysis of strategic relocation of the drivers that is

similar to ours, see Lagos [149] and Buchholz [56].

Next, we define the mean field equilibrium, and Theorem 1 establishes its existence; see Appendix A.9 for its proof.

Definition 1. *Given the problem primitives \bar{M} , M , N_i , c , F_{ij} , P_{ij} , τ_{ij} , d_{ij} , S_{ij} , A_{ij}^t , α , β , k , σ , and m_i^1 , the solution $(\lambda_i^t, m_i^t, m_{ij}^t, f_{ij}^t, q_{ij}^t, V(i, t))$ to Equations (1.3)-(1.9) is called the mean field equilibrium.*

Theorem 1. *Given the problem primitives \bar{M} , M , N_i , c , F_{ij} , P_{ij} , τ_{ij} , d_{ij} , S_{ij} , A_{ij}^t , α , β , k , σ , and m_i^1 , there exists a mean field equilibrium.*

1.4 Data

Our data set is the NYC yellow taxi trip record data spanning four years from January 2010 to December 2013.¹⁸ In this time-span, NYC yellow taxis offered an average of 477,497 rides per day (174.3 Million per year). NYC yellow taxis are only available through street hails.¹⁹ For each ride, the data set specifies the time-stamp (date and time up to the second) and location (longitude and latitude) of the pick-up/drop-off and itemized fares (fares, taxes, tolls, and tips) paid by the customers. The data set also includes identifiers for the taxi drivers that offered each ride. Therefore, we observe the status (full/empty) of the taxis at all times, the times at which taxis pick up or drop off customers, and the destinations of the customers. Note that we only observe the location of the taxis when they pick up or drop off a customer.

A preliminary look at the data indicates that weekdays and weekends as well as day and night shifts exhibit significant spatial and inter-temporal variations in the ride traffic. Figure 1.2 depicts the number of pick-ups per minute in the entire city on an average day. The number of pick-ups fluctuates between a minimum of 66 customers per minute (around 5:00

18. NYC taxi trip record data set is released by the NYC Taxi and Limousine Commission (TLC).

19. Prearranged services are offered by For-Hire-Vehicles (FHV).

{AM}) and a maximum of 507 customers per minute (around 7:30{PM}).²⁰ To focus on relatively busy hours of the day when the ride traffic is relatively stable, in the remainder of this chapter, we focus on the day shift on weekdays (i.e., 6_{AM} to 4_{PM}). Furthermore, we disregard trips with a duration more than three hours or a distance longer than a hundred miles. These instances are a combination of out-of-town trips and inaccurate data entries and account for less than 0.5% of the trips. Figure 1.3 depicts the average number of pick-ups during the day shift on a weekday in each month between January 2010 and August 2012. Prices did not change in this time-span. The dashed line in Figure 1.3, which depicts the regression of the average number of pick-ups on the month index k as in (1.1), highlights the trend in the average number pick-ups.

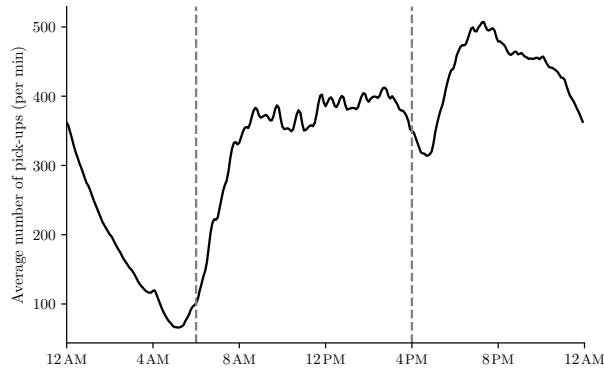


Figure 1.2: Average number of pick-ups per minute in NYC. The dashed vertical lines depict the beginning and the end of the day shift (6_{AM} and 4_{PM}, respectively).

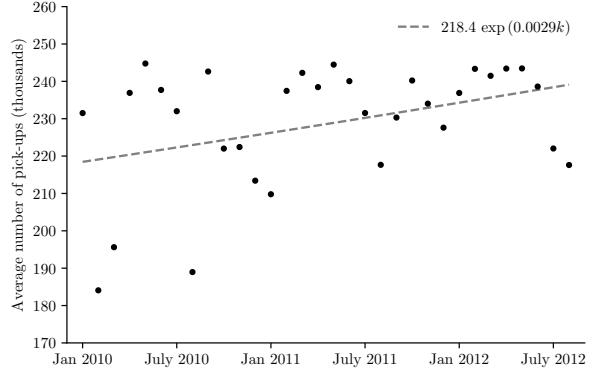


Figure 1.3: Average number of pick-ups in NYC during the day shift on a weekday of each month. The dashed line highlights the monthly trend.

There is significant spatial variation in the number of pick-ups across the city, with 93.8% of the rides (during the day shift) originating at Manhattan, 3.5% in the airports, 1.2% in Brooklyn, 1.4% in Queens, and less than 0.1% of the rides originating in Staten Island and Bronx. We focus our analysis on the rides taking place between the two airports (JFK and La Guardia) and parts of Manhattan, Brooklyn, and Queens that have sufficiently high traffic; see Appendix A.1. This area constitutes 99.4% of the pick-ups, 98.5% of the drop-offs, and

²⁰ A similar analysis can be carried out for the night shift.

97.9% of the pick-up, drop-off pairs in NYC.²¹ We divide this area into seventy five nodes and study trip statistics in each of them; see Figure A.3 for the nodes. This node definition provides sufficient separation between the major hubs of the city. Note that in the definition of the nodes, both size and density matter. This gives rise to a trade-off between size and density; for further details on the node definitions, see Appendix A.1. The average number of pick-ups and drop-offs in each of the seventy five nodes are depicted in Figure 1.4. Half of the pick-ups belong to the top ten nodes, with the top four nodes constituting a quarter of all pick-ups.

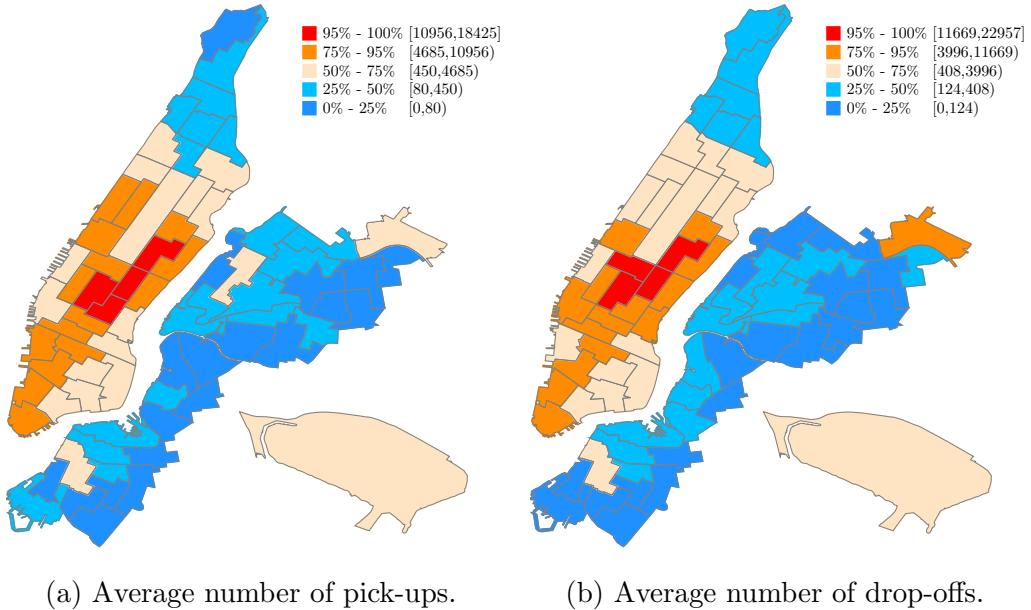


Figure 1.4: Average number of pick-ups and drop-offs during the day shift on weekdays in each node.

A trip on average is 2.7 miles and 13.2 minutes long. There is considerable spatial variation in the average trip duration and distance across the city; see Figures A.6-A.7 in Appendix A.3. Trips originating in Manhattan are the shortest with an average of less than three miles in most nodes and trips originating in the airports are the longest with an

²¹ Note that only 97.5% of the trips in the data set have pick-up coordinates that fall within the boundaries of NYC. Inaccurate pick-up coordinates (zero entry for longitude or latitude) account for 2.0% of the trips and out-of-town trips account for the remaining 0.5% of the trips.

average of nine miles in La Guardia and fifteen miles in JFK. Trips originating in Brooklyn and Queens fall in between with an average of three to five miles.

In summary, our analysis focuses on 217,051,494 taxi rides taken place between 6_{AM} and 4_{PM} on weekdays (excluding federal holidays) from January 2010 to December 2013, with pick-up and drop-off coordinates in the shaded area in Figure A.3.

1.5 Estimation

This section uses the nodes depicted in Figure A.3 to estimate the primitives of the model; see Appendix A.4 for a summary of the primitives. The procedure used to construct the nodes is outlined in Appendix A.1. The estimation is carried out in three steps. First, we directly estimate the problem primitives τ_{ij} , d_{ij} , S_{ij} , c , N_i , M , and m_i^1 . Second, given τ_{ij} , d_{ij} , S_{ij} , c , N_i , M , and m_i^1 , for each σ , we characterize the mean field equilibrium $(\lambda_i^t, m_i^t, m_{ij}^t, f_{ij}^t, q_{ij}^t, V(i, t))$ at each month. The data set provides the sequence of pick-ups and drop-offs of the taxi drivers. For each σ , we use the equilibrium relocation probabilities to calculate the likelihood of the observed sequence of pick-ups and drop-offs. Then, we use maximum likelihood estimation to estimate σ . Third, we use the estimated (potential) demand at the estimated σ to estimate the demand curve parameters, A_{ij}^t , α , and β . Appendix A.6 uses five-fold cross validation to examine the performance of our model in capturing the strategic relocations of the taxi drivers.

1.5.1 Offline Estimation of the Primitives τ_{ij} , d_{ij} , S_{ij} , c , N_i , M , and m_i^1

We use the trip duration and distances in the data set to estimate τ_{ij} and d_{ij} . We set $S_{ij} = 1$ if nodes i and j share a border or they are connected through a tunnel or bridge. Since JFK is geographically isolated and it does not share a border with any node, we add an arc between JFK and all nodes on the eastern border of Queens and Brooklyn. We use the fuel cost per mile and the depreciation cost per mile to estimate the cost of travel c , in each

month. The fuel cost can be estimated from the historical gas prices²² and the depreciation cost can be estimated from the salvage value of the taxis and the miles driven at the time of salvage.²³ Following Farber [87] and Frechette et al. [92], we define a shift of a taxi driver as a consecutive sequence of trips, where breaks between two trips cannot be longer than five hours. We use the shifts of the taxis drivers to calculate the number of active taxi drivers during the day. From 8_{AM} to 4_{PM}, the number of active taxi drivers shows little variation. Although the number of active taxi drivers between 6_{AM} and 8_{AM} is considerably lower than the rest of the day shift, for simplicity, we assume that the number of active taxis is fixed throughout the shift and let M equal to the average (across weekdays) of the maximum number of active taxi drivers during the day shift. Moreover, we use the distribution of the first pick-up of the taxis in the first fifteen minutes (average search time of the taxis until their next pick-up) of the day shift as the initial distribution of cars.

We estimate the number of regions in non-airport nodes as follows: Consider node i . Regions in node i are defined such that it takes a taxi exactly one period (five minutes) to explore each region. Therefore, it takes a taxi N_i periods ($5 \times N_i$ minutes) to search all the streets in node i . Thus, N_i can be estimated by dividing the average time it takes a taxi drivers to search all the streets of node i by the length of a period (five minutes). We use the average traffic speed in each zone and the shape file of NYC streets²⁴ to compute the average time it takes to explore all the streets in a node; see Table A.1 and Figure A.1 in Appendix A.1 for the average traffic speeds. At the airports, we set N_i equal to the number of taxi stands in that airport.²⁵

22. <https://www.nyserda.ny.gov/Researchers-and-Policymakers/Energy-Prices/Motor-Gasoline/Monthly-Average-Motor-Gasoline-Prices>.

23. In this calculation, we use the price of the dominant NYC Taxi model (Camry), \$24,000 including taxes, the average retired taxi salvage value of \$3,000, and the average mileage of the NYC taxis at the time of retirement, 350,000; See <https://nycitycab.com/Business/cabsforsalelist.aspx>.

24. <https://data.cityofnewyork.us/City-Government/NYC-Street-Centerline-CSCL-/exjm-f27b/data>.

25. Taxi operation in airports is accurately captured by (1.3b). In each stand, taxis and customers see each other and matching is perfect. However, taxis and customers in different stands cannot be matched.

1.5.2 Estimation of the Standard Deviation of Cost Shocks

Identification. We assume that taxi drivers in our dataset relocated rationally based on the model in Proposition 1. As proved in Proposition 1, taxi drivers relocate to those nodes for which the difference of their value function and the travel cost is higher with higher probabilities. However, the impact of the value function at the destination diminishes with σ . For small values of σ , taxi drivers almost always relocate to the destination with the highest value function minus travel cost. As σ increases, drivers relocate to destinations with lower value function minus travel cost with a higher probability, i.e., relocation probabilities become more uniform. Intuitively, the impact of σ on the relocation probabilities helps us identify σ .

Estimation. We estimate σ in two steps. First, for each σ , we characterize the equilibrium. We use the equilibrium relocation probabilities to calculate the likelihood of the sequence of pick-ups and drop-offs observed in the data set for each feasible σ . We then use maximum likelihood estimation to estimate σ .

Recall that we index the months with $k \in \{1, \dots, K\}$ and our data covers the months January 2010 through December 2013, i.e., $K = 48$. To emphasize the month dependence, we attach superscript k to various quantities of interest in the remainder of this section. We estimate λ_i^{tk} , the satisfied demand at node i in period t of month k , and $\pi_{ij}^{tk}(F^k, P^k)$, the distribution of the destinations of the customers at node i in period t of month k at fare structure (F^k, P^k) ²⁶, directly from the data. Given \bar{M} , M^k , N_i , m_i^{1k} , c^k , F_{ij}^k , P_{ij}^k , $\pi_{ij}^{tk}(F^k, P^k)$, and λ_i^{tk} , for each σ , we solve for the mean field equilibrium that satisfies the Equations (1.4)-(1.9) and $\lambda_i^{tk} \leq m_i^{tk}$.²⁷

26. Between January 2010 and August 2012, during the day shift, the fare of a ride between JFK and Manhattan was \$45 and the fixed portion of the fare and price per mile in the rest of the city were \$2.5 and \$2.0, respectively. On September 4th 2012, the fare structure of NYC taxi rides changed. The fare of a ride between JFK and Manhattan increased to \$52 and the price per mile increased to \$2.5. No price changes occurred between September 2012 and December 2013. We do not include Monday September 3rd 2012 in our analysis.

27. This approach and Definition 1 both characterize the mean field equilibrium. In Definition 1, λ_i^t satisfies

Next, we discuss how to use the equilibrium relocation probabilities and the sequence of pick-ups and drop-offs observed in the dataset to estimate σ . Let $m \in \{1, \dots, M^k\}$ and $d \in \{1, \dots, D^k\}$ index taxis and days, respectively, where D^k denotes the number of non-holiday weekdays in month k . Let $\mathcal{L}_{m,d,k}(\sigma)$ denote the likelihood of taxi m having the set of drop-off and subsequent pick-ups observed in the data on day d of month k that are not longer than two hours apart; see Appendix A.4 for the calculation of $\mathcal{L}_{m,d,k}(\sigma)$. In the dataset, a taxi driver on average searches for 2.5 periods after a drop-off until it picks up his next customer. Therefore, we omit drop-off and subsequent pick-up tuples that are longer than two hours apart from the likelihood function because they are likely due to breaks the drivers may be taking in their shifts. Indeed, the great majority of these occur around noon and likely correspond to lunch breaks. Such drop-off and pick-up tuples constitute less than 3% of all such tuples.

The likelihood of observing the set of drop-off and subsequent pick-ups (that are not longer than two hours apart) observed in the data set given σ is

$$L(\sigma) = \prod_{k=1}^K \prod_{m=1}^{M^k} \prod_{d=1}^{D^k} \mathcal{L}_{m,d,k}(\sigma),$$

where $\mathcal{L}_{m,d,k}(\sigma)$ is given in Equation (A.1) of Appendix A.4. Therefore, the maximum likelihood estimator of σ is the solution to

$$\text{maximize } \log(L(\sigma)) \text{ subject to (1.4)-(1.9)}, \quad (\text{P1})$$

given \bar{M} , M^k , N_i , m_i^{1k} , c^k , F^k , P^k , $\pi_{ij}^{tk}(F^k, P^k)$, and λ_i^{tk} (estimated directly from data) for all i, j, t, k . We solve Equations (1.4)-(1.9) for each month $k \in \{0, \dots, K\}$ and σ using the nonlinear optimization solver Knitro (Byrd et al. [60]). Then, we compute the likelihood of each σ and search \mathbb{R}^+ with a precision of 0.0001 for the solution to Problem (P1). We

equation (1.3) while in this characterization, λ_i^t is estimated from the data.

obtain the maximum likelihood estimate of $\sigma = 1.4048$; see Figure A.8 for a graph of log-likelihood values. To compute its standard error, we use the parametric bootstrap method (see e.g., Horowitz [127]). We generate 100 simulated datasets with the same size as our dataset using the estimated σ . We then estimate parameters of the simulated datasets and compute the standard error. This procedure results in the standard error of 0.001, a 95% confidence interval of (1.4015, 1.4056). We also conduct a Monte Carlo experiment to show that Problem (P1) can recover the true σ ; see Appendix A.5 for details.

1.5.3 Estimation of Demand Curve Parameters

In this section, we use the mean field equilibrium at the estimated σ to estimate the demand curve parameters A_{ij}^t , α , and β . Recall that demand model parameters are assumed to be the same in all non-holiday weekdays of the same month. Equation (1.3b) is monotone in Λ_i^t for all m_i^t . Therefore, given the satisfied demand λ_i^t and the mass of empty taxis m_i^t , there exists a unique estimate for (potential) demand $\Lambda_{ij}^{tk}(F^k, P^k)$ that satisfies (1.3b). Using this estimate, we set²⁸ $\Lambda_{ij}^{tk}(F^k, P^k) = \pi_{ij}^{tk}(F^k, P^k) \Lambda_i^{tk}(F^k, P^k)$ for all i, j, t, k . This gives the demand for rides from node i to node j in period t of a day in month k . Recall that $\pi_{ij}^{tk}(F^k, P^k)$ has been estimated directly from the data. Since there are D^k non-holiday weekdays in month k , the (realized) total demand for rides from node i to node j in period t across all non-holiday weekdays in month k is $Y_{ij}^{tk} = \bar{M} D^k \Lambda_{ij}^{tk}(F^k, P^k)$ customers.²⁹ Similarly, given the demand curve parameters A_{ij}^t , α , and β , by equation (1.1), the expected total demand for the same rides is $\bar{M} D^k A_{ij}^t [F_{ij}^k + P_{ij}^k d_{ij}]^\alpha \exp(\beta k)$ customers.

We use a negative binomial (NB-2) model³⁰ to estimate the primitives of the demand

28. Since customers who hailed a taxi are randomly chosen from all customers who needed a ride, the distribution of the destination of the served customers coincides with the distribution of the destinations of all customers.

29. The term $\Lambda_{ij}^{tk}(F^k, P^k)$ is the (normalized) mean field demand. By multiplying $\Lambda_{ij}^{tk}(F^k, P^k)$ by the average number of active taxis \bar{M} and the number of weekdays D^k in month k , we obtain the expected total number of customers who wanted a ride from node i to node j in period t on a weekday in month k .

30. A random variable Y is said to have a negative binomial distribution with parameters (μ, δ) if $Y \sim \text{Poisson}(\lambda u)$, where $u \sim \Gamma(\delta, 1/\delta)$. In other words, a negative binomial model can be thought of as a Poisson

model. To be more specific, we assume

$$Y_{ij}^{tk} \sim \text{Negative Binomial}\left(\bar{M}D^k A_{ij}^t [F_{ij}^k + P_{ij}^k d_{ij}]^\alpha \exp(\beta k), \delta\right) \quad (1.10)$$

and use maximum likelihood estimation to estimate its parameters. The parameter δ captures the over-dispersion of the estimated demands. As δ goes to zero, the negative binomial model approaches the Poisson model. For similar treatments and further details on negative binomial regression, see Hilbe [126], Hardin et al. [109, Chapter 13], and Lawless [153]. We obtain a price-elasticity estimate of $\alpha = -0.4735$ with standard error 0.0033, a monthly trend estimate of $\beta = 0.0024$ with standard error 1.7E-5 (equivalent to an annual trend of 2.92% with standard error 0.02%)³¹, and a dispersion parameter estimate of $\delta = 0.1082$ with standard error 1.5E-4. The estimated demand across the forty eight months do not show drastic changes. For brevity, we focus on September 2012 and provide an overview of the demand estimates in this month. The estimated potential and satisfied demand during the day in the entire city and Brooklyn are depicted in Figure 1.5. As discussed in Section 1.3, the efficiency of matching is positively correlated with the density of supply and demand. Therefore, in Manhattan, where the density of supply and demand is high (see Figure 1.6), the majority of customers are served. In Brooklyn, however, due to high friction as a result of low density of supply and demand, only a small fraction of customers are served. Increasing matching efficiency could add substantial value, particularly in locations with low density of supply/demand, by increasing the number of served customers. Furthermore, better matching could impact the power of spatial prices; see Section 1.6.2 for further details.

model with gamma heterogeneity, where the gamma noise has a mean of one. As a result, negative binomial is commonly referred to as the Poisson-Gamma mixture. In a negative binomial model, the ratio of variance to mean is $1 + \delta\mu$. As δ goes to zero, the gamma noise vanishes and the negative binomial model approaches the underlying Poisson model. We use the negative binomial model as opposed to the commonly used Poisson model because it can capture the over-dispersion (variance being greater than the mean) observed in the estimated demand.

31. Note that the trend of potential demand is not equal to the trend of satisfied demand; see Figure 1.3.

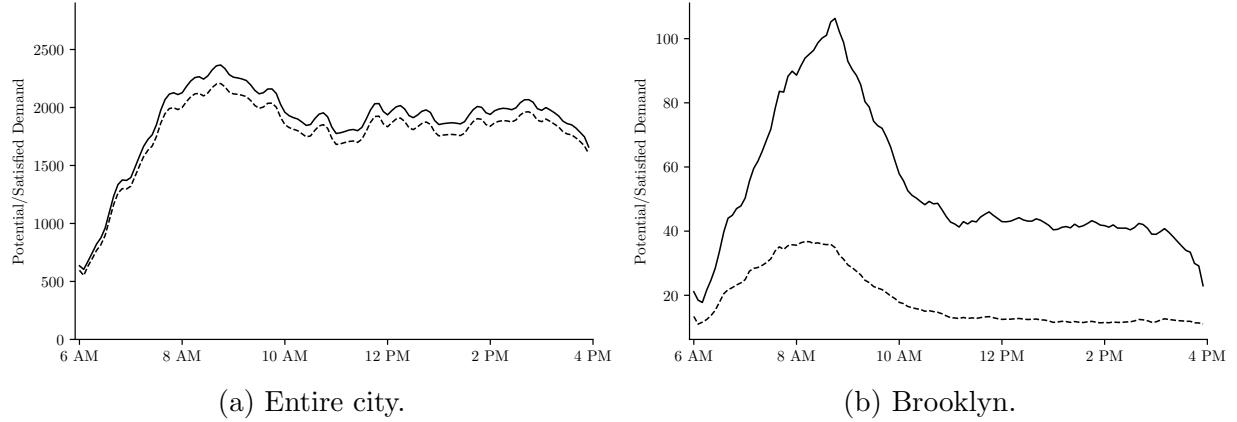
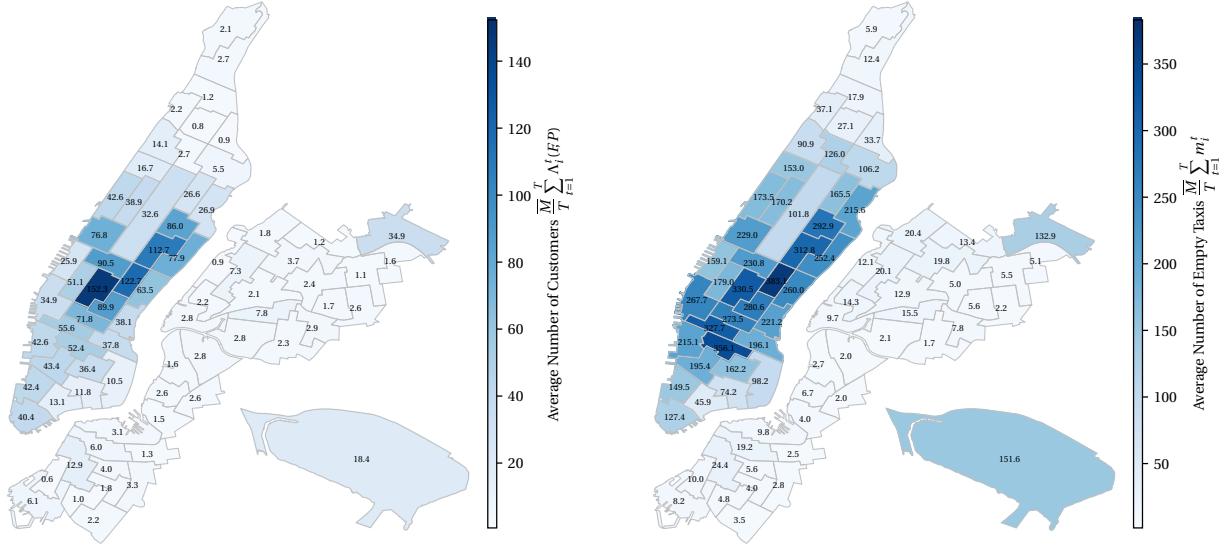


Figure 1.5: Estimated potential and satisfied demand (per period) during the day shift. The solid lines correspond to potential demand and the dashed lines correspond to satisfied demand.

1.6 Counterfactual Analysis

Using the model primitives estimated in Section 1.5, this section studies the impact of spatial pricing of taxi rides and removing the (local) search friction on the taxi market in NYC. Following Section 1.5, we focus our analysis on September 2012 and study spatial prices that maximize the consumer surplus; see Appendix A.8 for a similar analysis that focuses on maximizing drivers' profit. In what follows, we denote the fixed portion of the fare and the price per mile that were used in September 2012 by \bar{F}_{ij} and \bar{P}_{ij} , respectively, and refer to them as the base prices. We assume that $F_{ij} = \eta_{ij}\bar{F}_{ij}$ and $P_{ij} = \eta_{ij}\bar{P}_{ij}$, where η_{ij} denotes the ratio of the price/fare to the base price/fare for rides from node i to node j . We refer to η_{ij} as the price multiplier for rides from node i to node j . We assume this particular form because of computational simplicity. Moreover, for the case of origin-destination pricing, this form is without loss of generality.



(a) Average number of customers.

(b) Average number of empty taxis.

Figure 1.6: Average number of customers and empty taxis (in a five-minute period) during the day shift on weekdays in September 2012 (these are obtained through estimation; see Problem (P1) in Section 1.5).

1.6.1 Spatial Pricing

Since the demand curve introduced in Equation (1.1) has a constant price-elasticity and demand is (locally) inelastic,³² consumer surplus under this demand model is infinite. We use a truncation approach to tackle this issue (for a similar treatment, see Cohen et al. [76] and Buchholz [56]). We truncate the demand at some price multiplier $\hat{\eta} > 1$. This is equivalent to assuming that no customer is willing to pay more than $\hat{\eta}$ times the base fare for a ride. We refer to $\hat{\eta}$ as the maximum price multiplier. The next proposition derives the consumer surplus; see Appendix A.9 for its proof.

Proposition 2. *Consumer surplus resulting from the rides from node i to node j in period t under price multiplier η_{ij} is equal to*

$$CS_{ij}^t(\eta_{ij}) = \frac{\lambda_i^t}{\Lambda_i^t(F, P)} \Lambda_{ij}^t(\bar{F}, \bar{P}) [\bar{F}_{ij} + \bar{P}_{ij} d_{ij}] \frac{[\hat{\eta}^{(\alpha+1)} - \eta_{ij}^{(\alpha+1)}]}{1 + \alpha}. \quad (1.11)$$

³² The absolute value of the price-elasticity of rides around the base prices is less than one.

We start by studying pricing based on the origin of the rides. Let $\bar{V}(i, t)$ denote the value function of the taxi drivers under the base prices $(\bar{P}_{ij}, \bar{F}_{ij})$ and consider the following pricing formulation.

$$\underset{\eta_i}{\text{maximize}} \quad \sum_{i=1}^n \sum_{j=1}^n \sum_{t=1}^T \text{CS}_{ij}^t(\eta_i) \quad (\text{P2})$$

subject to (1.3) - (1.9) and

$$|\eta_i - 1| \leq \bar{\eta} \quad (\text{P2a})$$

$$(F_{ij}, P_{ij}) = \eta_i (\bar{P}_{ij}, \bar{P}_{ij}) \quad (\text{P2b})$$

$$\sum_{i=1}^n m_i^1 V(i, 0) \geq \sum_{i=1}^n m_i^1 \bar{V}(i, 0). \quad (\text{P2c})$$

We call the solution $\{\eta_i\}_{i \in \mathcal{V}}$ to Problem (P2) the optimal origin-only price multipliers (or equivalently the optimal origin-only prices) with a maximum price variation of $\bar{\eta}$. Note that in Problem (P2), price multipliers only depend on the origin of the ride. Constraints (1.3) - (1.9) ensure that equilibrium conditions are satisfied. Constraints (P2a)-(P2b) ensure that prices do not deviate from the base prices by more than a factor of $\bar{\eta}$. Constraint (P2c) ensures that the drivers' profit is not hurt by the new pricing scheme (i.e., the average value function of the drivers in the begining of the day shift under the new pricing scheme is greater than or equal to the average value function of the taxi drivers in the begining of the day shift under the base prices.).

The increase in consumer surplus for four plausible values of maximum price multiplier $\hat{\eta}$ is presented in Table 1.2. As shown in Table 1.2, using an origin-only pricing scheme with a maximum price variation of 20% results in an increase in consumer surplus of \$94,000 - \$325,000 on every day shift on weekdays. This is equivalent to an increase in consumer surplus of \$0.44 - \$1.53 per ride (4.9% - 17.0% of the average fare paid by the customers). Although the increase in consumer surplus is sensitive to maximum price multiplier $\hat{\eta}$ (as shown in Table 1.2), the pattern of prices is not sensitive to $\hat{\eta}$; see Figure A.11 in Section

A.7.3. Therefore, in the remainder of this section we use maximum price multiplier $\hat{\eta} = 5$, which is equivalent to assuming that for a ride of \$9 (average fare before tax and tips in September 2012), no customer is willing to pay more than \$45. This is a conservative choice for $\hat{\eta}$ considering its impact on consumer surplus (see Table 1.2).³³

Table 1.2: Consumer surplus under origin-only optimal spatial prices, with maximum price variation $\bar{\eta} = 0.2$.

Maximum price multiplier ($\hat{\eta}$)	5	10	20	50
Total CS under the base prices	\$4.8M	\$8.4M	\$13.7M	\$24.4M
Total increase in CS	\$94K	\$130K	\$196K	\$325K
Increase in CS per ride	\$0.44	\$0.61	\$0.90	\$1.53
Increase in CS in terms of average fare	4.9%	6.8%	10.0%	17.0%

Table 1.3 presents the increase in consumer surplus, number of served customers, and the miles traveled by customers for different maximum price variations, $\bar{\eta}$. We observe that larger values of $\bar{\eta}$ results in larger increases in consumer surplus, as one would expect. However, the returns to the change in $\bar{\eta}$ are diminishing. The reason for the diminishing returns is illustrated in Figure 1.7. As the maximum price variation $\bar{\eta}$ increases, fewer nodes require a price variation higher than $\bar{\eta}$. Table 1.3 indicates that the optimal origin-only pricing scheme (for maximizing consumer surplus) results in a 2.6% increase in the number of served customers and a 3.9% increase in the miles traveled by customers for a maximum price variation of $\bar{\eta} = 0.5$.³⁴

Given the decomposition of the consumer surplus provided in Proposition 2, we expect the average fraction of customers served, demand, and the fare to impact the optimal price multiplier of a node. In what follows, we study the impact of these parameters on the optimal origin-only price multipliers.

The price multiplier of rides from node i to node j impacts consumer surplus only through

33. In a similar setting, Cohen et al. [76, Section 4] makes the conservative assumption that no customer is willing to pay more than 4.9 times the base prices for an Uber ride.

34. By using the total number of served customers or miles traveled by customers as objectives in Problem (P2), we can find spatial prices that result in larger increases in these metrics. However, such spatial prices result in a lower consumer surplus.

Table 1.3: Impact of maximum price variation on various performance metrics (maximum price multiplier $\hat{\eta} = 5$).

Maximum price variation ($\bar{\eta}$)	10%	20%	30%	40%	50%
Consumer surplus					
Total increase	\$62K	\$94K	\$113K	\$126K	\$135K
Per ride	\$0.29	\$0.44	\$0.53	\$0.59	\$0.63
In terms of average fare	3.3%	4.9%	5.9%	6.6%	7.0%
Number of served customers	1.1%	1.7%	2.1%	2.4%	2.6%
Miles traveled by customers	1.8%	2.7%	3.2%	3.6%	3.9%

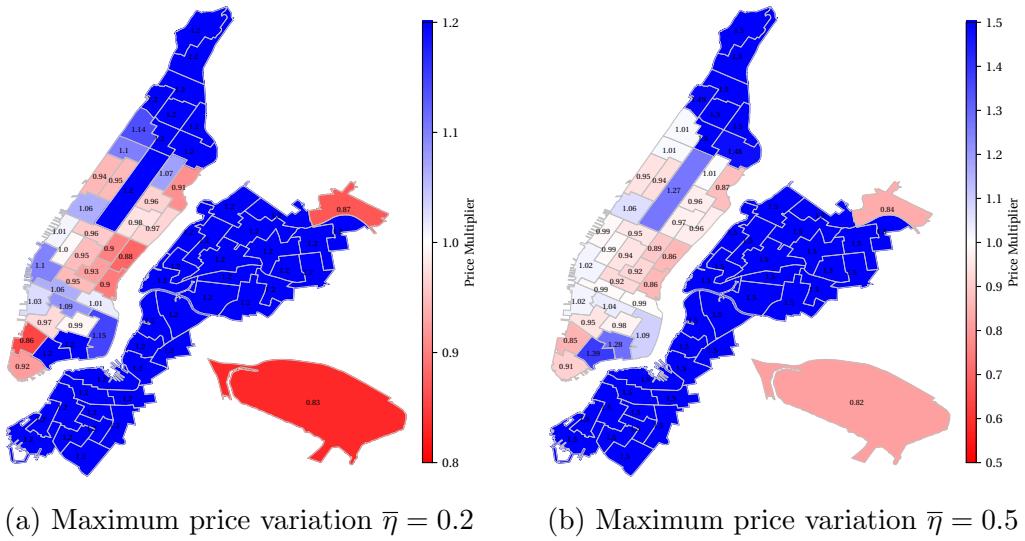


Figure 1.7: Optimal origin-only price multipliers for different maximum price variations $\bar{\eta}$ (maximum price multiplier $\hat{\eta} = 5$).

the fraction of customers served, $\lambda_i^t / \Lambda_i^t(F, P)$, and the term $(\hat{\eta}^{(\alpha+1)} - \eta^{(\alpha+1)}) / (1 + \alpha)$. As higher prices attract more drivers, they increase the fraction of customers served. The term $(\hat{\eta}^{(\alpha+1)} - \eta^{(\alpha+1)}) / (1 + \alpha)$, however, decreases in price. Depending on which term dominates, the consumer surplus of a node could be increasing or decreasing in its price multiplier.

In well-served nodes, a high fraction of customers are served and the impact of price multiplier on the fraction of customers served is minimal. Therefore, in such nodes, the last term dominates and the consumer surplus is decreasing in the price multiplier. In contrast, in under-served nodes, a lower fraction of customers are served and the impact of price multiplier on the fraction of customers served is substantial. In under-served nodes, the first

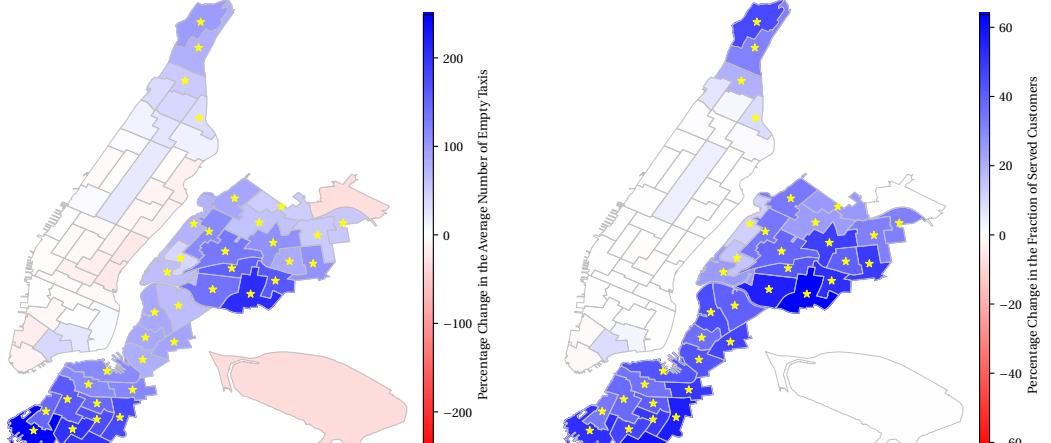
term dominates and the consumer surplus is locally increasing in the price multiplier. In such nodes, both customers and drivers benefit from higher prices. Therefore, the optimal price multiplier of these nodes is expected to be greater than one.

Nodes in which less than 80% of the customers are served (under the base prices) are marked with a star in Figure 1.8. We refer to these as under-served nodes. These nodes belong to upper Manhattan, Brooklyn and Queens, in which the supply of taxis is low; see Figure 1.6 for the map of the average supply and demand under the base prices. These nodes have a price multiplier greater than one. In response to a increase in prices in these nodes, more taxis relocate to them, which increases the supply of rides at these nodes; see Figure 1.8a. As a result of this increase in supply, a higher fraction of customers are served; see Figure 1.8b for the change in the fraction of customers served. This indicates that in these nodes the first term on the right hand-side of Equation (1.11) is dominant.

The fact that prices increase in less affluent neighborhoods in the city (upper Manhattan, Brooklyn, and Queens) highlights the necessity to use other mechanisms in addition to spatial prices (such as facilitating matching between customers and taxis through mobile applications) or subsidizing these neighborhoods. This issue is further explored in Section 1.6.2.

We refer to nodes in which at least 80% of the customers are served under the base prices as well-served nodes. For such nodes, the terms $\Lambda_{ij}^t(\bar{F}, \bar{P})$ and $\bar{F}_{ij} + \bar{P}_{ij}d_{ij}$ are the critical drivers of consumer surplus. To see this, consider the potential revenue that can be collected at a node, $\sum_{t=1}^T \sum_{j=1}^n \Lambda_{ij}^t(\bar{F}, \bar{P}) [\bar{F}_{ij} + \bar{P}_{ij}d_{ij}]$. Note that both demand and fare impact the potential revenue of a node. In the majority of the well-served nodes with high potential revenue prices decrease while in the majority of well-served nodes with low potential revenue prices increase; see Figure 1.9.

Next, we study origin-destination pricing; see Appendix A.7.1 for its mathematical formulation. Similar to origin-only pricing, we use maximum price variation $\bar{\eta} = 0.5$ and maximum price multiplier $\hat{\eta} = 5$; see Appendix A.7.2 for further details. Origin-destination



(a) Average number of empty taxis.

(b) Fraction of served customers.

Figure 1.8: Percentage change in the average number of empty taxis and the fraction of served customers under the optimal origin-only prices compared to the base prices (maximum price variation $\bar{\eta} = 0.5$ and maximum price multiplier $\hat{\eta} = 5$).

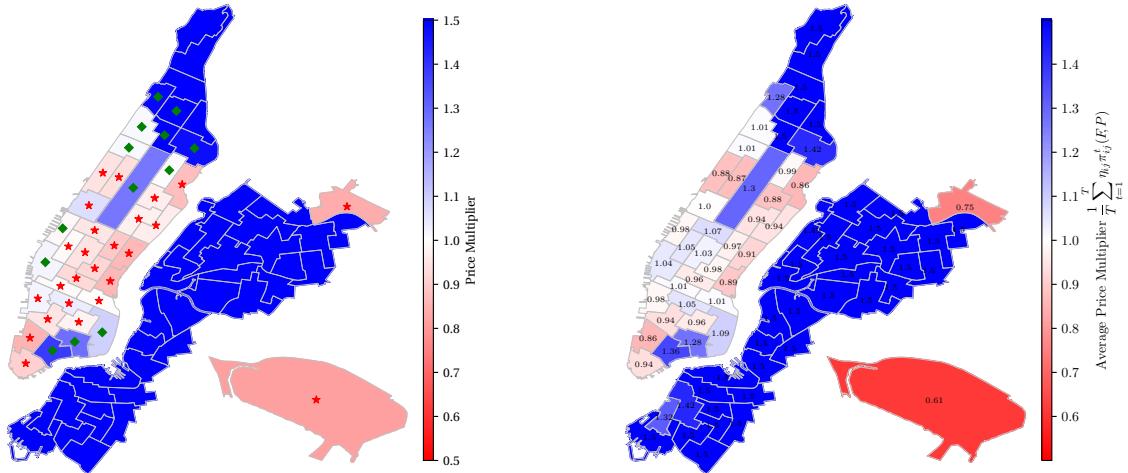
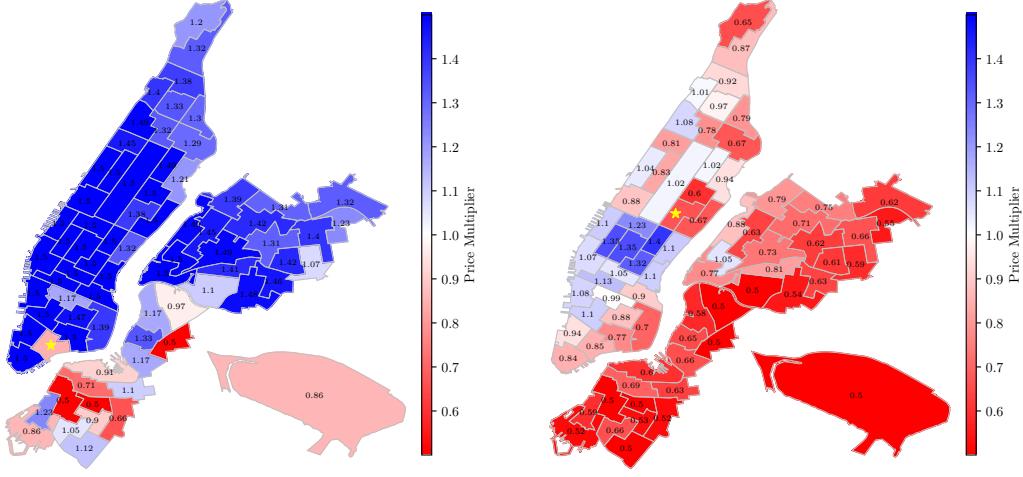


Figure 1.9: Well-served nodes with high (star) and low (diamond) potential revenue. Prices decrease in 88% of the nodes with high potential revenue and increase in 92% of the nodes with low potential revenue.

prices result in a \$0.79 increase in consumer surplus per ride, which is 24.4% higher than the increase in consumer surplus from origin-only pricing. They also results in a 3.2% increase in the number of served customers and an 7.4% increase in the miles traveled by customers.

The average optimal origin-destination price multipliers for the maximum price variation of 50% are depicted in Figure 1.10. Comparing Figures 1.7b and 1.10, we observe that the pattern of the average optimal origin-destination prices looks similar to the pattern of optimal origin-only prices. In the (majority of) under-served nodes, the same price is used for all destinations. In these nodes, the benefit of attracting more drivers by using higher average prices outweighs the benefit of price discrimination based on destinations. In the well-served nodes, however, price discrimination based on destinations is valuable. Figure 1.11 provides two examples of how destinations could impact the prices. Let us start with the node marked by a star in Figure 1.11a. Under origin-only pricing (see Figure 1.7b), the price of all rides from this node increase by 39%. However, under origin-destination pricing, the price of rides from this node to (parts of) Brooklyn are reduced. This reduction in prices increases the demand for rides to Brooklyn, which in turn, helps further increase the average number of the empty taxis in this region. That is, this increases the supply of taxis in Brooklyn for future rides. Since only 15% of the customers are headed to these destinations, this pricing pattern maintains (almost) the same average price as the origin-only pricing scheme. The pattern of reducing the price of rides headed to (some) under-served neighborhoods while increasing the price of rides headed to (some) well-served nodes is observed frequently. Figure 1.11b depicts another example of this pattern. Under origin-only pricing (see Figure 1.7b), the price of all rides from this node decrease by 3% (almost no change in prices). Under origin-destination pricing, the price of rides to midtown and parts of downtown increase while the price of all other rides decrease. This decreases the future supply in midtown and parts of downtown and increases the future supply in the rest of the city, particularly in Brooklyn and Queens. The price multiplier of the rides originating at this node range between 0.5 and 1.35 (a substantial change in prices). These examples exhibit how using destinations alongside origins in the pricing scheme allows us to maintain an average price similar to the origin-only pricing scheme while adjusting supply and demand in a more refined manner.



(a) A node in downtown Manhattan.

(b) A node in upper east side.

Figure 1.11: Optimal price multipliers under origin-destination pricing for the rides originating at the node denoted by a star (maximum price variation $\bar{\eta} = 0.5$ and maximum price multiplier $\hat{\eta} = 5$).

1.6.2 Removing Local Search Friction

Mobile application (such as Arro and Curb) can remove the local search friction by matching customers and taxi drivers in each neighborhood. The impact of such applications can be modeled by replacing the matching model (1.3b) with that in (1.3a).

Table 1.4 shows removing the local search friction alone results in a \$268,000 increase in consumer surplus in every day shift, a 4.3% increase in the number of served customers, and an 6.2% increase in customer miles. It also compares the improvement with those from spatial pricing. Removing the (local) search friction is on par with spatial pricing in terms of the number of served customers and customer miles. However, its impact on consumer surplus is considerably higher. Furthermore, although spatial prices introduced in Section 1.6.1 have no impact on drivers' profit by design, since removing the local search friction increases the total number of served customers without changing the prices, drivers' profit increase by \$48,000 in every day shift when search friction is removed.

For brevity, we focus on origin-only pricing in the remainder of this section. Figure 1.12 depicts the change in consumer surplus from origin-only spatial pricing and removing the

(local) search friction alone. The majority of the change in consumer surplus from spatial pricing is due to nodes with high demand and average fare, such as midtown Manhattan and the airports. Contrary to spatial pricing, when the local search friction is removed, well-served nodes (the majority of the nodes in Manhattan and the airports) do not benefit much, whereas under-served nodes enjoy the highest increase in consumer surplus. Under-served nodes benefit more from removing the (local) search friction than from spatial prices. In such nodes, matching is the key. In contrast, well-served nodes benefit more from spatial pricing than from removing the local search friction. In such nodes, spatial pricing is the key. This suggests that the impact of spatial prices and removing local search friction are complementary to each other. Therefore, we study a hybrid mechanism.

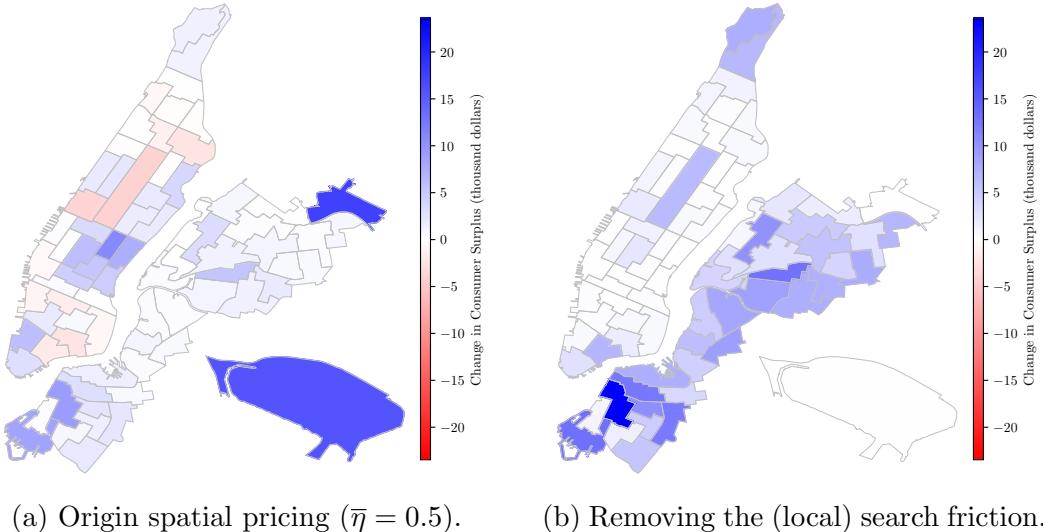


Figure 1.12: Change in consumer surplus from spatial pricing and removing the search friction (maximum price multiplier $\hat{\eta} = 5$).

To be specific, the hybrid mechanism removes (local) search friction in under-served nodes and uses spatial prices in well-served nodes; see Appendix A.7.1 for its mathematical formulation. Table 1.4 provides a comparison of the impact of spatial prices, removing the (local) search friction, and the hybrid mechanism on consumer surplus and drivers' profits. The hybrid mechanism with maximum price variation $\bar{\eta} = 0.2$ captures 99.4% of the increase in consumer surplus from a hybrid mechanism with $\bar{\eta} = 0.5$. This mechanism needs only a

little price variation, because the optimal price multipliers of the well-served nodes are close to one (see Figure 1.7).

Table 1.4: Comparison between spatial prices, removing local search friction, and the hybrid mechanism ($\bar{\eta} = 0.5$).

	Origin-Only Pricing	Origin-Destination Pricing	Removing Local Search Friction	Hybrid Mechanism	Proposed Mechanism	Citywide Origin-Only Pricing & Friction Removal
Consumer Surplus						
Total increase	\$135K	\$168K	\$268K	\$322K	\$417K	\$433K
well-served nodes	\$79K	\$109K	\$33K	\$82K	\$182K	\$186
under-served nodes	\$56K	\$59K	\$235K	\$240K	\$235K	\$247K
Per ride	\$0.63	\$0.79	\$1.26	\$1.52	\$1.96	\$2.04
In terms of average fare	7.0%	8.5%	13.9%	16.7%	21.5%	22.4%
Number of served customers	2.6%	3.2%	4.3%	5.7%	8.7%	8.9%
Miles traveled by customers	3.9%	7.4%	6.2%	8.1%	11.7%	12.0%
Drivers' Profit						
Total increase	\$0	\$0	\$48K	\$0	\$0	\$0
Per ride	\$0	\$0	\$0.22	\$0	\$0	\$0
In terms of average fare	0%	0%	2.5%	0%	0%	0%

Next, we study citywide spatial pricing and friction removal. The optimal price multipliers of this mechanism are the solution to Problem (P2), where (1.3b) is replaced with (1.3a). As shown in Table 1.4, using citywide spatial pricing and friction removal generates \$433,000 of consumer surplus. This is equivalent to a 34.4% improvement over the hybrid mechanism. The majority of this improvement is in well-served nodes. The optimal price multipliers under this mechanism are depicted in Figure 1.13. Although the pattern of optimal prices in well-served nodes (under the base prices) is similar to the pattern of optimal prices in presence of friction (see Figure 1.7a), once (local) search friction is removed in well-served nodes, it is optimal to use more aggressive price variations in these nodes (i.e., we observe a larger deviation from the base-prices, both in nodes that have a price multiplier greater than one and those that have a price multiplier less than one). This implies that in the absence of friction, spatial prices are more powerful.

Policy makers may prefer to avoid price discrimination in less affluent neighborhoods of the city. As such, we propose a mechanism in which friction is removed in the entire city

while spatial pricing is used only in well-served neighborhoods; see Appendix A.7.1 for its mathematical formulation. The proposed mechanism increases consumer surplus by \$1.96 per ride (96.3% of the consumer surplus generated by citywide spatial pricing and friction removal), serves 8.7% more customers, and increases customer miles by 11.7%.

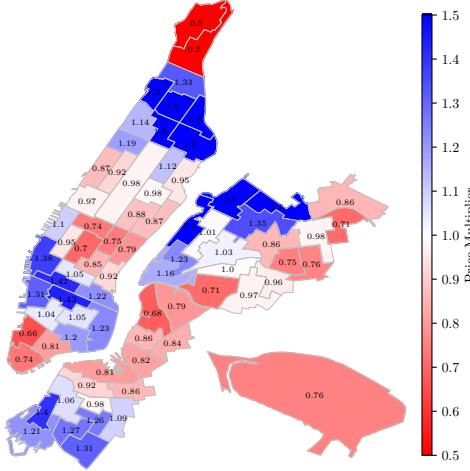


Figure 1.13: Optimal price multipliers under citywide origin-only pricing and friction removal (maximum price variation $\bar{\eta} = 0.5$ and maximum price multiplier $\hat{\eta} = 5$).

1.7 Concluding Remarks

We study the impact of spatial pricing on the taxi market in New York City. We use a mean field model, in which taxi drivers strategically search for customers in different neighborhoods across the city taking into account the spatial and temporal distribution of the supply and demand as well as the prices across the city. Our analysis reveals that spatial prices that only use origin information can increase consumer surplus by \$0.63 per ride, 7.0% of the average fare, and serve 2.6% more customers without hurting the drivers' profit. Similarly, spatial prices that utilize both origin and destination information can increase consumer surplus by \$0.79 per ride, 8.5% of the average fare, and serve 3.2% more customers. The optimal spatial prices increase in the under-served areas (e.g., Brooklyn, Queens, and upper Manhattan) whereas they decrease in some well-served areas (e.g., midtown), encouraging

the taxis to relocate accordingly, and hence, redistributing supply so as to increase consumer welfare.

We find that removing the local spatial search friction alone can increase consumer surplus by \$1.26 per ride, 13.9% of the average fare, and serve 4.3% more customers while simultaneously increasing drivers' profit by \$0.22 per ride. Removing the local search friction primarily impacts the under-served neighborhoods, whereas spatial prices primarily impact the well-served neighborhoods. This underscores the value of a hybrid mechanism. We propose a mechanism in which (local) search is eliminated in all neighborhoods while spatial pricing is only used in well-served neighborhoods. This mechanism increases consumer surplus by 21.5% of the average fare and serves 8.7% more customers, while avoiding price discrimination in less affluent neighborhoods of the city. The proposed mechanism achieves 96.3% of the benefits of a citywide spatial pricing and friction removal mechanism.

Our analysis has some limitations. First, a social planner might be most concerned with the impact of spatial pricing on the customer that is hurt by it the most. We limit the worst case impact of spatial pricing by introducing a maximum price constraint. Moreover, although we show that both customers and taxi drivers benefit from higher prices in the under-served areas in the presence of search friction, we do not study the implications of this phenomena fully. This phenomena leads to interesting work on subsidies and voucher mechanisms that is beyond the scope of our work. However, our study of the hybrid mechanisms is motivated by this concern. Lastly, this dissertation uses the NYC taxi dataset from an era, 2010-2013, in which taxi substitutes were limited.³⁵ Future work in this area is needed to study the impact of spatial pricing on customers and taxi drivers in the presence of competition.

35. Ride-sharing platforms and street hail liveries (green taxis) had less than 5% market share in May 2014, five months after the period this dissertation studies ends; see TLC [2] and Schneider [186].

CHAPTER 2

DYNAMIC DISPATCH AND CENTRALIZED RELOCATION OF CARS IN RIDE-HAILING PLATFORMS

2.1 Introduction

This chapter studies a ride-hailing platform that offers on-demand transportation services. A major operational concern of the ride-hailing platform is to match cars and customers. This is challenging because customers arrive randomly over time at various areas of the city, the distribution of cars across the city evolves randomly over time, the geography of the city imposes constraints on the matching of cars and customers, and the value generated by a match can depend on features of both the car and the customer involved. Motivated by this operational concern, this chapter considers a ride-hailing platform in steady state that seeks to maximize its profit by making dynamic dispatch and centralized relocation decisions. To be more specific, we study a platform that dynamically decides the following: First, which car, if any, to dispatch to each arriving customer. Second, which car, if any, to relocate from its current area to another area. The joint consideration of dispatch and relocation decisions makes the ride-hailing platform more responsive to the variability in the demand process and better equipped to tackle any inherent asymmetries in the system. For example, merely considering dispatch decisions can result in an undesirably high demand dropping probability independent of the number of cars due to possible asymmetries in demand; see e.g., Figure 2.12 in Section 2.9. Throughout the chapter, we apply our model to study the case of New York City (NYC) using the NYC yellow taxi dataset; see Chapter 1 for a description and empirical analysis of the NYC taxi market. Section 2.9 describes our results for the NYC application.

In determining the joint dynamic dispatch and relocation policy, the platform takes into account the distribution of cars, the demand patterns, the uncertainty of the demand, the set of available dispatch and relocation activities, the values of different dispatch activities, and

the costs of different relocation activities. We divide the city into disjoint areas and classify the cars into car classes based on their current area as well as other relevant information such as their size and luxury status. Similarly, we classify the customers into classes based on their pickup area as well as other relevant information such as the list of preferred car sizes and luxury status. In the NYC application, we divide the city into nine areas as depicted in Figure 2.1 and classify the cars and customers based on their location, i.e., we consider nine car classes and nine customer classes. Also, we assume the platform can dispatch cars to pick up customers who are less than fifteen minutes away. This corresponds to nineteen dispatch activities: nine local dispatch activities in which cars are dispatched to pick up the next customer in the same area, and ten non-local dispatch activities in which cars are dispatched to pick up the next customer in another area. The non-local dispatch activities are depicted with arrows in Figure 2.1; an arrow from area i to area k depicts the activity of dispatching a car from area i to pick up the next customer in area k . For concreteness, we assume the platform in the NYC application keeps a quarter of the fare paid by the customers¹ and incurs a dispatch cost of twenty cents per mile.² We assume the platform can also relocate cars from any area to any other area for a fee. This corresponds to seventy two relocation activities. For visual clarity, the relocation activities are not depicted in Figure 2.1. For concreteness, we assume the platform in the NYC application incurs a one dollar and twenty five cents per mile relocation cost for relocating each car.³

We assume cars move between areas instantaneously; see Banerjee et al. [40], Kanoria and Qian [136], and Besbes et al. [50] for similar assumptions. This modeling assumption allows us to obtain a tractable formulation by avoiding a control problem whose dimension grows quadratically with the number of areas in the city; see Alwan and Ata [15] for further details. The solution to our formulation provides structural insights into the dynamic control

1. This assumption is motivated by Uber’s practice of collecting 25% of the fares paid by customers as a service fee; for further details, see <https://www.uber.com/gh/en/drive/basics/tracking-your-earnings>.

2. The dispatch cost is solely used in the NYC application and the simulation study in Section 2.9.

3. This relocation cost is one half of the price per mile for NYC yellow taxis.

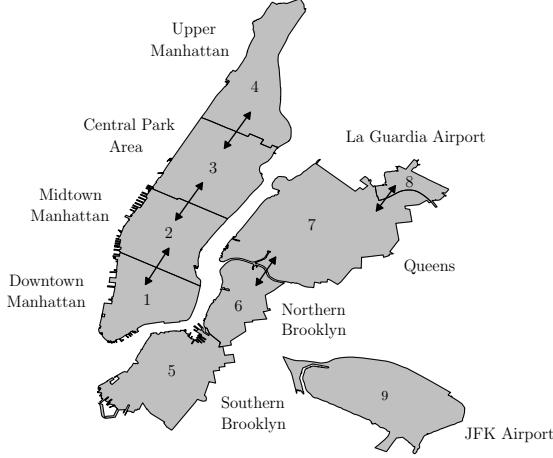


Figure 2.1: A map of NYC and the nine areas considered. The arrows represent the non-local dispatch activities. The local dispatch activities are not depicted for visual clarity. An arrow from area i to area k represents the activity of dispatching a car in area i to pick up the next customer in area k .

of ride-hailing platforms and its prescribed control policy performs well in the presence of non-zero travel times as we illustrate through a simulation study in Section 2.9.

We model the ride-hailing platform as a closed stochastic processing network as described in Harrison [114]. A stochastic processing network has three basic elements: jobs, servers, and activities. In our setting, jobs correspond to cars (job classes correspond to car classes), servers correspond to customer classes, and activities correspond to either the dispatch or the centralized relocation of cars. Each job class has a dedicated buffer, in which the jobs of this class are stored. In the context of the stochastic processing network, a dispatch activity corresponds to a server processing a job, followed by the job probabilistically transitioning into another job class. Dropping demand corresponds to idling a server. Similarly, a relocation activity corresponds to moving a job from one buffer to another. No server is involved in this activity.

The closed stochastic processing network model of the NYC application is depicted in Figure 2.2. The customers arriving at each of the nine areas in the ride-hailing system are modeled as the nine servers. There is a buffer corresponding to each area, and the cars waiting in an area are modeled as the jobs waiting in the corresponding buffer. The jobs

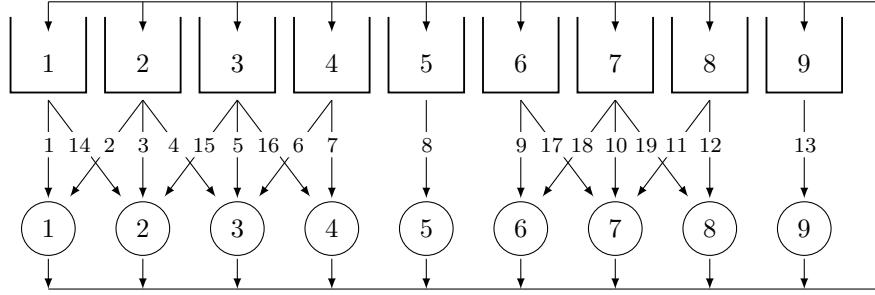


Figure 2.2: The closed stochastic processing network model of the NYC application and the numbering of the dispatch activities.

waiting in each buffer constitute a job class. There are nine job classes in total. The activity of dispatching a car waiting in an area to pick up the next customer arriving in another area is modeled as the activity of the server corresponding to the customer's area serving the buffer corresponding to the car's area. The nineteen dispatch activities are depicted as arrows in Figure 2.2. An arrow from buffer i to server k depicts the dispatch activity of serving buffer i with server k . We number the dispatch activities as shown in Figure 2.2. The activities are numbered in a manner consistent with the numbering of basic and non-basic activities introduced in Section 2.4. The activity of relocating a car from one area to another is modeled as the activity of moving a job from one buffer to the other. For visual clarity, the relocation activities are not depicted in Figure 2.2; see Table B.1 in Appendix B.1.2 for a list of the relocation activities.

In the absence of randomness, a linear program called the static planning problem can be used to solve for the optimal control policy; see Figure 2.3 for the optimal solution to the static planning problem of the NYC application. This optimal solution partitions the set of activities into two: basic (or efficient) activities and non-basic (or inefficient) activities. The basic activities are used under the optimal solution, whereas the non-basic activities are not. A key assumption of our work is that the optimal solution to the static planning problem perfectly matches the supply and demand in the absence of randomness. In other words, the value generated by dispatch activities, the relocation costs, and the demand flow are such that it is worth serving the entire demand in the absence of randomness. This assumption holds for

the NYC application and it facilitates the formulation of an approximating Brownian control problem that is far more tractable than the stochastic processing network model. Under uncertainty, however, a stochastic mismatch between supply and demand arises. Because the solution to the static planning problem cannot address this stochastic mismatch, it is not optimal in that case. The approximating Brownian control problem seeks to find the deviations from the solution to the static planning problem that provide the best match between the supply and demand under uncertainty.

The optimal solution to the static planning problem helps the platform pool the demand of different customer classes using the basic dispatch and relocation activities. In the context of the stochastic processing network, this corresponds to pooling servers or resource pooling. The NYC example has three such pools; see Figure 2.4. The four areas of Manhattan (downtown Manhattan, midtown Manhattan, central park area, and upper Manhattan) and southern Brooklyn constitute the first pool, northern Brooklyn is the second pool, and Queens, La Guardia airport, and JFK airport constitute the third pool. The pooling leads to a workload formulation that is equivalent to the Brownian control problem, but has a significantly lower dimension. Due to the dimension reduction, it suffices for purposes of optimal control to keep track of only the total number of cars that can be served by each resource pool using basic activities. The equivalent workload formulation is a singular stochastic control problem. Although the dimension reduction significantly simplifies the problem, multidimensional singular control problems are not analytically tractable in general. However, we are able to solve this problem numerically for the NYC example. To do so, we extend the computational method of Kumar and Muthuraman [145] to control problems with state space constraints and arbitrary number of oblique pushing directions. This could be of interest in its own right.

Interpreting the solution to the equivalent workload formulation in the context of the stochastic processing network model of the ride-hailing problem leads to a simple, intuitively appealing, and easily implementable policy: When there are sufficiently many cars to meet

the demand in each pool using basic activities, that is, when the workload is balanced, the policy uses the optimal solution to the static planning problem. That is, it uses the basic dispatch and relocation activities at rates prescribed by the static planning problem. When there are not sufficiently many cars to meet the demand in a pool using basic activities, that is, when the workload is sufficiently imbalanced, the platform intervenes to balance the workload. The interventions are in the form of dropping demand and using a non-basic dispatch or relocation activity. For the NYC application, they are dropping demand and using non-basic dispatch activities; see Figure 2.8c and Table 2.1 for a list of the interventions and when they are undertaken. For example, when no cars are in pool 2 (northern Brooklyn), depending on the distribution of cars in the other pools, the platform makes one of the following interventions: When a higher fraction of cars are in pool 3 (Queens, La Guardia, and JFK), cars from Queens are dispatched to pick up the customers in pool 2. However, when a lower fraction of cars are in pool 3, customers in pool 2 are dropped.

This work builds on and extends the existing literature and makes the following contributions. First, it studies the joint dynamic dispatch and centralized relocation of cars in ride-hailing platforms, thereby extending the literature that has focused on either the dispatch or the relocation control of ride-hailing platforms. Second, it contributes to the modeling and analysis of the ride-hailing platforms by modeling the platform as a closed stochastic processing network, which we study in the heavy traffic asymptotic regime and arrive at a far more tractable formulation. Third, it derives a simple and intuitive policy that is verifiably good for the NYC application. To be specific, we show through a simulation study for NYC that when the supply of cars is nearly adequate, our proposed policy outperforms the closest-driver policy, the MaxWeight policy, and the mirror backpressure policy. It also provides the following structural insights from the derived policy: The platform can pool the demand of different customer classes and the information on the total number of cars in each pool is sufficient to dynamically control the platform. Moreover, only when the distribution of cars is (sufficiently) imbalanced, e.g., there are not sufficiently

many cars to meet the demand in a pool using basic activities, the platform deviates from its static control policy to drop demand or to engage in a non-basic dispatch or relocation activity. Furthermore, the deviations from the static control policy cannot be determined solely based on the local information; that is, global information, such as the demand distribution and the distribution of cars across different pools, is required. Fourth, it develops a computationally efficient and numerically robust algorithm to solve multi-dimensional singular stochastic control problems with state space constraints, which have an arbitrary number of (oblique) pushing control directions at each face. This algorithm can be of interest beyond ride-hailing platforms. For example, it could be useful to study other stochastic network models or portfolio optimization problems that often give rise to multidimensional singular control problems.

The rest of the chapter is organized as follows. Section 2.2 reviews the literature. Section 2.3 introduces the control problem for the ride-hailing platform. Section 2.4 introduces the static planning problem and the approximating Brownian control problem. Section 2.5 introduces the reduced Brownian control problem. Section 2.6 introduces the equivalent workload formulation. Section 2.7 develops a numerical solution method for the equivalent workload formulation. Section 2.8 interprets the solution to the equivalent workload formulation in the context of the control problem laid out in Section 2.3. Section 2.9 provides a simulation study to demonstrate the effectiveness of the proposed policy. Section 2.10 concludes. Appendices B.1-B.3 provide auxiliary technical material.

2.2 Literature Review

This chapter is related to three streams of literature. The first stream studies the control of ride-hailing platforms. The second stream studies the control of queueing systems. The third stream studies singular stochastic control problems. The control of ride-hailing systems has received a great deal of attention. The majority of this literature has focused on pricing, dispatch, and relocation as the primary levers of control. Banerjee et al. [41] is one of the early

papers that study the control of ride-hailing platforms. The authors focus on the dynamic pricing of a single-location ride-hailing platform. They show that although the optimal static pricing policy is asymptotically optimal in the fluid-scale, dynamic pricing policies are more robust to fluctuations in the system parameters. Banerjee et al. [39] designs an approximately optimal pricing policy and shows that the approximation ratio of the pricing policy improves as the number of cars in each region grows. Despite the negative publicity of surge pricing, Cachon et al. [61] demonstrates that surge pricing in ride-hailing platforms can make both the platform and the customers better off. Besbes et al. [50] uses a stylized model on a line, where drivers make myopic decisions based on the next ride and can relocate instantaneously. The authors show the pricing problem can be spatially decomposed based on the attraction regions. The platform can use prices to create regions in which driver congestion is artificially high in order to motivate drivers to relocate to more profitable regions. Bimpikis et al. [52] focuses on the spatial pricing of rides in ride-hailing platforms. The authors show that if the demand pattern is not balanced, spatial pricing is beneficial and optimal prices can be written in terms of the optimal dual variables corresponding to the flow balance equations. Under the assumption that prices and compensations can be decoupled, the authors solve for the optimal prices by focusing on the mass balance equations only.

Selcuk and Gokpinar [188] studies the optimal design of spatial prices and commissions in ride-hailing platforms. The authors show that a flexible commission policy is a more effective tool than a flexible pricing policy. In particular, unless the commissions are flexible, the platform runs into bottlenecks at locations associated with short trip lengths. Garg and Nazerzadeh [94] studies surge pricing using a dynamic stochastic model. The authors show that multiplicative surge is not incentive compatible in a dynamic setting. They propose an incentive-compatible pricing mechanism, which is well-approximated by Uber’s additive surge pricing mechanism.

Hu et al. [129] studies the impact of surge pricing in response to a demand shock on customers and drivers in ride-hailing platforms. The authors focus on two pricing policies:

The skimming policy, where an initial high price is followed by a lower price, and the penetration policy, where a low price is followed by a high price. The authors show that the penetration policy is superior to the skimming policy in terms of the number of matches and the social efficiency and can often result in higher platform profits. However, the penetration policy requires the platform to share demand and supply information with the drivers. Afeche et al. [11] studies the optimal dynamic pricing and dispatch control of a ride-hailing platform in the presence of demand shocks and obtains structural insights into the impact of demand shocks on the optimal prices. For example, if the demand shock is predictable, the change in prices could start prior to the demand shock. Chen et al. [71] focuses on the spatial-intertemporal pricing of ride-hailing platforms. The authors propose a static pricing policy as well as a dynamic node-based and a dynamic arc-based pricing policy. They show that all three pricing policies are asymptotically optimal in a large market regime, where the demand and supply are large. However, they have different optimality gaps. Examples of other papers that use pricing as a lever to control ride-hailing platforms include Chen and Sheldon [70], Luo and Saigal [158], Ming et al. [166], Castillo et al. [66], Jacob and Roet-Green [134], Castillo [65], Abhishek et al. [3], Lu et al. [157], Balseiro et al. [38], and Hu and Zhou [130].

Özkan and Ward [172] studies the dispatch control of ride-hailing platforms. The authors use a non-stationary open queueing network model and propose a dispatch policy based on the solution to a continuous linear program. They demonstrate that this policy is asymptotically optimal in a large market regime where the arrival rates of drivers and customers are large. Varma et al. [200] studies the joint dispatch and pricing control of ride-hailing platforms using an open queueing network model. The authors propose a dynamic pricing policy as well as a modified MaxWeight policy for ride-hailing platforms. They show that these policies are asymptotically optimal among a broad class of control policies. To be specific, they show that in an asymptotic regime where arrival rates are scaled by n , all control policies have an optimality gap that is lower bounded by $O(n^{1/3})$. Moreover, their proposed

policies achieve this lower bound. Özkan [171] studies the joint dispatch and pricing control of ride-hailing platforms using a closed network model. Using a fluid model, the author demonstrates that a joint dispatch and pricing policy significantly outperforms a control policy that solely uses pricing or dispatch control.

Banerjee et al. [40] studies the dynamic dispatch control of ride-hailing platforms using a closed queueing network model. The authors propose a family of state-dependent dispatch policies called the Scaled MaxWeight policies.⁴ They show that under the complete resource pooling assumption, as the number of cars increases, the rate at which demand is dropped under each Scaled MaxWeight policy decays exponentially. Banerjee et al. [40] and this work have similarities in their modelling assumptions. However, our work allows for joint dynamic dispatch and relocation control of ride-hailing platforms. Moreover, unlike to Banerjee et al. [40], our work does not assume the complete resource pooling assumption necessarily holds. Indeed, our numerical solution approach is most useful when the complete resource pooling assumption does not hold. Furthermore, in contrast to Banerjee et al. [40] that focuses on minimizing the demand dropping probability, we assign a value to each dispatch and relocation activity and consider the objective of maximizing the total discounted value generated. Kanoria and Qian [136] studies the dynamic control of a ride-hailing platform, where besides dispatch decisions, the platform can use either an entry control or dynamic prices to control the demand. The authors propose a family of Mirror-Backpressure policies to dynamically control the platform. They show that this family of policies has an optimality gap of $O(K/T + 1/K)$ per customer, where K denotes the number of cars and T denotes the total number of customers (who will arrive over a given time horizon).

Bertsimas et al. [48] develops an optimization framework, coupled with a backbone algorithm, to dispatch cars. The authors use the NYC taxi dataset to show their proposed

4. Dai and Lin [83] shows that the maximum pressure policies are asymptotically optimal for stochastic networks under the complete resource pooling assumption. Ata and Lin [29] studies the maximum pressure policies in heavy traffic when the complete resource pooling assumption does not hold, i.e. the workload formulation is multidimensional.

policy does well. Examples of other papers that use dispatch as a control lever include Wang et al. [206], Lam and Liu [150], Feng et al. [88], Yan et al. [217], Guda and Subramanian [102], Chen et al. [69], Lyu et al. [159], Wang et al. [205], Rheingans-Yoo et al. [179], Aouad and Saritac [17], Blanchet et al. [53], and Feng et al. [89].

Examples of papers that focus on the relocation of cars in ride-hailing platforms include [55], Afèche et al. [10], and Benjaafar et al. [47]. Braverman et al. [55] focuses on centralized empty car routing in ride-hailing platforms and uses a closed queueing network model. The authors propose a static policy based on a fluid model and show that this policy is asymptotically optimal in the fluid-scale. Afèche et al. [10] studies the admission control of customers and the re-positioning of cars in ride-hailing platforms. The authors develop several insights into the interplay between centralized and de-centralized admission and re-positioning control and show that the value of these control policies are largest at moderate capacity and they increase with demand imbalances. Benjaafar et al. [47] studies the relocation problem in a multi-location rental network with a fixed number of rental units. The authors show that the optimal relocation policy can be described in terms of a desired set of states. Within this set, it is optimal not to relocate, whereas, outside this set, it is optimal to relocate. The relocation is undertaken in a minimal amount to move the state onto the boundary of the desired set. This observation is consistent with the region-of-inaction type policy we propose for ride-hailing platforms. However, contrary to Benjaafar et al. [47], our proposed policy uses both dispatch and relocation activities to move the state onto the boundary of our desired set. Examples of other papers that use relocation as a control lever include Hao et al. [108], He et al. [125], and Hosseini et al. [128].

Contrary to the previous papers, Chu et al. [73] focuses on information sharing as a control lever. To be specific, the authors study a ride-hailing platform that can reveal the destination of customers to idle drivers. They show that providing this information may hurt the platform's profits. They propose a routing policy that can align the incentives of the platform and the driver to achieve the first-best solution. Besbes et al. [49] is another

related paper that studies the capacity (required number of drivers) planning of a ride-hailing system. The focus in this paper is the impact of travel times in the capacity planning problem. The authors show that the platform should use a safety factor proportional to the offered load to the power of $2/3$, which is higher than the commonly used square-root safety factor in the queueing theory literature. Yu et al. [220] is another related paper that studies the impact of capacity controls in ride-hailing platforms. The authors show that a carefully designed regulatory policy can perform better than strict and no capacity control policies by striking a balance between the objectives of the customers, the drivers, and the platform.

The second stream of literature studies the control of queueing systems; see Stidham [190, 192] for a survey. Our work is closely related to the literature on the dynamic control of queueing systems in heavy traffic, pioneered by Harrison [110, 113, 114]. In this stream of literature, one approximates the queueing system with a Brownian model, which is easier to analyze. Early examples of this approach include Harrison and Wein [121], that studies an optimal sequencing problem for a criss-cross network, and Harrison and Wein [122], that studies a multiclass two-station closed queueing network. In both cases, the limiting Brownian models admit pathwise optimal solutions that have straightforward interpretations in the original problem. Harrison and Wein show that their policies perform well. Since then, many other researchers followed the heavy traffic approach to study manufacturing and queueing systems; see for example Wein [211, 212, 214], Chevalier and Wein [72], Krichagina and Wein [142], Reiman and Wein [178], Kumar [144], Bell and Williams [44], Plambeck et al. [175], Markowitz and Wein [163], Plambeck [176], Ata and Olsen [30, 31], [198], and Dai and Tezcan [80].

The heavy traffic approximations often give rise to singular stochastic control problems, which is the third stream of literature we review. Examples of papers that study a one-dimensional singular stochastic control problem include Harrison and Taksar [118], Taksar [195], and Dai and Yao [81, 82]. Harrison and Taksar [118] considers a singular control problem that seeks to minimize discounted costs of control that are linear and convex hold-

ing costs. The authors establish the optimality of a two-sided barrier policy. Taksar [195] extends the results to the average cost formulation. Dai and Yao [81, 82] extend the literature by considering a general impulse (and singular) control problem under the average and discounted cost objectives, respectively. They show that the impulse (singular) control problem admits a four-parameter control band (two-sided barrier) policy.

Most problems successfully analyzed in the literature by following Harrison's approach either have limiting formulations that admit a pathwise optimal solution, e.g., Harrison [112], Bell and Williams [44, 45], Ata and Kumar [27], Dai and Lin [83], or they involve solving a one-dimensional HJB equation, see for example Ata et al. [26], Ata [19], Ward and Kumar [208], Rubino and Ata [181], Ghamami and Ward [98], and Ata et al. [28]. Examples of papers that consider multidimensional singular stochastic control problems include Kushner and Martins [148], Taksar [196], and Kumar and Muthuraman [145]. Kushner and Martins [148] proposes an approximate solution to multidimensional singular stochastic control problems by considering a Markov decision process whose value function approximates the value function of the singular stochastic control problem; see Appendix B.1.8 for details. This approach is quite general. However, because it requires the solution of a Markov decision problem, it is computationally expensive, and therefore slow, for high dimensional control problems. Moreover, this method corresponds to a finite difference scheme for solving the partial differential equations (PDE), commonly referred to as the HJB equations, associated with singular control problems. Finite difference schemes tend to have a larger numerical error than other schemes that solve PDEs, such as the finite element method; see e.g., Le Dret and Lucquin [154, Sections 2.6 and 5.3]. As such, we develop a numerical method that can use the finite element method for solving the PDEs. Because the finite element method has a lower numerical error and is computationally more efficient, the method we develop is numerically robust and computationally efficient.

Taksar [196] formulates a multidimensional singular stochastic control problem as an infinite-dimensional linear program. Then, using the dual linear program, the author charac-

terizes the optimal cost function as a maximal solution to a variational inequality, commonly referred to as the HJB inequality. Kumar and Muthuraman [145] develops a numerical solution to multidimensional singular stochastic control problems with convex holding costs with two pushing directions along each axis, in the positive and negative directions, respectively. The authors limit their search to region-of-inaction policies. In these policies, there exists a region, called the region of inaction, in which no control is exerted. Once the state process reaches the boundary of the region, control is exerted to keep the process inside the region. The numerical procedure developed in Kumar and Muthuraman [145] starts with an initial guess at the region of inaction and the controls used on each point on the boundary of the region of inaction. Then, using the convexity of the value function and HJB equations, which are motivated by the HJB inequalities of Taksar [196], the boundary of the region of inaction is updated. This procedure is continued until no further updates are required. For the case of a one-dimensional control problem, the authors prove their numerical procedure converges to the optimal solution. The authors demonstrate the effectiveness of their solution in multi-dimensional settings through well-known numerical examples. The numerical solution developed in this chapter differs from the one developed in Kumar and Muthuraman [145] in three aspects: First, Kumar and Muthuraman [145] does not have state space constraints, unlike this dissertation. Second, the pushing directions in Kumar and Muthuraman [145] are restricted to be parallel to the axis, whereas we allow for general (oblique) pushing directions. Third, the numerical procedure developed in this dissertation can change the pushing directions as it updates the boundary of the region of inaction, whereas the procedure developed in Kumar and Muthuraman [145] only allows for updating the boundary and pushing directions are fixed.

2.3 Model

We divide the city into disjoint areas. An example of this for the NYC application is given in Figure 2.1. Then, we classify the cars into car classes based on their current area as well

as other relevant information such as their size and luxury status. Similarly, we classify the customers into customer classes based on their pickup area as well as other relevant information such as the list of preferred car sizes and luxury status. The number of car and customer classes need not be equal. We model the ride-hailing platform as a closed stochastic processing network as described in Harrison [113, 114]. A stochastic processing network has three basic elements: jobs, servers, and activities. In our setting, jobs correspond to cars (job classes correspond to car classes), servers correspond to customer classes, and activities correspond to either the dispatch or the centralized relocation of cars; see Figure 2.2 for the stochastic processing network model of the NYC application. Each job class has a dedicated buffer, in which the jobs of this class are stored.

A dispatch activity starts with the platform dispatching a car to pick up the next customer of a class and ends with the transition of the car from its current area to the customer's destination upon her arrival. We assume the ride-hailing platform does not know *a priori* the destination of the customers. However, it is aware of the distribution of the destinations for each customer class.⁵ In the context of the stochastic processing network, a dispatch activity corresponds to a server processing a job after which the job probabilistically transitions into another job class.⁶ A centralized relocation activity involves the platform relocating a car from one area to another for a fee.⁷ No customer is involved in this activity. In the context of the stochastic processing network, a relocation activity corresponds to moving a job from one buffer to another. Hereafter, we use the terminology of jobs and servers as opposed to cars and customers, respectively. We let Λ denote the total number of jobs circulating in the system, q denote the number of job classes, m denote the number of servers

5. Note that this model allows for the case in which the customers destinations are known. This can be modeled by using a larger set of customer classes and a routing matrix of zeros and ones.

6. That is, the arrival of a customer in the ride-hailing platform corresponds to the completion of a dispatch activity (by the server corresponding to the customer class) in the stochastic processing network.

7. We assume there is a known relocation fee that drivers will accept to relocate to the platform's desired area. This fee could depend on the car's current area as well as the platform's intended area; see Garg and Nazerzadeh [94] for an example of such a mechanism.

(resources), n denote the number of different dispatch activities, and \tilde{n} denote the number of different relocation activities available to the platform. For the NYC application, $q = 9$, $m = 9$, $n = 19$, and $\tilde{n} = 72$. We start by discussing the activities: dispatch and centralized relocation.

Dispatch Activities. For dispatch activity j , let $i(j)$ and $k(j)$ denote the unique job class and unique server associated with the activity, respectively.⁸ Denote the rate at which class k customers arrive at the ride-hailing system by μ_k for $k = 1, \dots, m$. Dispatch activity j corresponds to dispatching a class $i(j)$ car to pick up the next class $k(j)$ customer in the ride-hailing system and it is conducted at rate $\mu_{k(j)}$. In particular, one unit of dispatch activity j corresponds to the processing of class $i(j)$ jobs by server $k(j)$ for one time unit. Thus, letting $\{S_j(t) : t \geq 0\}$ be a Poisson process of rate $\mu_{k(j)}$, $S_j(t)$ denotes the number of jobs served by dispatch activity j during $[0, t]$ if it is used continuously on that time interval.

We describe the transition of jobs between buffers upon service competitions of dispatch activity j by the cumulative routing process $\{\Phi^j(t) : t \geq 0\}$ for $j = 1, \dots, n$. In particular, $\Phi_i^j(l)$ is the cumulative decrease in the number of class i jobs as a result of the first l jobs processed by dispatch activity j ; see Appendix B.1.1 for its derivation. Now, we define the flow process $\{F^j(t) : t \geq 0\}$ associated with dispatch activity j as $F^j(t) = \Phi^j(S_j(t))$ for $t \geq 0$. Then, $F_i^j(t)$ is the cumulative decrease in the number of class i jobs as a result of the first t units of dispatch activity j undertaken. In definitions of both Φ^j and F^j , a negative value indicates a net increase in the number of jobs.

We assume a class $i(j)$ job processed by dispatch activity j next becomes a class l job with probability P_{jl} . In particular, $\Phi_l^j(k)/k \rightarrow P_{jl}$ as $k \rightarrow \infty$. Also, we refer to the $n \times q$ stochastic matrix $P = (P_{jl})$ as the routing matrix. Now, we define R as the $q \times n$ matrix

8. Bramson and Williams [54] introduces the class of unitary networks, where at most one activity is associated with each buffer-server pair and each activity serves exactly one buffer and uses at most one server. The stochastic processing network introduced in this chapter is a unitary network because each dispatch activity is associated with a unique buffer and a unique server, the matrix P defined below is stochastic, and each relocation activity relocates a job from a unique origin buffer to a unique destination buffer without using any servers.

whose (i, j) -th component is given as follows:

$$R_{ij} = \mu_{k(j)}(\delta_{i,i(j)} - P_{ji}), \quad (2.1)$$

where $\delta_{il} = 1$ if $i = l$, and it is zero otherwise. Harrison [110, Equation 4.11] argues that $\mathbb{E}[F^j(t)] \sim R^j t$ and $\text{Cov}[F^j(t)] \sim \Gamma^j t$, where

$$\Gamma^j = \mu_{k(j)}\Omega^j + R^j(R^j)'/\mu_{k(j)}, \quad (2.2)$$

the vector R^j is the j -th column of R , and Ω^j is an $q \times q$ matrix whose (i, l) -th component is equal to $\Omega_{il}^j = P_{ji}(\delta_{il} - P_{jl})$. In fact, by the Strong Law of Large Numbers (SLLN) for renewal processes, it follows that $F^j(t)/t \rightarrow R^j$ as $t \rightarrow \infty$. Thus, we refer to R^j as the asymptotic drift of F^j and interpret R_{ij} as the rate at which activity j decreases the number of class i jobs. Moreover, one can deduce from the Functional Central Limit Theorem (FCLT) for renewal processes that as $t \rightarrow \infty$,

$$\frac{F^j(t \cdot) - tR^j}{\sqrt{t}} \Rightarrow \text{BM}(0, \Gamma^j),$$

where $\text{BM}(0, \Gamma^j)$ denotes an q -dimensional driftless Brownian motion with covariance matrix Γ^j . Thus, we refer to Γ^j as the asymptotic covariance of F^j for dispatch activity j .

To facilitate the analysis to follow, we let A denote the $m \times n$ capacity consumption matrix for the dispatch activities, where

$$A_{kj} = \begin{cases} 1 & \text{if resource } k \text{ undertakes dispatch activity } j, \text{ i.e., } k = k(j), \\ 0 & \text{otherwise.} \end{cases} \quad (2.3)$$

Centralized Relocation Activities. For relocation activity j , let $o(j)$ and $d(j)$ denote the origin and destination job classes associated with the relocation, respectively. We let

$\tilde{\mu}$ denote the number of jobs relocated per unit of activity. In other words, one unit of relocation activity j corresponds to the relocation of $\tilde{\mu}$ jobs from buffer $o(j)$ to buffer $d(j)$. To formalize this, we let the flow process $\{\tilde{F}^j(t) : t \geq 0\}$ associated with relocation activity j be the q -dimensional deterministic process with

$$\tilde{F}_i^j(t) = \begin{cases} \lfloor \tilde{\mu}t \rfloor & \text{for } i = o(j), \\ -\lfloor \tilde{\mu}t \rfloor & \text{for } i = d(j), \\ 0 & \text{otherwise,} \end{cases} \quad (2.4)$$

for $t \geq 0$; $\tilde{F}_i^j(t)$ is the cumulative decrease in the number of class i jobs as a result of the first t units of relocation activity j undertaken. The units of relocation and dispatch activities differ in the sense that the argument of $\tilde{F}^j(\cdot)$ is the cumulative units of relocation activity j undertaken, whereas that for $F^j(\cdot)$ is the total time the platform has engaged in dispatch activity j . Note that $\tilde{\mu}$ does not depend on j , the index of the relocation activity. We need not set $\tilde{\mu} = 1$. As the reader will see in Section 2.4, this facilitates our derivation of the approximating Brownian control problem. Recall that in our model, cars travel instantaneously. Thus, we let \tilde{R} be the $q \times \tilde{n}$ matrix, whose (i, j) -th component is

$$\tilde{R}_{ij} = \begin{cases} \tilde{\mu} & \text{for } i = o(j), \\ -\tilde{\mu} & \text{for } i = d(j), \\ 0 & \text{otherwise.} \end{cases} \quad (2.5)$$

It follows from Equation (2.4) that $\{\tilde{F}^j(t) : t \geq 0\}$ is a deterministic flow process and $|\tilde{F}_i^j(t) - \tilde{R}_{ij}t| \leq 1$ for $i = 1, \dots, q$ and $t \geq 0$. Thus, $\tilde{F}^j(t)/t \rightarrow \tilde{R}^j$ as $t \rightarrow \infty$ and we refer to \tilde{R}^j as the asymptotic drift of \tilde{F}^j . We interpret \tilde{R}_{ij} as the rate at which relocation activity j decreases the number of class i jobs.

Because relocation activities do not consume resources, we need not define a capacity consumption matrix for them. In other words, the capacity consumption matrix is only

defined for the dispatch activities. Moreover, we assume that each job class can be served by at least one dispatch activity and each server can undertake at least one dispatch activity.

Let $Q_i(t) \geq 0$ denote the number of jobs in class i and $Q(t) = (Q_i(t))$ denote the q -dimensional queue length vector at time t . The process $Q = \{Q(t), t \geq 0\}$ is called the queue length process. A dynamic control policy takes the form of a pair of processes (T, \tilde{T}) , where $T = \{T(t), t \geq 0\}$ and $\tilde{T} = \{\tilde{T}(t), t \geq 0\}$ are non-decreasing n and \tilde{n} dimensional stochastic processes, respectively. We refer to T and \tilde{T} as the allocation processes for dispatch and relocation activities, respectively. To be more specific, we interpret $T_j(t)$ and $\tilde{T}_j(t)$ as the cumulative amount of dispatch and relocation activity j undertaken up to time t , respectively. To repeat, the units of T is time while the units of \tilde{T} is cars. As a result, relocation activity j for $j = 1, \dots, \tilde{n}$ can be undertaken in any amount as long as there are jobs at the origin buffer $o(j)$ to relocate. Under policy (T, \tilde{T}) , the queue length process evolves as follows:

$$Q(t) = Q(0) - \sum_{j=1}^n F^j(T_j(t)) - \sum_{j=1}^{\tilde{n}} \tilde{F}^j(\tilde{T}_j(t)), \quad t \geq 0, \quad (2.6)$$

where $Q(0)$ is the initial queue length vector. We let $q_0 = Q(0)$. Note that the total number of jobs in the system remains fixed. That is,

$$Q(t) \geq 0 \quad \text{and} \quad \sum_{i=1}^q Q_i(t) = \Lambda, \quad t \geq 0. \quad (2.7)$$

Throughout the chapter, we let e denote the vector with appropriate dimension, which will be clear from the context. The r -dimensional idleness process associated with dispatch control T is defined as follows:

$$I(t) = te - AT(t), \quad t \geq 0. \quad (2.8)$$

The allocation pair (T, \tilde{T}) is said to be feasible if it is non-anticipating and

$$\text{Allocation processes } T \text{ and } \tilde{T} \text{ are non-decreasing with } T(0) = 0 \text{ and } \tilde{T}(0) = 0, \quad (2.9)$$

$$\text{Idleness process } I \text{ is continuous and non-decreasing.} \quad (2.10)$$

Let $\alpha > 0$ denote the interest rate, let $c \in \mathbb{R}_+^{\tilde{n}}$ denote the cost vector for the relocation activities, and let $v \in \mathbb{R}_+^n$ denote the value vector for the dispatch activities. In other words, c_j denotes the cost incurred by moving a car using relocation activity j and v_j denotes the value obtained by dispatching a car using dispatch activity j ; see Appendix B.1.2 for the vectors c and v of the NYC application. Given (c, v) , we define the cumulative net value generated by policy (T, \tilde{T}) up to time t as $V(t) = \sum_{j=1}^n v_j \mu_{k(j)} T_j(t) - \sum_{j=1}^{\tilde{n}} c_j \tilde{\mu} \tilde{T}_j(t)$ for $t \geq 0$. The platform seeks to choose a control policy (T, \tilde{T}) to

$$\text{maximize } \liminf_{t \rightarrow \infty} \mathbb{E} \left[\int_0^t e^{-\alpha s} dV(s) \right] \text{ subject to (2.6) -- (2.10).} \quad (2.11)$$

This problem appears analytically intractable. Therefore, we pursue an approximate analysis in the heavy traffic asymptotic regime below.

2.4 Approximating Brownian Control Problem

Following Harrison [114], this section advances a Brownian approximation for the problem described above. To do so, we consider a sequence of closely related systems indexed by a parameter r , whose formal limit is the Brownian control problem. We attach a superscript r to various quantities of interest corresponding to the r -th system in this sequence. The asymptotic regime we focus on is the one where the activity rates grow with r , the total number of jobs grows with \sqrt{r} , and the initial distribution of cars is the same. To be

specific, we assume that for all $r \in \mathbb{N}$ and $k = 1 \dots, m$,

$$\mu_k^r = r\mu_k, \quad \tilde{\mu}^r = r\tilde{\mu}, \quad \Lambda^r = \sqrt{r}, \quad \text{and} \quad q_0^r = \sqrt{r}q_0. \quad (2.12)$$

The assumption that $\mu_k^r = r\mu_k$ and $\tilde{\mu}^r = r\tilde{\mu}$ can be viewed as choosing a larger time unit for the r -th system so that they are the corresponding rates per time unit. It follows from Equations (2.1), (2.2), (2.5), and (2.12) that

$$R^r = rR, \quad \tilde{R}^r = r\tilde{R}, \quad \text{and} \quad \Gamma^r = r\Gamma. \quad (2.13)$$

See Appendix B.1.2 for the matrices A , R , and \tilde{R} of the NYC application. Next, we introduce the static planning problem, which helps us articulate the heavy traffic assumption and facilitates the derivation of the Brownian control problem.

Static Planning Problem. For choosing the long-run average activity rates, we consider the following linear program that ignores the randomness in the system: Find $(x, \tilde{x}) \in \mathbb{R}^n \times \mathbb{R}^{\tilde{n}}$ of activity rates so as to

$$\begin{aligned} & \text{maximize} \sum_{j=1}^n v_j \mu_{k(j)} x_j - \sum_{j=1}^{\tilde{n}} c_j \tilde{\mu} \tilde{x}_j \\ & \text{subject to} \end{aligned} \quad (2.14)$$

$$Rx + \tilde{R}\tilde{x} = 0, \quad (2.15)$$

$$Ax \leq e, \quad (2.16)$$

$$x, \tilde{x} \geq 0. \quad (2.17)$$

This formulation seeks to maximize the net rate at which value is generated by processing activities subject to three sets of constraints: Equation (2.15) ensures that the input and output rates for each buffer are equal. Equation (2.16) imposes the capacity constraints for the resources associated with dispatch activities. Equation (2.17) ensures activity rates are

non-negative. Note that the static planning problem could equivalently be stated using the matrices R^r, \tilde{R}^r in place of R, \tilde{R} , respectively, for the r -th system. That formulation has the same feasible and optimal activity rates.

Next, we state the heavy traffic assumption that is crucially used in deriving the approximating Brownian control problem.

Assumption 1 (Heavy-Traffic Assumption). *The static planning problem (2.14)-(2.17) has a unique optimal solution (x^*, \tilde{x}^*) and $Ax^* = e$.*

We refer to (x^*, \tilde{x}^*) as the platform's nominal plan. Dispatch activity j is called basic if $x_j^* > 0$. Similarly, relocation activity j is called basic if $\tilde{x}_j^* > 0$. Let b and \tilde{b} denote the numbers of basic dispatch and relocation activities, respectively. Without loss of generality, the basic dispatch activities are numbered $1, \dots, b$ and the basic relocation activities are numbered $1, \dots, \tilde{b}$. The remaining activities will be referred to as nonbasic activities. We partition x^* and \tilde{x}^* such that

$$x^* = \begin{pmatrix} x_B^* \\ x_N^* \end{pmatrix} \quad \text{and} \quad \tilde{x}^* = \begin{pmatrix} \tilde{x}_B^* \\ \tilde{x}_N^* \end{pmatrix}, \quad (2.18)$$

where $x_B^* > 0$ and $x_N^* = 0$ denote the vectors of basic and nonbasic dispatch activity rates, respectively, and $\tilde{x}_B^* > 0$ and $\tilde{x}_N^* = 0$ denote the vectors of basic and nonbasic relocation activity rates, respectively. Hereafter, we attach subscripts B and N to various vectors of interest to denote the components corresponding to basic and non-basic activities, respectively. Moreover, following Harrison [114], we partition the matrices R and A such that

$$R = [H \ J] \quad \text{and} \quad A = [B \ N], \quad (2.19)$$

where H and B have b columns; H and B are the sub-matrices of R and A that correspond to basic dispatch activities. Similarly, we partition the matrix \tilde{R} such that $\tilde{R} = [\tilde{H} \ \tilde{J}]$,

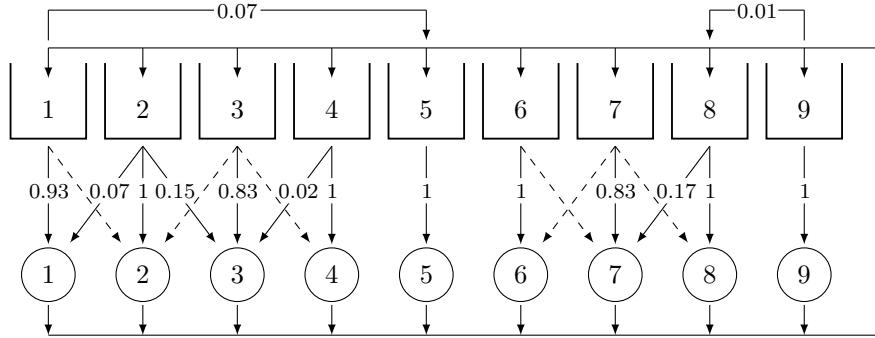


Figure 2.3: A graphical representation of the optimal solution to the static planning problem of the NYC application. The solid arrows from the buffers to the servers depict the basic dispatch activities and the rate at which they are conducted. The solid arrows between the buffers depict the basic relocation activities and the rate at which they are conducted. The dashed arrows depict the nonbasic dispatch activities. The non-basic relocation activities are not shown for visual clarity.

where \tilde{H} has \tilde{b} columns. For the NYC application, we have that $b = 13$, $\tilde{b} = 2$, and

$$x_B^* = (0.93 \ 0.07 \ 1.00 \ 0.15 \ 0.83 \ 0.02 \ 1.00 \ 1.00 \ 1.00 \ 0.83 \ 0.17 \ 1.00 \ 1.00)', \tilde{x}_B^* = (0.07 \ 0.01)';$$

see Figure 2.3 for a graphical representation of this solution. The solid arrows depict the basic dispatch and relocation activities and the rate at which these activities are conducted. The dashed arrows depict the nonbasic dispatch activities. The non-basic relocation activities are not shown for visual clarity. Server 1 serves buffers 1 and 2, server 3 serves buffers 1, 2, and 3, server 7 serves buffers 7 and 8, and all other servers only serve one buffer (the buffer with the same number). Moreover, jobs are relocated from buffer 9 to buffer 8 and from buffer 1 to buffer 5. This solution satisfies Assumption 1 and has $b = 13$ basic dispatch activities and $\tilde{b} = 2$ basic relocation activities.

If the platform were to ignore the randomness in the system, thus ignore all congestion-related costs, it would undertake the dispatch and relocation activities at rate (x^*, \tilde{x}^*) . However, because of the randomness, the platform may benefit from (dynamic) adjustments to the dispatch and relocation activity rates. To model these dynamic adjustment decisions, we let $\mathcal{S} = \{z \geq 0 : z'e = 1\}$ denote the q -dimensional simplex, which will eventually be the

state space of the approximating Brownian control problems. We denote the scaled initial queue length and the scaled queue length process by

$$z = \frac{q_0^r}{\sqrt{r}} \quad \text{and} \quad Z^r(t) = \frac{Q^r(t)}{\sqrt{r}}, \quad t \geq 0, \quad (2.20)$$

respectively, where q_0^r is the initial queue length vector (in the first system in the sequence) and z_i denotes the fraction of jobs that are initially in buffer i (in all systems in the sequence). Similarly, the stochastic process $\{Z_i^r(t) : t \geq 0\}$ captures the evolution of the fraction of jobs in buffer i in the r -th system over time for $i = 1, \dots, q$. Also, we define the centered and scaled dispatch and relocation control processes by

$$Y^r(t) = \sqrt{r}(x^{\star}t - T^r(t)) \quad \text{and} \quad \tilde{Y}^r(t) = \sqrt{r}(\tilde{x}^{\star}t - \tilde{T}^r(t)),$$

respectively. Lastly, we define the centered and scaled cumulative cost process as follows:

$$\xi^r(t) = \frac{(\sum_{j=1}^n v_j \mu_{k(j)} x_j^{\star} - \sum_{j=1}^{\tilde{n}} c_j \tilde{\mu} \tilde{x}_j^{\star}) t - V^r(t)}{\sqrt{r}}. \quad (2.21)$$

The scaling considered in this chapter is equivalent to one considered in Harrison [114]. Therefore, using an argument similar to Harrison [114], one can argue that in the heavy traffic limit, the scaled processes Z^r , Y^r , \tilde{Y}^r , and ξ^r converge weakly to Z , Y , \tilde{Y} , and ξ that

jointly satisfy the following for $t \geq 0$:

$$Z(t) = X(t) + RY(t) + \tilde{R}\tilde{Y}(t), \quad (2.22)$$

$$U(t) = KY(t), \quad (2.23)$$

$$\tilde{U}(t) = -\tilde{Y}_N(t), \quad (2.24)$$

$$U, \tilde{U} \text{ are nondecreasing with } U(0), \tilde{U}(0) \geq 0, \quad (2.25)$$

$$Z(t) \in \mathcal{S}, \quad (2.26)$$

$$\xi(t) = \sum_{j=1}^n v_j \mu_{k(j)} Y_j(t) - \sum_{j=1}^{\tilde{n}} c_j \tilde{\mu} \tilde{Y}_j(t), \quad (2.27)$$

where $X = \{X(t), t \geq 0\}$ is a $(0, \Sigma)$ Brownian motion with initial value z and covariance matrix $\Sigma = x_1^* \Gamma^1 + \dots + x_n^* \Gamma^n$, and

$$K = \begin{bmatrix} B & N \\ 0 & -I \end{bmatrix}, \quad (2.28)$$

see Appendix B.1.2 for Σ and K of the NYC application.

Remark 1. *The first m components of $U(t)$ correspond to the (scaled) idleness processes for the m servers that undertake the dispatch activities. The requirement that these components are non-decreasing ensures that dispatch controls do not use more resources than what is available at each of the m servers. The requirement that the remaining components of $U(t)$ are non-decreasing ensures that the non-basic dispatch activities are performed in non-decreasing amounts. Similarly, the requirement that $\tilde{U}(t)$ is non-decreasing ensures that the non-basic relocation activities are performed in non-decreasing amounts. Because the relocation activities do not use any resources, there are no constraints on $\tilde{Y}_B(t)$, deviations from the nominal plan corresponding to basic relocation activities. Also, note that $U(t)$ and $\tilde{U}(t)$ are \mathbb{R}^{m+n-b} and $\mathbb{R}^{\tilde{n}-\tilde{b}}$ -valued, respectively.*

Hereafter, we refer to (2.22)-(2.27) as the Brownian formulation and call the adapted

pair (Y, \tilde{Y}) an admissible control if there exists (Z, U, \tilde{U}) that satisfies (2.22)-(2.26) and the regularity condition

$$\mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(e' U)(t) \right] < \infty \quad \text{and} \quad \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(e' \tilde{U})(t) \right] < \infty. \quad (2.29)$$

We refer to (Z, U, \tilde{U}) as the extended state process associated with control (Y, \tilde{Y}) . Next, we make an assumption on the primitives of the Brownian formulation that ensures the existence of an admissible control.

Assumption 2. *We assume that $\{Ry + \tilde{R}\tilde{y} : Ky \geq 0, \tilde{y}_N \leq 0, y \in \mathbb{R}^n, \tilde{y} \in \mathbb{R}^{\tilde{n}}\} = \mathbb{R}^q$.*

This assumption ensures that the platform has controls available to impose an instantaneous displacement in any direction, although there are state constraints and a possible cost associated with such a displacement. Harrison and Williams [123, Theorem A.1 and Lemma A.2] proves that Assumption 2 is necessary and sufficient for the existence of an admissible control to (2.22)-(2.27). This assumption holds for the NYC application; see Appendix B.1.2.

Given the initial state vector z , we seek to minimize the expected discounted costs given by

$$J_z(Y, \tilde{Y}) = \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d\xi(t) \right]. \quad (2.30)$$

Using integration-by-parts, Harrison and Williams [123, Lemma 3.2] shows that

$$J_z(Y, \tilde{Y}) = \mathbb{E} \left[\alpha \int_0^\infty e^{-\alpha t} \left(\sum_{j=1}^n v_j \mu_{k(j)} Y_j(t) - \sum_{j=1}^{\tilde{n}} c_j \tilde{\mu} \tilde{Y}_j(t) \right) dt \right]. \quad (2.31)$$

Thus, the platform's problem is given as follows:⁹ Choose an adapted control (Y, \tilde{Y}) so as to

$$\inf J_z(Y, \tilde{Y}) \text{ subject to (2.22) -- (2.26).} \quad (2.32)$$

9. Proposition 8 in Appendix B.1.3 shows that there is no opportunity for arbitrage in this formulation.

We refer to this formulation as the Brownian Control Problem (BCP), and formulate an equivalent problem that has a lower-dimensional state descriptor in the next section.

2.5 Reduced Brownian Control Problem

This section develops an equivalent formulation that is of lower dimension. To this end, we define the space of reversible displacements, denoted by \mathcal{R} , as follows:¹⁰

$$\mathcal{R} = \{Hy_B + \tilde{H}\tilde{y}_B : By_B = 0\}, \quad (2.33)$$

and let the space $\mathcal{M} = \mathcal{R}^\perp$ denote its orthogonal complement. Then, let M be a matrix whose rows are a basis of space \mathcal{M} . We refer to M as the workload matrix and the dimension of the space \mathcal{M} , denoted by $\dim \mathcal{M}$, as the workload dimension, which is characterized in Lemma 1; see Appendix B.2 for its proof. Lemma 1 crucially uses the fact that each dispatch activity is associated with a unique server.

Lemma 1. *The workload dimension is equal to $p = q + m - (b + \tilde{b})$.*

We propose a natural choice for the workload matrix M based on the notion of communicating buffers, which is defined next.¹¹

Definition 2 (Communicating buffers). *Buffers i and i' communicate directly if either*

- (i) *there exists a server that serves both buffers using basic dispatch activities, or*¹²
- (ii) *there exists a basic relocation activity that moves jobs either from i to i' or from i' to i .*

10. If the difference between two state vectors is reversible, the platform can instantaneously exchange either of the state vectors for the other without affecting (U, \tilde{U}) .

11. Harrison and Williams [124, Theorem 7.2] proposes a canonical choice for the workload matrix. Because this canonical choice requires the computation of all dual extreme point optimal solutions, it is computationally expensive.

12. This condition, where buffers share a server, has been studied in the literature. Complete Resource Pooling (CRP) takes advantage of this; see e.g., Harrison and López [116] and Ata and Kumar [27].

Buffer i communicates with buffer i' if there exist buffers i_1, \dots, i_l such that $i_1 = i$, $i_l = i'$, and buffer i_s communicates directly with buffer i_{s+1} for $s = 1, \dots, l-1$.

Communication is an equivalence relation, so buffers can be partitioned into disjoint subsets such that each buffer communicates only with the buffers in the same subset. We call each communicating subset a buffer pool and denote the set of buffers in buffer pool l by \mathcal{P}_l for $l = 1, \dots, p$. We also define resource pool l as follows:

$$\mathcal{S}_l = \{k(j) : j = 1, \dots, b, \text{ such that } i(j) \in \mathcal{P}_l\}, \quad l = 1, \dots, p.$$

Each server in resource pools \mathcal{S}_l serves a buffer in \mathcal{P}_l using a basic dispatch activity. And, all servers that serve the buffers in \mathcal{P}_l are in the set \mathcal{S}_l . Moreover, resource pool \mathcal{S}_l for $l = 1, \dots, p$ are disjoint because \mathcal{P}_l for $l = 1, \dots, p$ are disjoint sets and because of Definition 2(i). Lemma 2 shows that the workload dimension is p and proposes a natural choice for the workload matrix; see Appendix B.2 for its proof. Lemma 2 crucially uses the fact that all dispatch activities associated with server k for $k = 1, \dots, m$ are conducted at rate μ_k . That is, the service rates depend on the servers only and not the activity types.

Lemma 2. *The workload dimension is equal to the number of buffer pools. Furthermore, the workload matrix $M = (M_{hi})$ is given as follows:*¹³

$$M_{hi} = \begin{cases} 1 & \text{if buffer } i \text{ belongs to buffer pool } h, \\ 0 & \text{otherwise.} \end{cases}$$

13. Buffers i and i' communicate if and only if any displacement of jobs between them is reversible; see Equation (2.33) and Harrison and Williams [124, Theorem 6.2]. Reversible displacements under Condition (i) of Definition 2 are well-understood; see e.g., Harrison and López [116] and Ata and Kumar [27]. A reversible displacement under Condition (ii) is achieved by increasing or decreasing the basic relocation activity rate.

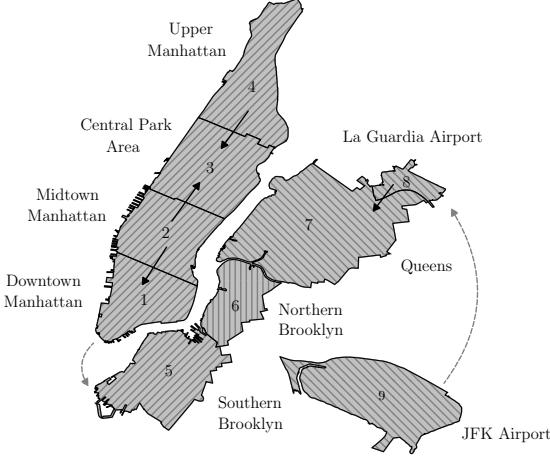


Figure 2.4: A graphical representation of the optimal solution to the static planning problem and the resulting pools. All local dispatch activities are basic. The basic non-local dispatch activities are depicted with solid black arrows. The basic relocation activities are depicted with dashed gray arrows. Areas with the same shading pattern belong to the same pool. There are three pools in total.

For the NYC application, the workload dimension is $p = 3$ and the workload matrix is

$$M = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix}.$$

See Figure 2.4 for a graphical representation of the pools in the context of the ride-hailing system. In the NYC application, the buffer pools and the server pools correspond to the same set of areas. The solid black arrows depict the basic non-local dispatch activities (activities in which a car is dispatched to pick up a customer in another area). The dashed gray arrows depict the basic relocation activities. Areas with the same shading pattern belong to the same pool. There are three pools in total. The four areas of Manhattan (downtown Manhattan, midtown Manhattan, central park area, and upper Manhattan) and southern Brooklyn create the first pool, i.e., $\mathcal{P}_1 = \{1, 2, 3, 4, 5\}$, northern Brooklyn is the second pool, i.e., $\mathcal{P}_2 = \{6\}$, and Queens, La Guardia, and JFK create the third pool, i.e., $\mathcal{P}_3 = \{7, 8, 9\}$. Given the workload matrix M characterized in Lemma 2, we define the set $\mathcal{W} = \{Mz : z \in \mathcal{S}\}$. Corollary 1 shows that \mathcal{W} is the p -dimensional simplex; see Appendix B.2 for its proof.

Corollary 1. *The set \mathcal{W} is the p -dimensional simplex, i.e., $\mathcal{W} = \{w : w'e = 1, w \in \mathbb{R}_+^p\}$.*

To facilitate the analysis to follow, we define

$$\tilde{G} = -M\tilde{J}, \quad (2.34)$$

and let the $p \times (m + n - b)$ dimensional matrix G be a solution¹⁴ to

$$GK = MR. \quad (2.35)$$

Next, we introduce the reduced Brownian formulation, where the system state is given by the workload process W and the control is the adapted pair (U, \tilde{U}) . To be more specific, W and (U, \tilde{U}) must jointly satisfy the following for $t \geq 0$,

$$W(t) = \chi(t) + GU(t) + \tilde{G}\tilde{U}(t), \quad (2.36)$$

$$U, \tilde{U} \text{ are nondecreasing with } U(0), \tilde{U}(0) \geq 0, \quad (2.37)$$

$$W(t) \in \mathcal{W}, \quad (2.38)$$

where $\chi = \{\chi(t), t \geq 0\}$ is a $(0, M\Sigma M')$ Brownian motion with initial value $w = Mz$. Recall that the matrix G is $p \times (m + n - b)$ dimensional. Its first r columns correspond to the displacements in the workload process resulting from idling a server, whereas its last $n - b$ columns correspond to displacements in the workload process resulting from non-basic dispatch activities. Similarly, \tilde{G} is $p \times (\tilde{n} - \tilde{b})$ dimensional and its columns correspond to displacements in the workload process resulting from non-basic relocation activities. The adapted control (U, \tilde{U}) is admissible for the reduced Brownian formulation if it satisfies the regularity condition (2.29) and there exists a workload process W that satisfies (2.36)-

14. Harrison and Williams [123, Lemma 4.2] proves that such a G exists. Given K , M , and R , Equation (2.35) corresponds to a system of linear equations and we can solve for G by solving this system of linear equations.

(2.38).¹⁵

Lemma 3 facilitates a cost structure for the reduced Brownian formulation (2.36)-(2.38) that is equivalent to the cost structure of the Brownian formulation (2.22)-(2.27).

Lemma 3 (Harrison and Williams [123, Lemmas 4.3 and 4.4]). *There exists $\pi \in \mathbb{R}^q$, $\kappa \in \mathbb{R}^{m+n-b}$, and $\tilde{\kappa} \in \mathbb{R}^{\tilde{n}-\tilde{b}}$ that solve*

$$\pi' R + \kappa' K = v', \quad \pi' \tilde{H} = -c'_B, \quad \text{and} \quad \pi' \tilde{J} - \tilde{\kappa}' = -c'_N.$$

Moreover, if $z \in \mathbb{R}^q$, $u \in \mathbb{R}^{m+n-b}$, and $\tilde{u} \in \mathbb{R}^{\tilde{n}-\tilde{b}}$ satisfy $Mz = Gu + \tilde{G}\tilde{u}$, then, there exists a unique $(y, \tilde{y}) \in \mathbb{R}^n \times \mathbb{R}^{\tilde{n}}$ such that $z = Ry + \tilde{R}\tilde{y}$, $u = Ky$, and $\tilde{u} = -\tilde{y}_N$. Furthermore, $\pi' z + \kappa' u + \tilde{\kappa}' \tilde{u} = v' y - c' \tilde{y}$.

Lemma 3 helps us express the costs and values of the activities in terms of buffer and idleness costs. Namely, we interpret π_i as the holding cost rate for class i queue. Similarly, we interpret each component of κ and $\tilde{\kappa}$ as the cost rate of an increase in the corresponding component of U and \tilde{U} , e.g., the increase in cost due to idling a server or engaging in a non-basic activity. We observe from Equations (2.1) and (2.5) that the columns of R and \tilde{R} sum to zero. Therefore, π is defined only up to an additive constant; see Appendix B.1.2 for the vectors π , κ , and $\tilde{\kappa}$ of the NYC application.

In what follows, without loss of generality, we assume $\pi \geq e$ and define the effective holding cost function as follows:

$$g(w) = \min\{\alpha\pi' z : Mz = w, z \in \mathcal{S}\}. \quad (2.39)$$

The effective holding cost function g is well defined since \mathcal{S} is compact. Also, we define

15. The reduced Brownian formulation introduced in Harrison and Williams [123] includes the extra condition that $U(t) \in \{Ky : y \in \mathbb{R}^n\}$ for $t \geq 0$. This condition is necessary for ensuring that there exists an admissible control (Y, \tilde{Y}) to (2.22)-(2.27) whose extended state process includes U . We can omit this condition because $\{Ky : y \in \mathbb{R}^n\} = \mathbb{R}^{m+n-b}$ in our case; see Lemma 12 in Appendix B.1.4.

$i^* : \{1, \dots, p\} \rightarrow \{1, \dots, m\}$ as follows:

$$i^*(l) = \operatorname{argmin} \{\pi_i : i \in \mathcal{P}_l\}, \quad l = 1, \dots, p. \quad (2.40)$$

The function $i^*(\cdot)$ identifies the cheapest buffer in each buffer pool to hold the workload in. If there is a tie in (2.40), the buffer with the smallest index is chosen. We then define the lifting map $\Delta : \mathcal{W} \rightarrow \mathcal{S}$ as follows:

$$\Delta(w)_i = \begin{cases} w_l & \text{if } i \in \mathcal{P}_l \text{ and } i = i^*(l), \\ 0 & \text{otherwise} \end{cases}$$

for $i = 1, \dots, q$. Note that $w = M\Delta(w)$.

Lemma 4 shows the effective holding cost function is linear and that for each buffer pool l it is optimal to keep all its workload w_l in the cheapest buffer in the pool (i.e., buffer $i^*(l)$) as prescribed by the lifting map Δ ; see Appendix B.2 for its proof.

Lemma 4. *The effective holding cost function g can be expressed as follows:*

$$g(w) = \alpha \pi' \Delta(w) = \alpha \sum_{l=1}^p \pi_{i^*(l)} w_l, \quad w \in \mathcal{W}.$$

For the NYC application, we have that

$$g(w) = \alpha (12.52 w_1 + 16.05 w_2 + w_3) \quad \text{and} \quad \Delta(w) = (0 \ 0 \ 0 \ w_1 \ 0 \ w_2 \ 0 \ 0 \ w_3)'.$$

Given the initial value of the workload process $W(0) = w$, we let

$$\check{J}_w(U, \tilde{U}) = \mathbb{E} \left[\int_0^\infty e^{-\alpha t} g(W(t)) dt \right] + \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(\kappa' U + \tilde{\kappa}' \tilde{U})(t) \right]. \quad (2.41)$$

The Reduced Brownian Control Problem (RBCP) can be stated as follows: Choose an

adapted control (U, \tilde{U}) so as to

$$\inf \check{J}_w(U, \tilde{U}) \text{ subject to (2.36) -- (2.38).} \quad (2.42)$$

Harrison and Williams [123] establishes an equivalence between the BCP (2.32) and the RBCP (2.42). In particular, given the state $Z(t)$ in the BCP, one sets $W(t) = MZ(t)$ to arrive at the equivalent state descriptor in the RBCP. To go in the other direction, given the state $W(t)$ in the RBCP, one sets $Z(t) = \Delta(W(t))$ to arrive at the equivalent state descriptor in the BCP; see Lemma 13 in Appendix B.1.5 for further details.

The analysis of the one-dimensional case, i.e., $p = 1$, is trivial.¹⁶ Henceforth, we assume $p > 1$, and derive a $(p - 1)$ -dimensional equivalent workload formulation in the next subsection.

2.6 Equivalent Workload Formulation

As a preliminary to introducing the equivalent workload formulation, we introduce the following notation. First, letting I denote the $(p - 1)$ -dimensional identity matrix, we define the $(p - 1) \times p$ dimensional matrix $L = [I \ 0]$. Then, we define the set $\hat{\mathcal{W}}$ on which the workload process lives as follows:

$$\hat{\mathcal{W}} = \{\hat{w} \in \mathbb{R}^{p-1} : \hat{w}'e \leq 1, \hat{w} \geq 0\}.$$

The matrix L can be viewed as a projection with domain \mathcal{W} , i.e., $L : \mathcal{W} \rightarrow \mathbb{R}^{p-1}$. In particular, it projects the set \mathcal{W} onto $\hat{\mathcal{W}}$ as shown in Lemma 5; see Appendix B.2 for its proof.

Lemma 5. *We have that $\hat{\mathcal{W}} = L(\mathcal{W})$. Moreover, the mapping $L : \mathcal{W} \rightarrow \hat{\mathcal{W}}$ is one-to-one and onto.*

16. To see this, note that when $p = 1$, $\mathcal{W} = \{1\}$ and X is a driftless Brownian motion with initial value $w = 1$ and covariance matrix $\Sigma = 0$. Therefore, it is optimal to use no controls.

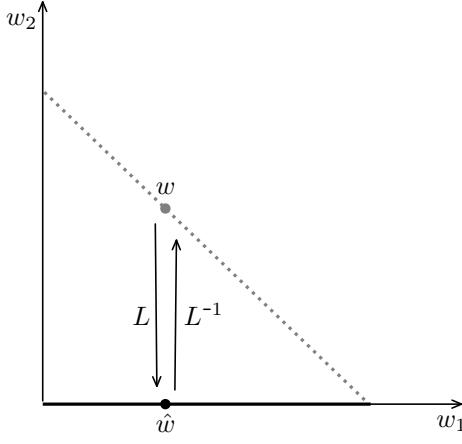


Figure 2.5: An illustration of the projection mapping L and its inverse L^{-1} for $p = 2$. The dashed gray and solid black lines represent \mathcal{W} and $\hat{\mathcal{W}}$, respectively; L maps $w \in \mathcal{W}$ to $\hat{w} \in \hat{\mathcal{W}}$.

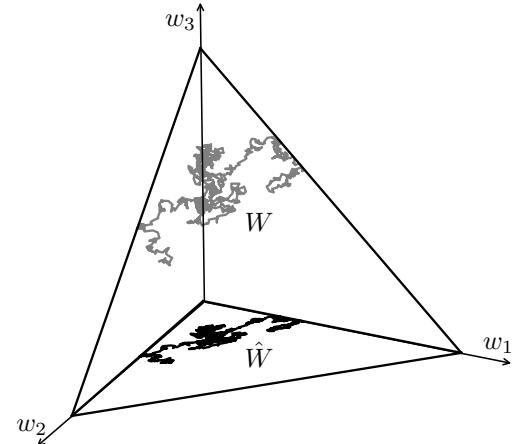


Figure 2.6: An illustration of the relationship between W and \hat{W} for $p = 3$. The gray and black curves depict a sample path of W and \hat{W} , respectively.

In essence, the projection Lw of w drops the last component of w ; see Figure 2.5. Then, defining

$$\hat{W}(t) = LW(t), \quad t \geq 0,$$

we note that for $t \geq 0$ and $k = 1, \dots, p-1$, we have $\hat{W}_k(t) = W_k(t)$. That is, $\hat{W}(t)$ is the projection of $W(t)$ on $\hat{\mathcal{W}}$; see Figure 2.6.

Next, we formulate a control problem with state process \hat{W} that is equivalent to the RBCP (2.42). For technical convenience, we replace g (defined on \mathcal{W}) with an equivalent holding cost function h defined on $\hat{\mathcal{W}}$. To be specific, we let

$$h(\hat{w}) = g(L^{-1}\hat{w}), \quad \hat{w} \in \hat{\mathcal{W}}, \quad (2.43)$$

where $L^{-1} : \hat{\mathcal{W}} \rightarrow \mathcal{W}$ is the inverse of the projection L . Note, in particular, that

$$h(\hat{w}) = \alpha \sum_{l=1}^{p-1} \pi_{i^*(l)} \hat{w}_l + \alpha \pi_{i^*(p)} \left(1 - \sum_{l=1}^{p-1} \hat{w}_l \right), \quad \hat{w} \in \hat{\mathcal{W}}.$$

Then, we define the $d = m + (n - b) + (\tilde{n} - \tilde{b})$ -dimensional vector $\hat{\kappa}$ as follows:

$$\hat{\kappa} = \begin{pmatrix} \kappa \\ \tilde{\kappa} \end{pmatrix}. \quad (2.44)$$

We also define the $(p - 1) \times d$ dimensional matrix \hat{G} as follows:

$$\hat{G} = L [G \ \tilde{G}]. \quad (2.45)$$

Note that \hat{G} is the sub-matrix of $[G \ \tilde{G}]$, which comprises of its first $p - 1$ rows.

Given a d -dimensional, adapted, non-decreasing stochastic process \hat{U} and a $(p - 1)$ -dimensional $(0, \hat{\Sigma})$ Brownian motion $\hat{\chi}$ with covariance matrix $\hat{\Sigma} = LM\Sigma M'L'$ and initial value $\hat{w} = LMz$, define \hat{W} as follows:

$$\hat{W}(t) = \hat{\chi}(t) + \hat{G}\hat{U}(t), \quad t \geq 0.$$

Moreover, we define the cost $\hat{J}_{\hat{w}}(\hat{U})$ associated with control \hat{U} as follows:

$$\hat{J}_{\hat{w}}(\hat{U}) = \mathbb{E} \left[\int_0^\infty e^{-\alpha t} h(\hat{W}(t)) dt \right] + \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(\hat{\kappa}' \hat{U})(t) \right]. \quad (2.46)$$

The Equivalent Workload Formulation (EWF) is defined as follows: Choose an adapted control \hat{U} so as to

$$\inf \hat{J}_{\hat{w}}(\hat{U}) \quad (2.47)$$

subject to

$$\hat{W}(t) = \hat{\chi}(t) + \hat{G}\hat{U}(t), \quad t \geq 0, \quad (2.48)$$

$$\hat{U} \text{ is nondecreasing with } \hat{U}(0) \geq 0, t \geq 0, \quad (2.49)$$

$$\hat{W}(t) \in \hat{\mathcal{W}}, \quad t \geq 0, \quad (2.50)$$

where $\hat{\chi}$ is a $(0, \hat{\Sigma})$ Brownian motion with initial value \hat{w} . We call the control \hat{U} admissible if there exists \hat{W} that satisfies (2.48)-(2.50) and the regularity condition

$$\mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(e' \hat{U})(t) \right] < \infty. \quad (2.51)$$

We refer to \hat{W} as the state process associated with control \hat{U} . Corollary 2 establishes an equivalence between the RBCP (2.42) and the EWF (2.47)-(2.50).

Corollary 2. *The EWF (2.47)-(2.50) is equivalent to the RBCP (2.42) in the following sense: Fix $\epsilon \geq 0$.*

(i) *Suppose that (U, \tilde{U}) is an ϵ -optimal control for the RBCP (2.42) with Brownian motion χ and state process W . Then, $\hat{U} = \begin{pmatrix} U \\ \tilde{U} \end{pmatrix}$ is an ϵ -optimal control for the EWF (2.47)-(2.50) with Brownian motion $\hat{\chi}$ and state process \hat{W} given by $\hat{\chi}_k(t) = \chi_k(t)$ and $\hat{W}_k(t) = W_k(t)$ for $k = 1, \dots, p-1$ and $t \geq 0$.*

(ii) *Suppose that \hat{U} is an ϵ -optimal control for the EWF (2.47)-(2.50) with Brownian motion $\hat{\chi}$ and state process \hat{W} . Let U denote the first $m+n-b$ components of \hat{U} and let \tilde{U} denote the remaining $\tilde{n}-\tilde{b}$ components. Then, (U, \tilde{U}) is an ϵ -optimal control for the RBCP (2.42) with Brownian motion χ and state process W given for $t \geq 0$ by*

$$\begin{aligned} \chi(t) &= \left(\hat{\chi}_1(t), \dots, \hat{\chi}_{p-1}(t), - \sum_{l=1}^{p-1} \hat{\chi}_l(t) \right), \\ W(t) &= \left(\hat{W}_1(t), \dots, \hat{W}_{p-1}(t), 1 - \sum_{l=1}^{p-1} \hat{W}_l(t) \right). \end{aligned}$$

Then, it follows from Corollary 2, by virtue of Lemma 13 (see Appendix B.1.5), that the EWF (2.47)-(2.50) is equivalent to the BCP (2.32). Recall that control \hat{U} is a d -dimensional stochastic process. Hereafter, for ease of exposition, we refer to $l = 1, \dots, d$ as pushing controls and refer to \hat{G}^l for $l = 1, \dots, d$ as the corresponding pushing directions.

Definition 3 (Dominated Pushing Control). *Pushing control j is dominated if there exist a subset $\mathcal{L} \subset \{1, \dots, d\} \setminus \{j\}$ of pushing controls and scalars $\hat{u}_l > 0$ for $l \in \mathcal{L}$ that jointly*

satisfy

$$\hat{G}^j = \sum_{l \in \mathcal{L}} \hat{G}^l \hat{u}_l \quad \text{and} \quad \hat{\kappa}_j \geq \sum_{l \in \mathcal{L}} \hat{\kappa}_l \hat{u}_l. \quad (2.52)$$

Lemma 6 shows that it suffices to focus on pushing controls that are not dominated; see Appendix B.2 for its proof.

Lemma 6. *In the EWF (2.47)-(2.50), we can set $\hat{U}_l \equiv 0$ for dominated pushing controls l without loss of optimality.*

In light of Lemma 6, the analysis below assumes that no pushing control is dominated in the sense of Definition 3. Furthermore, without loss of generality, we assume \hat{G}^l for $l = 1, \dots, d$ are normalized such that $\|\hat{G}^l\| = 1$, where $\|\cdot\|$ denotes the Euclidean norm.¹⁷ For the NYC application, we have

$$\hat{\mathcal{W}} = \{\hat{w} : \hat{w}_1 + \hat{w}_2 \leq 1, \hat{w} \in \mathbb{R}_+^2\} \quad \text{and} \quad h(\hat{w}) = \alpha(11.52\hat{w}_1 + 15.05\hat{w}_2 + 1), \quad \hat{w} \in \hat{\mathcal{W}},$$

and the covariance matrix is

$$\hat{\Sigma} = \begin{bmatrix} 0.21 & -0.01 \\ -0.01 & 0.02 \end{bmatrix}.$$

The NYC application has five non-dominated pushing controls with

$$\hat{\kappa} = (0.01, 0.70, 1.50, 0.41, 0.53)'$$

The pushing directions, \hat{G}^l , of the non-dominated pushing controls are depicted in Figure 2.7. Three of these pushing controls correspond to idling servers and two of them correspond to undertaking non-basic dispatch activities. Pushing control 1, depicted in yellow in Figure

17. The normalization is achieved by multiplying \hat{G}^l and $\hat{\kappa}_l$ by the same factor for each $l = 1, \dots, d$.

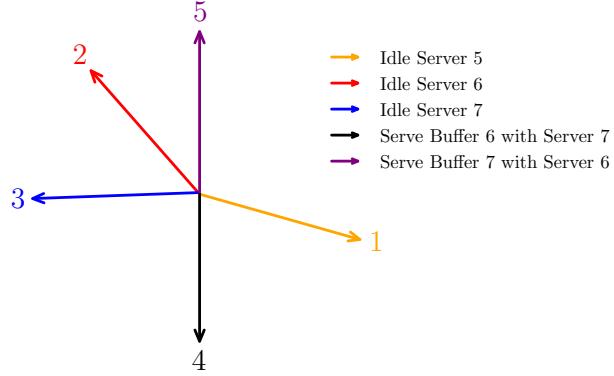


Figure 2.7: The non-dominated pushing controls and the corresponding pushing directions for the NYC application.

2.7, corresponds to idling server 5 (southern Brooklyn). Pushing control 2, depicted in red, corresponds to idling server 6 (northern Brooklyn). Pushing control 3, depicted in blue, corresponds to idling server 7 (Queens). Pushing control 4, depicted in black, corresponds to the non-basic dispatch activity of serving buffer 6 (northern Brooklyn) with server 7 (Queens). Pushing control 5, depicted in purple, corresponds to the non-basic dispatch activity of serving buffer 7 (Queens) with server 6 (northern Brooklyn).

2.7 Solving the Equivalent Workload Formulation

This section develops a computational method to solve the EWF (2.47)-(2.50). For notational simplicity, in the remainder of this section, we use dummy variables w and u instead of \hat{w} and \hat{u} . We let $\mathcal{A}(w)$ denote the set of admissible controls \hat{U} to the EWF (2.47)-(2.50) with initial workload $w \in \hat{\mathcal{W}}$ and define its (optimal) value function $\phi : \hat{\mathcal{W}} \rightarrow \mathbb{R}$ as follows:

$$\phi(w) = \inf_{\hat{U} \in \mathcal{A}(w)} \hat{J}_w(\hat{U}), \quad w \in \hat{\mathcal{W}}. \quad (2.53)$$

Proposition 3 shows that the value function is convex; see Appendix B.2 for its proof.

Proposition 3. *The (optimal) value function $\phi : \hat{\mathcal{W}} \rightarrow \mathbb{R}$ is convex.*

We let $C^2(\hat{\mathcal{W}})$ denote the space of functions $f : \hat{\mathcal{W}} \rightarrow \mathbb{R}$ that are twice continuously differ-

entiable up to the boundary. As it is often done in the singular stochastic control literature (see e.g., Kumar and Muthuraman [145]), we make the following two assumptions on the EWF (2.47)-(2.50).

Assumption 3. *There exist constants $C_1, C_2 > 0$ such that $C_1\|w\| \leq \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} w_i w_j \leq C_2\|w\|$ for $w \in \mathbb{R}^{p-1}$.*

Assumption 4. *For each initial workload $w \in \hat{\mathcal{W}}$, there exists an optimal control process \hat{U}_w , under which there is a unique Markovian state process. Moreover, the optimal value function is twice continuously differentiable, i.e., $\phi \in C^2(\hat{W})$.¹⁸*

We define the elliptic linear operator $\mathcal{L} = \frac{1}{2} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_{ij}$, where $\hat{\sigma}_{ij}$ denotes the (i, j) -th component of the covariance matrix $\hat{\Sigma}$ and ∂_{ij} denotes the differential operator $\partial^2/\partial w_i \partial w_j$ for $i, j = 1, \dots, p-1$. The next proposition introduces a set of inequalities that will be useful for our solution approach. It extends Proposition 1 of Kumar and Muthuraman [145] to the case with an arbitrary number of pushing controls, oblique pushing directions, and state space constraints; see Appendix B.2 for its proof.

Proposition 4. *For $\epsilon \geq 0$, suppose there exists a bounded $f \in C^2(\hat{\mathcal{W}})$ such that*

$$\mathcal{L}f(w) - \alpha f(w) + h(w) \geq -\epsilon, \quad w \in \hat{\mathcal{W}} \quad (2.54)$$

$$\nabla f(w)' \hat{G}^l + \hat{\kappa}_l \geq -\epsilon, \quad w \in \hat{\mathcal{W}} \text{ and } l = 1, \dots, d. \quad (2.55)$$

Then, there exists a function $\eta : \hat{\mathcal{W}} \rightarrow \mathbb{R}_+$ such that

$$f(w) \leq \phi(w) + \epsilon\eta(w), \quad w \in \hat{\mathcal{W}}. \quad (2.56)$$

¹⁸ Budhiraja and Ross [58] establishes the existence of an optimal control for the case in which $\hat{\mathcal{W}}$ is a closed convex cone.

Proposition 4 subsumes the usual "verification theorem" as follows: If $f \in C^2(\hat{\mathcal{W}})$ satisfies (2.54)-(2.55) with $\epsilon = 0$, then it follows from Proposition 4 that $\phi(w) = f(w)$ for $w \in \hat{\mathcal{W}}$. It is not known whether there exists $f \in C^2(\hat{\mathcal{W}})$ that satisfies (2.54)-(2.55) with $\epsilon = 0$. However, Taksar [196, Proposition 7.4] shows that when $\hat{\mathcal{W}} = \mathbb{R}^{p-1}$, for any $\epsilon > 0$, there exists a bounded $f \in C^2(\hat{\mathcal{W}})$ that satisfies (2.54)-(2.55).

Next, we introduce some preliminary definitions in order to describe our algorithm. First, for $w, w' \in \hat{\mathcal{W}}$, we let

$$C(w, w') = \underset{u}{\text{minimize}} \hat{\kappa}' u \quad (2.57)$$

$$\text{subject to } \hat{G}u = w' - w, \quad (2.58)$$

$$u \geq 0 \quad (2.59)$$

denote the cost of an instantaneous displacement from w to w' . In (2.57)-(2.59), the objective is to find $u \geq 0$ with the lowest cost such that it creates a displacement in the amount of $w' - w$. Second, restricting attention to the region-of-inaction type control policies and given a region of inaction \mathcal{B} along with the pushing directions on its boundary, we define the corresponding value function and the control policy achieving it. To this end, let \mathcal{B} be a convex subset of $\hat{\mathcal{W}}$ and assume there is a partition $\partial\mathcal{B}^1, \dots, \partial\mathcal{B}^d$ of its boundary, $\partial\mathcal{B}$, such that pushing control l is used for $w \in \partial\mathcal{B}^l$. Then, we associate with the region of inaction \mathcal{B} and the partition $\partial\mathcal{B}^1, \dots, \partial\mathcal{B}^d$ of its boundary $\partial\mathcal{B}$ the function $f : \hat{\mathcal{W}} \rightarrow \mathbb{R}$ that satisfies the following linear elliptic partial differential equation (PDE) on \mathcal{B} :

$$\mathcal{L}f(w) - \alpha f(w) + h(w) = 0, \quad w \in \text{int}(\mathcal{B}) \quad (2.60)$$

subject to the boundary condition

$$\nabla f(w)' \hat{G}^l + \hat{\kappa}_l = 0, \quad w \in \partial\mathcal{B}^l \text{ and } l = 1, \dots, d. \quad (2.61)$$

Moreover, f is defined on $\hat{\mathcal{W}} \setminus \mathcal{B}$ as follows:

$$f(w) = \min_{w' \in \partial\mathcal{B}} \{C(w, w') + f(w')\}, \quad w \in \hat{\mathcal{W}} \setminus \mathcal{B}. \quad (2.62)$$

We interpret this function as the value function corresponding to the following region-of-inaction type control policy \hat{U} : Exert no control in the interior of \mathcal{B} , $\text{int}(\mathcal{B})$, and keep the state process \hat{W} in \mathcal{B} by using pushing control l on $\partial\mathcal{B}^l$ minimally for $l = 1, \dots, d$. If the initial workload $w \in \hat{\mathcal{W}} \setminus \mathcal{B}$, we move the state process \hat{W} instantaneously from w to the point $\text{argmin}_{w' \in \partial\mathcal{B}} \{C(w, w') + f(w')\}$, i.e., the point on the boundary $\partial\mathcal{B}$ for which the cost associated with this instantaneous move plus the value function at the new point is the smallest. To be more specific, as shown in Appendix B.1.6, $f(w)$ can be interpreted as the expected discounted cost (see Equation (2.46)) starting at the initial workload $w \in \hat{\mathcal{W}}$ under the control policy \hat{U} .

As is commonly done in the singular stochastic control literature (see e.g., Kumar and Muthuraman [145] and Muthuraman and Kumar [167]), we limit our search to the region-of-inaction type control policies. In other words, we limit our search to functions f that satisfy (2.60)-(2.62) for some closed convex set \mathcal{B} and some partition $\partial\mathcal{B}^1, \dots, \partial\mathcal{B}^d$ of its boundary $\partial\mathcal{B}$ such that pushing control l is used for $w \in \partial\mathcal{B}^l$.

Initialization. The algorithm is initialized (step 0) with the region on inaction set $\mathcal{B}^0 = \hat{\mathcal{W}}$, a closed convex set, and a partition $\partial\mathcal{B}^{0,1}, \dots, \partial\mathcal{B}^{0,d}$ of its boundary, $\partial\mathcal{B}^0$, such that pushing control l is used for $w \in \partial\mathcal{B}^{0,l}$. Let \hat{U}^0 denote the region-of-inaction type control policy associated with the region of inaction \mathcal{B}^0 and the partition $\partial\mathcal{B}^{0,1}, \dots, \partial\mathcal{B}^{0,d}$. For the NYC application, we start with the partition depicted in Figure 2.8a. The arrows on each part of the boundary represent the pushing control used on that part of the boundary; see Figure 2.7 for the complete list of the non-dominated pushing controls for the NYC application. As depicted in Figure 2.8a, pushing control 1 is used on the vertical face, pushing control 2 is used on the horizontal face, and pushing control 3 is used on the the

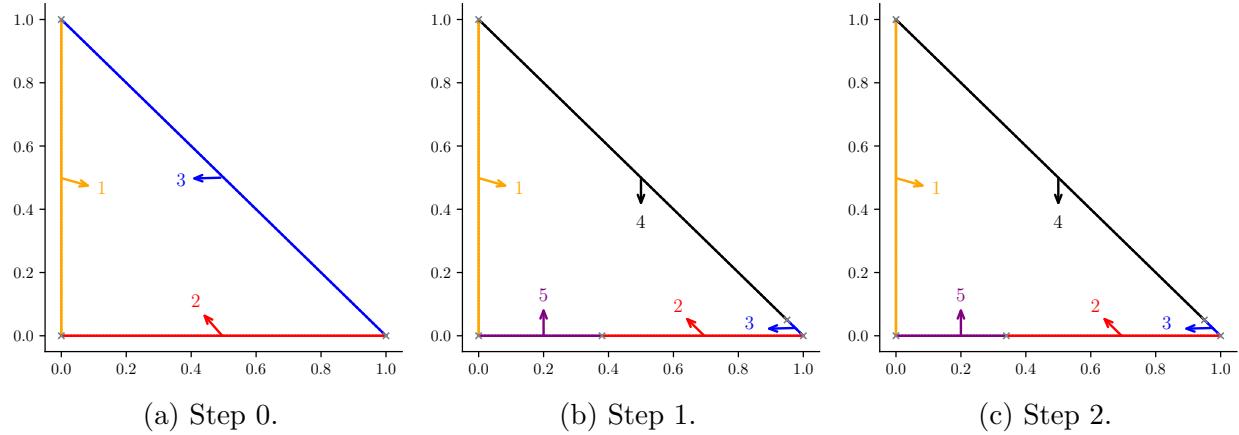


Figure 2.8: An illustration of the region of inaction and the pushing controls used on its boundary in each step of the algorithm for the NYC application.

diagonal face. Pushing controls 4 and 5 are not used under control policy \hat{U}^0 . In the interior of \mathcal{B}^0 , $\text{int}(\mathcal{B}^0)$, no control is exerted. As will be discussed in Section 2.8, in the context of the ride-hailing system, \hat{U}^0 corresponds to using the optimal solution to the static planning problem (2.14)-(2.17) in the interior $\text{int}(\mathcal{B}^0)$ and using the pushing controls depicted in Figure 2.8a on the boundary $\partial\mathcal{B}^0$.

To characterize the system performance under control policy \hat{U}^0 , we consider the function f_0 characterized by (2.60)-(2.62). We refer to $f_0 : \hat{\mathcal{W}} \rightarrow \mathbb{R}$ as the value function corresponding to control policy \hat{U}^0 . To be specific, as explained in Appendix B.1.6, $f_0(w)$ can be interpreted as the expected discounted cost (see Equation (2.46)) starting at the initial workload $w \in \mathcal{B}^0$ under the control policy \hat{U}^0 , depicted in Figure 2.8a. One can use any numerical method of choice to solve Equations (2.60)-(2.61) and the specific method used does not impact our algorithm. In Appendix B.1.7, we describe a numerical solution based on the finite element method, which is numerically robust and computationally efficient. To this end, we discretize the region of inaction as illustrated in Figure 2.9 and iterate over the boundary points as discussed next. Keeping all else the same as they were at the beginning of the iteration, we decide whether each boundary point should be moved inward, outward, or kept the same. If the boundary point is moved, the pushing control is not updated. If the

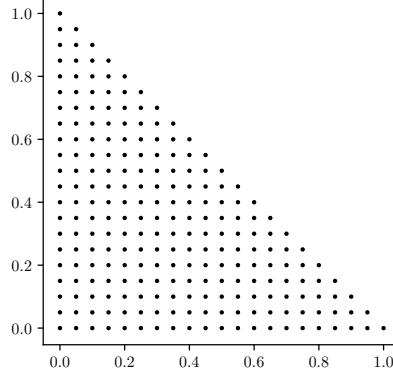


Figure 2.9: An example of the discretized region of inaction with a discretization step size of 0.05.

boundary point is not moved, we then decide whether the pushing control at the boundary point should be changed. Because each boundary point is considered in isolation while keeping others the same, the order in which we consider different boundary points is immaterial as the reader will see below. After each boundary point is considered, we check if the updated region of inaction is convex. If not, a convexification step follows. Next, we describe each step of the update procedure in detail.

The Update Procedure. Fix a pushing control l and a boundary point $w \in \partial\mathcal{B}^{0,l}$. We decide whether to move the boundary point w by checking the local behavior of the value function f_0 along the currently used pushing direction \hat{G}^l at w , i.e., whether it is locally concave, locally convex, or locally linear. Consequently, we have three cases to consider.

As a preliminary to discussing these three cases, consider the following scenario: The directional derivative of the value function f_0 along a feasible pushing direction \hat{G}^j is strictly smaller than $-\hat{\kappa}_j$, the negative of the control cost, in the vicinity of a boundary point w . That is, for some $\delta > 0$,

$$\nabla f_0(w + z\hat{G}^j)' \hat{G}^j < -\hat{\kappa}_j, \quad z \in (0, \delta); \quad (2.63)$$

see Figures 2.10a and 2.10b for an illustration of the two cases for which (2.63) holds. By

integrating both sides of (2.63) along the pushing direction \hat{G}^j , we arrive at the following:

$$f_0(w + \delta \hat{G}^j) - f_0(w) = \int_0^\delta \nabla f_0(w + z \hat{G}^j)' \hat{G}^j dz < -\hat{\kappa}_j \delta. \quad (2.64)$$

We consider the control policy \check{U} that moves the process from w to $w + \delta \hat{G}^j$ instantaneously and then follows \hat{U}^0 . It follows from (2.64) that

$$\hat{J}_w(\check{U}) = \hat{\kappa}_j \delta + f_0(w + \delta \hat{G}^j) < f_0(w) = \hat{J}_w(\hat{U}^0).$$

In other words, the initial deviation reduces the discounted cost and \check{U} is an improvement over \hat{U}^0 . This observation is used crucially in the updating procedure as discussed below.

Our algorithm updates the region of inaction in one of two instances: either the control policy can be improved with an initial deviation as discussed above, or the boundary of the region of inaction can be moved to smoothly paste the value function. The first instance is illustrated in Figures 2.10a and 2.10b. The second instance is illustrated in Figure 2.10c. Therefore, the three panels of Figure 2.10 are the only relevant instances.

Case 1. The value function is locally concave along the pushing direction \hat{G}^l , i.e., $(\hat{G}^l)' \nabla^2 f_0(w) \hat{G}^l < 0$. In this case, we move the boundary point w inward along \hat{G}^l by

$$\delta^* = \min \{ \delta \in \mathbb{R}_+ : \delta \text{ is a local minimizer of } \nabla f_0(w + \delta \hat{G}^l) \}.$$

Because the second derivative of the value function f_0 along the pushing direction \hat{G}^l at $w \in \partial \mathcal{B}^{0,l}$ is negative, i.e., $(\hat{G}^l)' \nabla^2 f_0(w) \hat{G}^l < 0$, as illustrated in Figure 2.10a, the directional derivative of the value function f_0 along the pushing direction \hat{G}^l is decreasing and strictly smaller than $-\hat{\kappa}_l$ in the vicinity of w , i.e., (2.63) holds, because f_0 satisfies the boundary condition (2.61) at w . Therefore, we can improve the control policy by initially deviating from \hat{U}^0 and moving w inward along the pushing direction \hat{G}^l . Moving the boundary point w to $w + \delta^* \hat{G}^l$ not only yields the lowest directional derivative as illustrated in Figure 2.10a

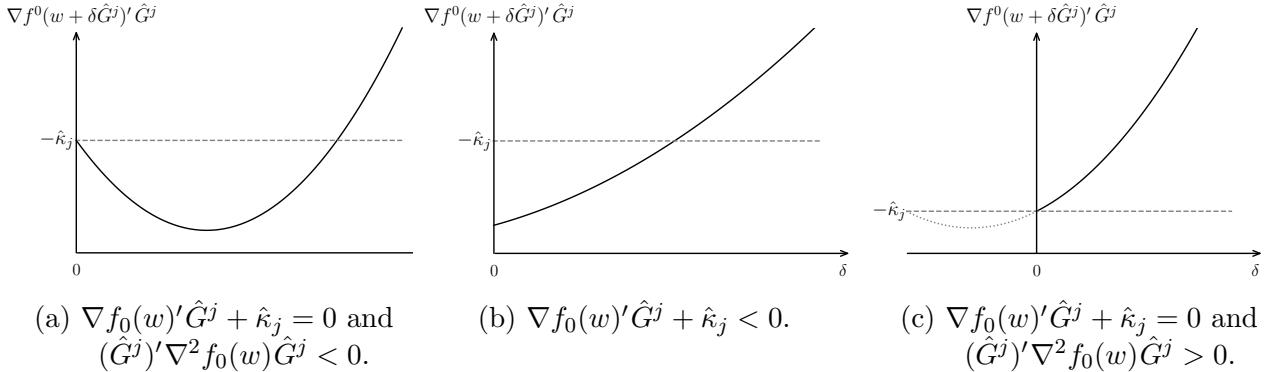


Figure 2.10: An illustration of the three instances in which the control policy can be improved.

(with index j replaced by l), but also sets the second directional derivative to zero. Recall that the value function can be extended linearly outside the region of inaction, thus, having a second derivative of zero there. Consequently, the choice of δ^* is also motivated by the principle of smooth pasting.

Case 2. The value function is locally linear along the pushing direction \hat{G}^l , i.e., $(\hat{G}^l)' \nabla^2 f_0(w) \hat{G}^l = 0$. In this case, we do not move the boundary point w . However, we change the pushing control if there exists another feasible pushing direction \hat{G}^j such that the directional derivative of f_0 along \hat{G}^j is smaller than $-\hat{\kappa}_j$ in an interval $\{w + z\hat{G}^j : z \in (0, \delta)\}$, which can happen in two ways as illustrated in Figures 2.10a and 2.10b. First, there exists a feasible pushing control j such that $\nabla f_0(w)' \hat{G}^j + \hat{\kappa}_j < 0$ (see Figure 2.10b). Second, for all feasible pushing controls i , we have $\nabla f_0(w)' \hat{G}^i + \hat{\kappa}_i \geq 0$. However, there exists a feasible pushing control j such that $\nabla f_0(w)' \hat{G}^j + \hat{\kappa}_j = 0$, but $(\hat{G}^j)' \nabla^2 f_0(w) \hat{G}^j < 0$, i.e., f_0 is locally concave along the pushing direction \hat{G}^j (see Figure 2.10a). In both cases, the directional derivative of the value function along the pushing direction \hat{G}^j is strictly smaller than $-\hat{\kappa}_j$ in the vicinity of the boundary point w and we can improve the performance by instantaneously moving the process from w inward. Motivated by this observation, we keep the boundary point the same, but change the pushing control to j^* . Namely, in the first

case, we choose the pushing control

$$j^* = \operatorname{argmin} \{ \nabla f_0(w)' \hat{G}^j + \hat{\kappa}_j : j \text{ is a feasible pushing control at } w \}.$$

In the second case, we choose the pushing control

$$j^* = \operatorname{argmin} \{ (\hat{G}^j)' \nabla^2 f_0(w) \hat{G}^j : j \text{ is a feasible pushing control at } w \text{ with } \nabla f_0(w)' \hat{G}^j + \hat{\kappa}_j = 0 \}.$$

In both cases, tie are broken in favor of smaller indices.

Case 3. The value function is locally convex along the the pushing direction \hat{G}^l , i.e., $(\hat{G}^l)' \nabla^2 f_0(w) \hat{G}^l > 0$. The directional derivative of f_0 along the pushing direction \hat{G}^l is shown in Figure 2.10c (with index j replaced by l). We consider two subcases in general: (a) w can be moved outwards, i.e., $w \in \text{int}(\hat{\mathcal{W}})$; and (b) w cannot be moved outwards, i.e., $w \in \partial \hat{\mathcal{W}}$. In case (a), we move the boundary point w outward along the pushing direction \hat{G}^l by δ_0 , where δ_0 is a tunable parameter, or until the boundary of $\hat{\mathcal{W}}$ is reached. In case (b), we cannot move the boundary outward. Therefore, we consider changing the pushing control. To do so, we proceed as in Case 2 described above. Note that at the initiation step, only case 3(b) is relevant if $\mathcal{B}^0 = \hat{\mathcal{W}}$. However, in future iterations or with a different region of inaction, case 3(a) can be relevant as well. Next, we discuss the convexification of the updated region of inaction. This step is not necessary for the NYC application, but it is included for completeness.

Convexification. Since the boundary points may move in the update procedure, the updated region of inaction may not be convex. Therefore, we convexify the updated region of inaction by taking the convex hull of the updated boundary points. Then, we use the pushing controls used on the updated (before convexification) boundary points to assign a pushing control to each point on the boundary of the convexified region of inaction. We do so by projecting each point on the boundary of the convexified region of inaction to the set

of updated (before convexification) boundary points. To be specific, letting \mathcal{U} denote the set of updated (before convexification) boundary points and \mathcal{B}^1 denote the convexified region of inaction, to $w \in \partial\mathcal{B}^1$, we assign the pushing control used at $\operatorname{argmin}\{\|w - w'\| : w' \in \mathcal{U}\}$, where $\|\cdot\|$ denotes the Euclidean norm. Namely, for each boundary point w of the convexified region of inaction \mathcal{B}^1 , we first find the closest point $w' \in \mathcal{U}$ to w , and assign the pushing control at w' to w .

We denote the partition of $\partial\mathcal{B}^1$ by $\partial\mathcal{B}^{1,1}, \dots, \partial\mathcal{B}^{1,d}$, where pushing control l is used for $w \in \mathcal{B}^{1,l}$ for $l = 1, \dots, d$; see Figure 2.8b for the updated region of inaction \mathcal{B}^1 and the partition of its boundary for the NYC application.

Step $k \geq 1$. For expositional clarity, we first discuss step 1. We use the new region of inaction \mathcal{B}^1 and the partition of its boundary $\partial\mathcal{B}^{1,1}, \dots, \partial\mathcal{B}^{1,d}$ to define the region-of-inaction type control policy \hat{U}^1 as done in step 0. Then, we solve a linear elliptic PDE similar to the one in (2.60)-(2.61) to compute the value function associated with \hat{U}^1 and proceed in the same way as described above. Step k for $k > 1$ is identical to step 1.

Termination. The algorithm terminates once $\partial\mathcal{B}^{k,1}, \dots, \partial\mathcal{B}^{k,d}$ are sufficiently close to $\partial\mathcal{B}^{k+1,1}, \dots, \partial\mathcal{B}^{k+1,d}$ for some k . More specifically, we terminate the algorithm when the set of mesh nodes (points) that belong to $\partial\mathcal{B}^{k+1,l}$ is identical to the set of mesh nodes that belong to $\partial\mathcal{B}^{k+1,l}$ for $l = 1, \dots, d$. For the NYC application, given the initialization depicted in Figure 2.8a and a discretization step size of 0.01, the algorithm terminates after step 3. Of course, with a different initialization, it could take longer. In all examples considered, the algorithm converges in fewer than four steps. The solution depicted in Figure 2.8c outperforms the solution derived by the Markov chain approximation method; see Appendix B.1.8 for details. For the NYC application, the control policy obtained in each step is an improvement over the previous control policy. Lemma 17 in Appendix B.1.9 establishes this property under certain assumptions.

2.8 Proposed Policy

This section proposes a dynamic control policy for the problem introduced in Section 2.3 by interpreting the solution to the EWF (2.47)-(2.50) in the context of the original control problem. Recall that we consider a sequence of systems indexed by the system parameter r , whose formal limit is the BCP (2.32). The original control problem introduced in Section 2.3 can be viewed as a specific element of this sequence of problems, which is determined by the particular choice of r . Therefore, first, we determine the system parameter r . In our case, by Equation (2.12), we have $r = (\Lambda^r)^2$, where Λ^r is the total number of cars; see Section 2.3. The underlying assumption of the Brownian approximation is that the system parameter r is large enough that the various (scaled) performance-relevant processes of the original control problem can be approximated by the corresponding processes of the Brownian control problem. Recall that the scaled workload process of the EWF (2.47)-(2.50) can be expressed in terms of the queue length process of the original control problem (formulated in Section 2.3) as follows: $\hat{W}(t) = LMQ(t)/\sqrt{r}$ for $t \geq 0$.

It has long been observed in the heavy traffic literature that the idleness can be avoided for all servers serving a buffer pool that has positive workload. As articulated in Harrison [111], this behavior can be achieved by keeping positive but small queues in each buffer, e.g., safety stocks, for further details see Harrison [111, 112], Harrison and López [116], Maglaras [160], and Ata and Kumar [27]. Therefore, as is customary in the heavy traffic literature, we put small safety stocks in various buffers. The safety stock for buffer i is denoted by s_i for $i = 1, \dots, q$; the values s_i can be calibrated via simulation as done in Wein [213].

To complete the policy description, we consider two cases: $\hat{W}(t) \in \text{int}(\mathcal{B})$ and $\hat{W}(t) \notin \text{int}(\mathcal{B})$. Recall that in the solution to the equivalent workload formulation, no action is taken when the workload process is in the interior of the region of inaction. In the context of the original problem formulation introduced in Section 2.3, this can be interpreted as undertaking the dispatch and relocation activities at the rates dictated by the solution to the static planning problem, (x^*, \tilde{x}^*) , while avoiding server idleness, when $\hat{W}(t) \in \text{int}(\mathcal{B})$.

In the solution to the equivalent workload formulation, when the workload process reaches the boundary of the region of inaction, pushing controls are exerted. Next, we elaborate on these two cases further in the context of the problem introduced in Section 2.3:

Case 1. $\hat{W}(t) \in \text{int}(\mathcal{B})$. For $k = 1, \dots, m$, when server k becomes available, it chooses a buffer to serve as follows: First, it only considers those buffers that are above their safety stock thresholds and that it can serve using basic dispatch activities only. Then, it randomizes among them with probabilities proportional to the nominal activity rates given by the static planning problem. That is, the server engages in one of dispatch activities $\{j : k(j) = k, Q_{i(j)}(t) \geq s_i, j = 1, \dots, b\}$ with probabilities proportional to x_j^* . If no such activity exists, then the server engages in one of the basic dispatch activities available to it that serve a buffer that is non-empty currently, i.e., $\{j : k(j) = k, Q_{i(j)}(t) > 0, j = 1, \dots, b\}$, with probabilities proportional to x_j^* . If no such activity exists either, the server idles. Moreover, the platform undertakes basic relocation activities $j = 1, \dots, \tilde{b}$ at rates \tilde{x}_j^* for $j = 1, \dots, \tilde{b}$ as long as the number of cars at the origin buffer is greater than or equal to the safety stock, i.e., $Q_{o(j)}(t) \geq s_{o(j)}$.

Case 2. $\hat{W}(t) \notin \text{int}(\mathcal{B})$. In this case, the platform modifies the policy described in Case 1 by using the pushing controls prescribed by the solution to the EWF (2.47)-(2.50). Namely, at each boundary point, the platform modifies the policy described in Case 1 by undertaking the activity corresponding to the pushing control prescribed by the solution to the EWF (2.47)-(2.50). Each such activity may be in addition to the basic activities or it may replace a basic activity as discussed next. For example, if the solution prescribes idling server k , server k idles instead of undertaking the basic dispatch activity prescribed by the policy described in Case 1. If the solution prescribes taking part in non-basic dispatch activity j , server $k(j)$ serves buffer $i(j)$ as long as there are jobs in buffer $i(j)$ to serve, i.e., $Q_{i(j)}(t) > 0$, instead of undertaking the basic dispatch activity assigned to it by the policy described in Case 1. If the solution prescribes taking part in non-basic relocation activity j , the platform relocates jobs from buffer $o(j)$ to buffer $d(j)$ as long as there are jobs available at the origin

Table 2.1: Pushing controls used in the solution (derived by our computational method) to the EWF (2.47)-(2.50) for the NYC application.

State	Pushing Control	
	Number	Action
$\{w : w_1 = 0, 0 < w_2 < 1\}$	1	Idle Server 5
$\{w : w_2 = 0, w_1 < 0.35\}$	5	Serve Buffer 7 with Server 6
$\{w : w_2 = 0, w_1 \geq 0.35\}$	2	Idle Server 6
$\{w : w_1 + w_2 = 1, w_1 < 0.95\}$	4	Serve Buffer 6 with Server 7
$\{w : w_1 + w_2 = 1, w_1 \geq 0.95, w_2 > 0\}$	3	Idle Server 7

to relocate, i.e., $Q_{o(j)}(t) > 0$, in addition to the policy described in Case 1. These will be illustrated further for the NYC application in Section 2.9.

The solution to the EWF (2.47)-(2.50) for the NYC application is depicted in Figure 2.8c. In this solution, $\mathcal{B} = \hat{\mathcal{W}}$. Since the solution to the EWF (2.47)-(2.50) exerts no control when $\hat{W}(t) \in \text{int}(\hat{\mathcal{W}})$, the optimal solution to the static planning problem (2.14)-(2.17) is implemented as discussed in Case 1 above. Similarly, since the solution to the EWF (2.47)-(2.50) exerts pushing controls when $\hat{W}(t) \in \partial\hat{\mathcal{W}}$, pushing controls are used as discussed in Case 2. Table 2.1 lists the pushing controls used on each part of $\partial\hat{\mathcal{W}}$. For example, on the horizontal face of $\partial\hat{\mathcal{W}}$, which corresponds to zero workload in pool 2, the following pushing controls are exerted in the solution to the EWF (2.47)-(2.50): If the scaled workload of pool 1 is less than 0.35, pushing control 5 is exerted. Otherwise, pushing control 2 is exerted. This corresponds to server 6 serving buffer 7 in the first case and server 6 idling in the second case.

2.9 Simulation Study and Discussion

This section undertakes a simulation study to illustrate the effectiveness of the proposed control policy. The simulation study focuses on the NYC application introduced in Section 2.1. In doing so, we crucially relax the assumption of zero travel times. Instead, we use the travel times estimated from the NYC data mentioned in Section 2.1. We compare the proposed policy with nine other policies: the static policy, the closest-driver policy and its

modified version, the MaxWeight policy and its modified version, the penalized MaxWeight policy and its modified version, and the mirror backpressure policy and its modified version. The static policy implements the optimal solution to the static planning problem (2.14)-(2.17); see Özkan and Ward [172] for a similar policy. The closest-driver policy is a natural policy that dispatches a car from the closest area from which cars can be dispatched to the customer (i.e., the dispatch activity with the highest value is undertaken).¹⁹ The closest-driver policy is commonly used as a benchmark in the ride-hailing literature; see e.g., Özkan and Ward [172]. The MaxWeight policy, also known as maximum pressure or backpressure policy, dispatches a car from the area with the largest number of empty cars among the areas from which cars can be dispatched to the customer. This is a popular policy in the queueing literature; see e.g., Dai and Lin [83] and Ata and Lin [29]. Banerjee et al. [40] shows that the MaxWeight policy has an exponentially decaying demand dropping probability under the complete resource pooling assumption, i.e., $p = 1$. The penalized MaxWeight policy dispatches a car from the area with the largest number of empty cars minus a linear travel cost among the areas from which cars can be dispatched to the customer. This policy is proposed as an improvement over the original MaxWeight policy in Banerjee et al. [40]. The mirror backpressure policy dispatches a car from the area where the immediate payoff of serving the customer plus the difference of the congestion costs at the car's current location and the customer's destination is maximized. The mirror backpressure policy differs from our proposed policy as well as the other policies considered in terms of the information structure, since it requires knowledge of the destinations of the customers. Kanoria and Qian [136] shows that the mirror backpressure policy is a near-optimal dispatch policy.

The demand flow in the NYC application is not symmetric, i.e., the rate at which customers depart an area is not equal to the rate at which they enter the area. Consequently, under the closest-driver, MaxWeight, penalized MaxWeight, and mirror backpressure poli-

19. As stated in Section 2.1, in the NYC application, the platform incurs a dispatch cost of \$0.2 per mile. Consequently, the platform generates a larger (immediate) value by dispatching a car from the closest areas.

cies, empty cars tend to pile up in certain areas. To alleviate this issue, we also consider suitably modified versions of the closest-driver, MaxWeight, penalized MaxWeight, and mirror backpressure policies. In the modified policies, the closest-driver, the MaxWeight, penalized MaxWeight, and the mirror backpressure policies are enhanced by incorporating (static) relocation decisions adopted from the static planning problem as explained below; see Appendix B.1.10 for a detailed description of the nine policies considered.

As mentioned above, we relax the assumption that cars move between different areas instantaneously and incorporate travel times. This allows us to test the performance of the proposed policy in a realistic setting. In relaxing this assumption, we include the travel time to pick up the customer, the travel time to deliver the customer, and the travel time to relocate from one area to another in the simulation. The number of (empty) cars at various areas in the simulation study can differ from those in the model studied above because of the travel times. To map the proposed policy to an implementable policy in the setting with travel times, we use $\hat{W}(t) = LMQ(t) / \sum_{i=1}^9 Q_i(t)$ for $t \geq 0$; $\hat{W}_l(t)$ denotes the fraction of empty cars in pool l at time t for $l = 1, 2$ and $t \geq 0$.

In the simulation study, the annual interest rate is $\alpha = 0.05$. Moreover, as done by Wein [213], for simplicity, we focus on safety stock parameters of the form $s_i = s_1$ for $i = 2, \dots, 9$. That is, we choose identical safety stocks for all buffers. We assume that all cars are initially empty and waiting for customers to arrive. We also assume that the initial number of cars in each area is proportional to the demand in that area.²⁰ As a preliminary to reviewing the simulation results, let Λ_0 denote the average number of traveling cars just to sustain the basic activity rates (x^*, \tilde{x}^*) derived from the static planning problem. To be specific, letting

20. The assumption that all cars are initially empty and waiting for customers to arrive has a negligible impact on our results as the ride-hailing system reaches steady-state in less than an hour. Note that we simulate the ride-hailing system under each policy 100 times for 300 years. Preliminary checks show that costs incurred beyond this point account for less than 0.0001% of the discounted cost.

τ_{il} denote the average travel time from area i to area l for $i, l = 1, \dots, 9$, we have that

$$\Lambda_0 = \sum_{j=1}^{19} \mu_{k(j)} x_j^* \tau_{i(j)k(j)} + \sum_{j=1}^{19} \sum_{l=1}^9 \mu_{k(j)} x_j^* P_{jl} \tau_{k(j)l} + \sum_{j=1}^{72} \tilde{x}_j^* \tau_{o(j)d(j)}, \quad (2.65)$$

where the first term on the right-hand side is the average number of cars en route to pick up customers, the second term is the average number of cars en route to deliver customers, and the third term is the average number of relocating cars (under the static policy if the entire demand is satisfied).²¹ For the NYC application, $\Lambda_0 = 8200$.²² To assess the impact of the number of cars Λ on the system performance, we vary Λ from $0.5\Lambda_0 = 4100$ cars to $1.5\Lambda_0 = 12300$ cars, corresponding to a $\pm 50\%$ range around Λ_0 . This wide range of values for Λ helps us thoroughly compare the performance of the different policies considered. Hereafter, we refer to $\Lambda - \Lambda_0$ as the (percentage) excess supply. To compare the performance of the aforementioned policies, we report the expected discounted cost of each policy in Figure 2.11 below as well as in Tables B.4-B.5 in Appendix B.1.11. To be more specific, Figure 2.11 and Tables B.4-B.5 report the unscaled expected discounted costs ΛJ_z for each policy; see Equation (2.30) for the definition of J_z as well as Equations (2.11), (2.21), (2.30), and (2.31) for further details on Equation (2.30). Moreover, Tables B.6-B.7 in Appendix B.1.11 provide a comparison of the scaled expected discounted costs, i.e. J_z or the expected discounted cost per car. Tables B.8-B.9 in Appendix B.1.11 report the percentage difference of the various policies considered from the proposed policy. We also report the demand dropping probabilities in Figure 2.12 below as well as in Tables B.10-B.11 in Appendix B.1.11. For brevity, among possible aforementioned mirror backpressure policies corresponding to different congestion functions, we report the simulation results only for the best performing

21. Equation (2.65) follows from Little's law. Let us illustrate this for the first term on its right-hand side. Note that $\mu_{k(j)} x_j^*$ is the rate at which customers are picked up under basic dispatch activity j for $j = 1, \dots, 13$. Therefore, $\mu_{k(j)} x_j^* \tau_{i(j)k(j)}$ is the average number of cars en route to pick up customers under basic dispatch activity j for $j = 1, \dots, 13$. Summing over j for $j = 1, \dots, 13$ gives the average number of cars en route to pick up customers, which is the first term on the right-hand side of Equation (2.65).

22. From the 8200 traveling cars, about 4800 are en route to deliver customers, 3300 are en route to pick up customers, and 100 are in the relocation process.

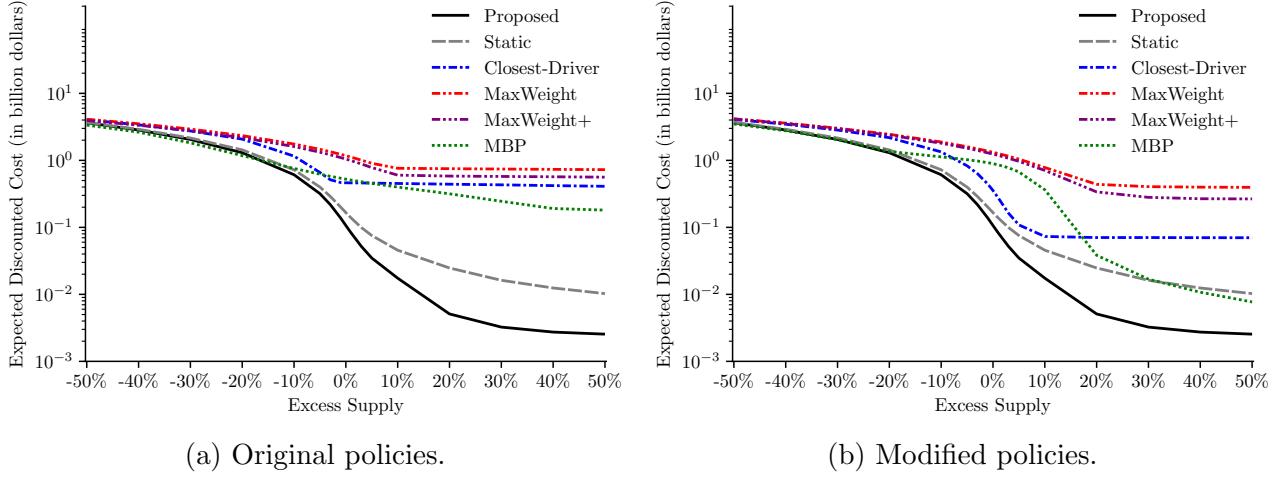


Figure 2.11: The expected discounted cost under various policies.

one, which uses the logarithmic congestion function. Moreover, for visual clarity, we refer to the penalized MaxWeight policy as MaxWeight+ in the figures and tables.

Figures 2.11-2.12 show that the performance of all policies improves as the total number of cars increases. The proposed policy outperforms the static policy as well as the original and the modified closest-driver, MaxWeight, and penalized MaxWeight policies both in terms of the cost and the demand dropping probability. This demonstrates the effectiveness of the proposed policy. Interestingly, the static policy has a lower expected discounted cost than the following policies: the original and the modified closest-driver policies, the original and the modified MaxWeight policies, and the original and the modified penalized MaxWeight policies. The static policy also has a lower demand dropping probability than all the policies mentioned above except for the modified MaxWeight policy. When excess supply is large, i.e., $\Lambda = 1.4\Lambda_0$, 0.16% of the demand is dropped under the modified MaxWeight policy, whereas 0.11% of the demand is dropped under the static policy.

The proposed policy is slightly outperformed by both the original and the modified mirror backpressure policies when supply is too limited, e.g., the proposed policy is outperformed by the original mirror backpressure policy by up to 11% when the excess supply is less than -10% . In this range of supply, demand significantly exceeds supply and the mirror backpressure policies can not only drop the least profitable customers, but also avoid the

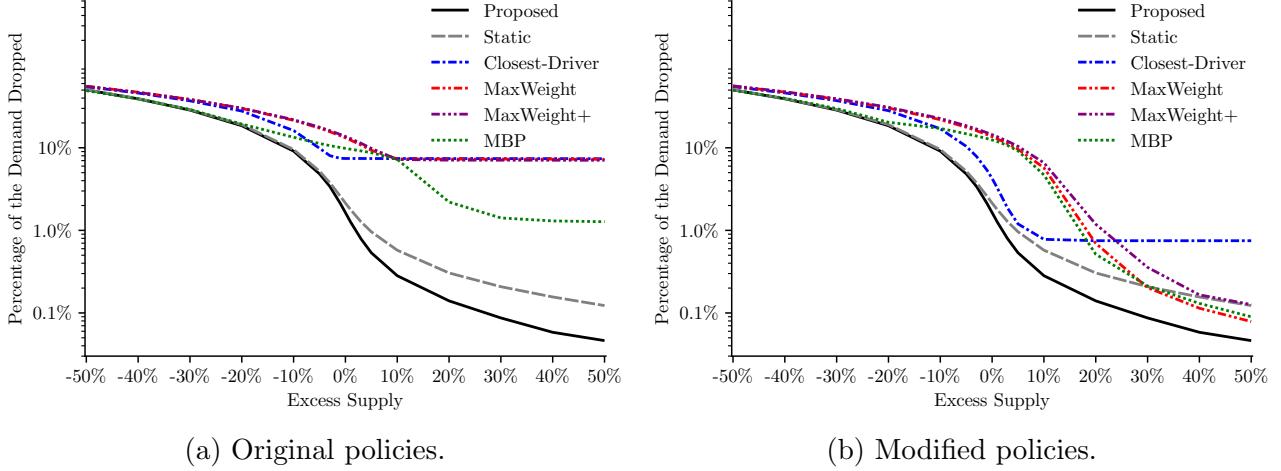


Figure 2.12: Percentage of the demand dropped under various policies.

accumulation of empty cars in an area, because they use the customer's destination information. The proposed policy outperforms the original and the modified mirror backpressure policies otherwise, i.e. when the excess supply is greater than or equal to -10% . This demonstrates that the proposed policy performs well in the regime for which it is designed. Namely, when the system is critically loaded, i.e., it is in heavy traffic, but it is not overloaded. Moreover, the fact that our proposed policy (which uses the demand distribution and not the destination information) outperforms the mirror backpressure policy (which uses the destination information and not the demand distribution) suggests that knowing the demand distribution may be more valuable to ride-hailing platforms than the destination information, although they are both valuable.

The original closest-driver, MaxWeight, and penalized MaxWeight policies, which focus only on dispatch control, cannot meet the entire demand even when there is an abundant supply of cars. This is due to their inability to address the asymmetry of the demand flow. Although the original mirror backpressure policy performs better than the original closest-driver, MaxWeight, and penalized MaxWeight policies, it also has the same shortcoming. The modified closest-driver, the modified MaxWeight, the modified penalized MaxWeight, and the modified mirror backpressure policies, which take advantage of the static relocation policy to offset the asymmetric flow of demand, alleviate this concern especially when there

is sufficient supply. We also observe that the fraction of the demand dropped appears to decay faster under the modified MaxWeight policy, the modified penalized MaxWeight policy, and the mirror backpressure policies in comparison to the closest-driver and static control policies.²³ Lastly, we observe that our proposed dynamic control policy outperforms all other policies considered in terms of the fraction of the demand dropped.²⁴

As reported in Chapter 1, on an average day during the period our NYC application focuses on, there were approximately 11200 active taxis in NYC. However, there is a major difference between our setting and the taxi market in NYC. The NYC taxi and limousine commission did not centrally control the taxi market during the period on which our NYC application focuses. During this period, taxis had to search for customers street by street. As shown in Chapter 1, this street-by-street search for customers introduced substantial (search) friction. In contrast, our paper focuses on a platform that centrally controls the market. Given this major difference, the reader should not associate the 11200-car case of the NYC taxi market with the +40% excess supply case in our simulation study. As shown in Chapter 1, about 4.3% of the demand was dropped during the period on which our NYC application focuses. Our proposed control policy, however, achieves a 4.3% demand dropping probability with about 7950 cars. This corresponds to 3250 fewer taxis than the NYC taxi and limousine commission was using. This highlights the value of centrally controlling the market or the cost of lack-of-coordination.

23. Banerjee et al. [40] shows that the original MaxWeight policy has an exponentially decaying demand dropping probability under the complete resource pooling assumption, i.e., $p = 1$. Our findings for the modified MaxWeight policy shown in Table B.11 appear to support its finding beyond the complete resource pooling case, i.e., for $p > 1$.

24. Note that the safety stock in our proposed dynamic control policy as well as the tuning parameter in the penalized MaxWeight policies and the mirror backpressure policies are calibrated such that the expected discounted cost is minimized.

2.10 Concluding Remarks

This chapter studies the control of a ride-hailing platform that offers on-demand transportation to heterogeneous customers by dynamically making dispatch and centralized relocation decisions. An offline linear program, called the static planning problem, is solved. The optimal solution to the static planning problem, which can be interpreted as the optimal static control policy, partitions the set of activities into two: basic or efficient activities and non-basic or inefficient activities. The basic activities are used under the optimal solution to the static planning problem, whereas the non-basic activities are not. The platform partitions the areas of the city into pools of areas and only uses the information on the fraction of cars in the various pools as opposed to the fraction of cars in each area. This significantly reduces the state space dimension, making the solution of the associated HJB equations computationally feasible.

We propose a dynamic control policy that involves a so-called region of inaction, a convex set of states, where the state of the system is defined as the fraction of cars in the various pools. When the system state is in the interior of the region of inaction, the policy uses the optimal solution to the static planning problem, i.e., it uses the basic or efficient dispatch and relocation activities. When the system state reaches the boundary of the region of inaction, the platform exerts additional controls by dropping demand or undertaking a non-basic dispatch or relocation activity to keep the system state in the region of inaction. It is demonstrated through a simulation study that the proposed policy performs well.

CHAPTER 3

DYNAMIC PRICING OF A MULTICLASS MAKE-TO-STOCK QUEUE

3.1 Introduction

This chapter studies the dynamic control of a make-to-stock manufacturing system that sells multiple products to price-sensitive customers. In make-to-stock manufacturing systems, the products are manufactured based on a forecast of the demand and stored in finished goods inventory. Customers are dynamically quoted a price. When a customer order is received, it is served from the finished goods inventory if the item requested is in stock. Otherwise, it is backordered. The manufacturer incurs a holding cost for each product in inventory and a backorder cost for each request on backorder (per unit of time). The make-to-stock manufacturing system is modeled as a multi-class single-server queueing system. The system manager seeks to maximize the long-run average profit by making dynamic pricing, outsourcing, and scheduling decisions: First, she controls the demand by dynamically adjusting the prices. Second, when the backlog of work is judged excessive, she outsources (or rejects) customer orders. Finally, she dynamically decides when and what products to manufacture. The joint consideration of pricing and operational decisions, such as scheduling and outsourcing, leads to a manufacturing system that is more responsive to the variability in the demand and production processes. Moreover, dynamic prices provide the system manager with the opportunity to take advantage of the heterogeneity of customers in their price sensitivity to increase profits. In determining the prices, she takes into account the price-sensitivity of demand, the holding, backlog, and outsourcing costs, the variability in the demand and production processes, and the production capacity of the manufacturing system.

Because the joint dynamic pricing, outsourcing, and scheduling control problem appears intractable in its exact form, we follow Harrison [110] and approximate it by a far more

tractable formulation referred to as the Brownian control problem. We do so in the so-called heavy traffic asymptotic regime which assumes both the demand and the system capacity are large, and the average server utilization is near one. We derive an effective and intuitive joint dynamic pricing, outsourcing, and scheduling policy for the manufacturing system by solving the approximating Brownian control problem explicitly and interpreting its solution in the context of the original control problem.

The approximating Brownian control problem is equivalent to a one-dimensional drift rate control problem, whose state process, called the workload process, represents the total (scaled) inventory (or backlog) in the system measured in hours of total work for the server. An important feature of this drift rate control problem is the presence of state costs that capture the holding and backorder costs in the original control problem. The drift rate control problem is solved explicitly by exploiting the solution to a Riccati equation. The optimal policy comprises of the drift rate control policy and a two-sided barrier policy. The outsourcing and idling processes are used to keep the workload process between the lower and upper reflecting barriers. Between the two barriers, a state-dependent drift rate is used to control the workload process, which is ultimately interpreted as a dynamic pricing policy.

The optimal solution to the approximating Brownian control problem leads to an intuitively appealing and easily implementable policy and reveals several structural insights. First, the system manager's pricing, outsourcing, and sequencing decisions depend primarily on the (aggregate) workload. In particular, the prescribed prices tend to vary on a slower time-scale than the demand and production processes. Second, in large systems, the magnitude of the prescribed price changes are small. Third, as the backlog of work increases and the inventory of products decreases, prices should be adjusted to decrease the (effective) demand rate. Finally, only when the backlog of work is too excessive, outsourcing is worthwhile.

The rest of the chapter is organized as follows. Section 3.2 reviews the relevant literature. Section 3.3 introduces the control problem for a make-to-stock manufacturing system.

Section 3.4 formally derives the approximating Brownian control problem. Section 3.5 develops the one-dimensional equivalent workload formulation. Section 3.6 derives an optimal solution to the equivalent workload formulation by solving the associated Bellman equation. Section 3.7 interprets the solution to the equivalent workload formulation in the context of the control problem laid out in Section 3.3, proposes a policy in that context, and provides a simulation study to show the effectiveness of the proposed policy. Section 3.8 concludes the chapter.

3.2 Literature Review

This chapter is related to three streams of literature. The first stream studies the scheduling of make-to-stock manufacturing systems. The second stream studies the control of queueing systems. The third stream studies the revenue management of manufacturing and queueing systems. The literature on the scheduling of manufacturing systems is vast; see Nahmias and Olsen [168, Chapter 9] for a brief overview. To the best of our knowledge, the scheduling of a multi-class make-to-stock manufacturing system was first studied by Zheng and Zipkin [224]. Zheng and Zipkin [224] showed that in the case of two symmetric products with exponential production times, the policy that serves the longest queue outperforms the First-Come-First-Served (FCFS) policy. Zipkin [225] extended the results to the case with multiple products and general production times. Another early paper in this area and an important antecedent of our work is Wein [213], that studies the dynamic scheduling of a multiclass single-machine make-to-stock manufacturing system in heavy traffic. Wein [213] proposes a scheduling policy to minimize the long-run average holding and backorder costs. We build on Wein [213] by introducing dynamic pricing and outsourcing. These extra features lead to a different workload formulation, where the analytical solution of the associated Bellman equation dictates both a specific dynamic pricing policy and a novel discretionary outsourcing threshold in addition to the scheduling policy that is similar to Wein [213]'s. Ha [104] provides theoretical justification for some of the ideas proposed in Wein [213]. In

particular, it shows that in a two-product system, it is optimal to produce the product with the larger $b\mu$ index if it backordered (independent of the inventory of the other product).

Typically, Brownian models fall short of providing complete scheduling prescriptions under linear state costs. Rather they would reveal the cheapest buffers to keep the inventory and backlog. The prioritization of the other buffers are done through safety stocks, which are in turn determined by ad hoc methods such as simulation studies. Veatch and Wein [202] seeks complete scheduling prescriptions that are effective and computationally tractable. The authors limit their attention to index policies, where an index is dynamically calculated for each product and the product with the smallest index is manufactured. Rubio and Wein [182] considers a multi-product/multi-server model of make-to-stock manufacturing system using a Jackson queueing network model. They focus on base-stock policies and show that the optimal base-stock level of each product corresponds to a critical fractile of the steady state distribution of the product's total work-in-progress inventory. Then, using results from product-form queueing networks, the authors derive closed-form expressions for the optimal base-stock levels.

Véricourt et al. [203] focuses on the optimal scheduling of a two-product make-to-stock manufacturing system. It characterizes the hedging point for a particular region of the state space and shows that in this region, the monotone switching curve, that separates the priority regions of the two products, is a straight line. Veatch and Véricourt [201] studies the scheduling of a two-part-type single-product make-to-stock queue using the model of Wein [213]. For the case with exponential inter-arrival and production times, it identifies the necessary and sufficient conditions for a zero-inventory policy to be optimal. Sanajian et al. [184] is another relevant paper in this literature that compares the performance of the static FCFS, preemptive, and non-preemptive priority policies for a make-to-stock system with general production times. It provides conditions under which it is optimal to produce according to a make-to-order policy. Kim and Van Oyen [138] studies the scheduling of a make-to-stock manufacturing system with setup and switching costs. Karaesmen et al. [137]

studies a system with advance order information.

The second stream of literature studies the control of queueing systems; see Stidham [190, 192] for surveys. Our work is closely related to the literature on the dynamic control of queueing systems in heavy traffic, pioneered by Harrison [110, 113, 114]. In this stream of literature, one approximates the queueing system with a Brownian system, which is simpler to analyze. Early examples of this approach include Harrison and Wein [121], that studies an optimal sequencing problem for a criss-cross network, and Harrison and Wein [122], that studies a multiclass two-station closed queueing network. In both cases, the limiting Brownian models admit pathwise optimal solutions that have straightforward interpretations in the original problem. Harrison and Wein show that their policies perform well. Since then many other researchers followed the heavy traffic approach to study manufacturing and queueing systems; see for example Wein [211, 212, 214], Chevalier and Wein [72], Krichagina and Wein [142], Reiman and Wein [178], Kumar [144], Bell and Williams [44], Plambeck et al. [175], Markowitz and Wein [163], Maglaras and Zeevi [161, 162], Plambeck [176], Ward and Kumar [208], Reed and Ward [177], Ata and Olsen [30], Tezcan and Dai [198], Dai and Tezcan [80], and Ward and Armony [207].

The heavy traffic approximations often give rise to singular control problems. Harrison and Taylor [119] considers the problem of minimizing the discounted cost of control plus linear holding costs while ensuring that inventory stays non-negative. First, the authors consider a singular control formulation and show that a single-barrier policy is optimal in this case. The authors also consider a variation where the control is a step function and a fixed cost is incurred each time the control is exercised. Harrison and Taylor show that a control band policy is optimal in this case. Harrison et al. [117] consider the impulse control problem where the system manager can increase (or decrease) the inventory by any amount by incurring a fixed plus a proportional cost for displacement. The authors show that a three-parameter control band policy is optimal. Harrison and Taksar [118] considers a singular control problem that seeks to minimize discounted costs of control that are linear

and convex holding costs. The authors establish the optimality of a two-sided barrier policy. Taksar [195] extends the results to the average cost formulation. Dai and Yao [81, 82] extend the literature by considering a general impulse (and singular) control problem under the average and discounted cost objectives, respectively. They show that the impulse (singular) control problem admits a four-parameter control band (two-sided barrier) policy. In a related problem, Weerasinghe [210] considers a one-dimensional Brownian motion plus a process of bounded variation, where the controller can control the drift and variance of the Brownian motion as well as a bounded variation pushing process. The objective is to minimize the long-run average holding cost plus linear cost of using the bounded variation process. It is shown that the optimal solution results in a reflected Brownian motion on a bounded interval.

Control of queueing systems often boils down to admission or service rate control, which have been tackled either by heavy traffic approximations or Markov decision processes in the literature. The heavy traffic approximations often result in drift rate control problems for Brownian models, which we review next. Krichagina and Taksar [141] and Krichagina et al. [140] study service rate control problems and establish the optimality of threshold-type policies. An important antecedent of our work is Ata et al. [26], that studies a (one-dimensional) drift rate control problem of a reflected Brownian motion on a bounded interval. The formulation of Ata et al. [26] is motivated by a service rate control problem in wireless communication. Under the average cost criterion, the authors derive a closed-form solution. Although Ata et al. [26] is closely related to our work methodologically, it differs in important ways, which render its solution technique incapable in our setting: First, its solution technique hinges on the absence of state costs, whereas our formulation crucially has state costs, i.e., holding and backlog costs. Second, Ata et al. [26] considers the drift rate control problem on a bounded interval which simplifies the solution of the Bellman equation considerably. In contrast, we consider the problem on the entire real line and part of our analysis involves choosing the discretionary lower and upper reflecting barriers using the smooth pasting

technique. Third, our problem requires combining the solutions on the negative and positive parts of the state space in a smooth fashion while the problem studied in Ata et al. [26] does not involve such intricacies. Ata [19] considers an admission control problem for a multiclass single-server queueing system as a model of make-to-order production system with due date leadtimes and builds on Ata et al. [26] methodologically. Its novel feature is that all but one class are "thin". Consequently, the admission control problem can be approximated by a drift-rate control problem in the heavy traffic asymptotic regime. Solving this problem yields a nested threshold policy as the optimal admission control policy. Another methodologically relevant paper is Ghosh and Weerasinghe [99], that extends Ata et al. [26] by incorporating holding costs and allowing the system manager to choose the upper barrier endogenously. The authors establish that an optimal solution exists and that it can be characterized by solving the Bellman equation. We build on Ghosh and Weerasinghe [99] by considering a drift rate control problem on the entire real line and derive an explicit solution by solving the Bellman equation in closed-form.

Ghosh and Weerasinghe [100] extends the earlier work Ghosh and Weerasinghe [99] by introducing abandonments. Budhiraja et al. [57] studies the joint admission and service rate control in a queueing network and constructs asymptotically optimal control policies. Ata and Tongarlak [34] studies the dynamic scheduling in a multiclass queueing system with abandonments and general costs of delay. It considers convex, concave, and convex-concave costs of delay. Its solution to the Brownian control problem yields a non-greedy dynamic index policy as the optimal scheduling policy. Rubino and Ata [181] and Ghamami and Ward [98] both study a dynamic scheduling problem for a parallel-server system with abandonments and derive the optimal policy by solving the Bellman equation for the approximating Brownian control problem. Ghamami and Ward [98] also shows that its proposed policy is asymptotically optimal for the original system. Cao and Yao [63] studies the joint drift rate and impulse control for a one-dimensional diffusion process subject to the condition that the diffusion process must remain non-negative. It shows that the optimal solution

is a three-parameter control band policy with drift control. Weerasinghe [209] studies the service rate control of a single class $G/M/n + GI$ queue. The objective is to minimize the discounted convex cost of service rate control plus linear delay and server idling costs and linear abandonment cost. It is shown that the solution to the Brownian control problem is of the feedback type. An asymptotically optimal policy is obtained by interpreting this solution in the context of the original problem. Ata et al. [28] models the gleaning operation as a novel queueing system and formulates an approximating drift rate control problem. The authors derive a nested threshold policy as the optimal staffing policy.

As mentioned above, various authors tackled admission or service rate control problems using Markov decision processes. Crabill [77] studies the control of an $M/M/1$ queueing system with a finite number of service rates. Crabill [78] considers a machine repair problem where a finite number of failure-prone machines require repair once they fail and the repair facility can decide the speed of the repair. Stidham and Weber [193] considers the problem of controlling the arrival and service rate of a queueing system with the objective of minimizing the total expected cost to reach state zero and shows that the optimal policy is monotone in the number of customers in the system. George and Harrison [97] studies the service rate control of a Markovian single-server queue with the objective of minimizing the long-run average holding cost plus cost-of-effort. It solves for the optimal state-dependent service rate and shows that the service rate is increasing in the queue length. Ata [18] considers a service rate control problem for a finite buffer Markovian queue motivated by a wireless communication application. The objective is to minimize the average service rate cost subject to an upper bound on the blocking probability. The author solves this problem explicitly, deriving the optimal service rate in closed-form. Ata and Zachariadis [35] extends this to multiple queues and channel fading, whereby the cost of service rate is modulated by a Markov chain. The authors derive the optimal service rates for each buffer, explicitly. Adusumilli and Hasenbein [7] studies the admission (through rejecting customers) and service rate control in a single-server queueing system. The authors propose an algorithm to solve

for the optimal policy that minimizes the long-run average holding cost, service control cost, and customer rejection cost. Kumar et al. [143] studies the dynamic service rate control of a single-server queue with a finite-state Markov-modulated Poisson arrival process. It is shown that the optimal service rate is non-decreasing in the number of customers in the system.

The third stream of literature studies the revenue management and pricing of queueing and manufacturing systems; see Elmaghraby and Keskinocak [85] and Gallego and Topaloglu [93] for overviews of this literature. The revenue management of manufacturing systems has been studied using various techniques. Caldentey and Wein [62] considers a single-product make-to-stock queue that uses two alternative selling channels: long-term contracts and a spot market of electronic orders. Under heavy traffic conditions, the problem is approximated by a diffusion-control problem, and analytical approximations are used to derive an effective policy. The closest paper to our work in this stream of literature is Çelik and Maglaras [67], that studies the problem of dynamic pricing, lead-time quotation, outsourcing, and scheduling in a make-to-order manufacturing system. It approximates the original problem with a drift rate control problem. The solution to the drift rate control problem is used to propose an effective dynamic control policy. Ata and Olsen [31] is another related paper that considers the problem of dynamically quoting price and leadtime menus for customers in a system where two classes of customers compete for a given resource. Customers have convex-concave delay costs. The authors derive a policy that is asymptotically optimal in the heavy traffic limit.

Many papers, as done in this chapter, study the revenue management of manufacturing and queueing systems by mapping them to admission and service rate control problems. Stidham [191] is among the first papers in this literature. Yoon and Lewis [219] studies congestion control in a non-stationary queueing system. It is shown that under the infinite horizon discounted and average reward optimality criteria, for each fixed time, optimal pricing and admission control strategies are non-decreasing in the number of customers in the system. Vulcano [204] studies the problem of dynamic pricing in a single-class make-to-stock

queue. The system manager can set the prices and a threshold for the admissible backlog. For the linear demand model, Vulcano [204] obtains closed-form expressions for the optimal demand rates, which are then used as an approximate solution to the original problem. Another related paper, Xu and Chao [216], studies the production rate control and pricing in a single-product make-to-stock queueing system. The controller aims to maximize the long-run average profit, which consists of the sales profit minus holding, backlog, and the linear cost of production effort. The controller can dynamically choose between a high and a low price and choose a production rate within a bounded interval. A closed-form solution for the prices is obtained and an algorithm is proposed to compute the optimal state-dependent price and production rates. A related paper to this stream of literature is Mendelson [164], that studies the static pricing of a queueing system with the objective of maximizing social welfare. In a related setting, Ata and Shneorson [33] studies the problem of dynamic arrival and service rate control in a state-dependent M/M/1 queueing system with the objective of maximizing social welfare. Mendelson and Whang [165] extends Mendelson [164] by considering a price menu design problem where the objective of the controller is to design a mechanism that is both (socially) optimal and incentive-compatible. Afèche [8] builds on this literature by considering a price/delay menu design problem with the objective of revenue maximization for a manufacturing system with two classes of customers. A novel problem formulation and solution method that combines the achievable region approach with mechanism design is developed. It is also shown that the work conserving $c\mu$ rule, known to be delay cost minimizing, incentive-compatible, and socially optimal, need not be revenue maximizing. A strategic delay policy may be optimal. Necessary and sufficient conditions for optimal strategic delay are also identified.

Other related papers in this literature include Ziya et al. [226, 227], Chao and Zhou [68], Gayon and Dallery [95], and Afèche et al. [9]. Ziya et al. [226] considers the problem of setting the optimal static price for a G/G/s/m queueing system. A lower bound on the optimal price is proved and the optimal price expressions for M/M/1/m and M/GI/s/s systems are

obtained. It is also shown that the optimal prices are monotone in the buffer size. Whether the prices are increasing or decreasing depends on whether the offered load is above or below a certain critical threshold. Ziya et al. [227] extends the results of Ziya et al. [226] to the multiclass case. Gayon and Dallery [95] studies the pricing and production control in a single-product make-to-stock manufacturing system with lost sales. Afèche et al. [9] studies revenue-maximizing tariffs that depend on realized leadtimes for a provider serving multiple time-sensitive customer types.

3.3 Model

This section advances a model of a make-to-stock manufacturing system selling K different products to price-sensitive customers. We model the manufacturing system as a multiclass single server queue. Class k products have a general production time distribution with mean m_k and squared coefficient of variation ν_{sk}^2 for $k = 1, \dots, K$; we refer to $\mu_k = m_k^{-1}$ as the production rate of class k products. Define $m = (m_k)$ and let $S_k(t)$ for $k = 1, \dots, K$ denote the number of class k products manufactured until time t if the system were to continuously work on class k products up to time t . We model the demand over time as a non-homogeneous Poisson process whose intensity depends on the prices the system manager charges. To be specific, demand for product k arrives according to a non-homogeneous Poisson process with instantaneous rate $\lambda_k(t)$ for $t \geq 0$ and $k = 1, \dots, K$. That is, the demand for product k up to time t is

$$N_k \left(\int_0^t \lambda_k(s) ds \right),$$

where N_k is a unit rate Poisson process. We assume that for $k = 1, \dots, K$, the processes N_k and S_k are independent of each other. Letting $\lambda(t) = (\lambda_k(t))$ denote the instantaneous demand rate vector at time t , we refer to $\lambda = \{\lambda(t) : t \geq 0\}$ as the instantaneous demand rate process.

The system manager chooses the price vector $p(t) = (p_k(t)) \in \mathcal{P}$ for $t \geq 0$, where $p_k(t)$ denotes the price of product k at time t for $k = 1, \dots, K$ and $\mathcal{P} \subset \mathbb{R}_+^K$ denotes the set of admissible price vectors.¹ We refer to $p = \{p(t) : t \geq 0\}$ as the price process. The price-sensitivity of demand is captured by a non-negative demand function Λ that maps the price vector to an instantaneous demand rate vector. That is,

$$\lambda(t) = \Lambda(p(t)), \quad t \geq 0.$$

The set of admissible instantaneous demand rate vectors, denoted by \mathcal{L} , is given as follows:

$$\mathcal{L} = \{\Lambda(p) \text{ for some } p \in \mathcal{P}\}.$$

Note that $\mathcal{L} \subset \mathbb{R}_+^K$ because Λ is non-negative by definition. Following Çelik and Maglaras [67], we make the following regularity assumptions on the demand function.

Assumption 5. *The demand function Λ satisfies the following:*

- (i) *It is bounded and continuously differentiable on \mathcal{P} .*
- (ii) *For each product $k = 1, \dots, K$, Λ_k is strictly decreasing in p_k for all $p \in \mathcal{P}$.*
- (iii) *The set of admissible instantaneous demand rate vectors \mathcal{L} is convex.*
- (iv) *There exists a unique inverse demand function $\Lambda^{-1} : \mathcal{L} \rightarrow \mathcal{P}$ that maps each admissible instantaneous demand rate vector to the price vector that induces it.*

We denote the variable cost of manufacturing for product k by δ_k for $k = 1, \dots, K$ and let $\delta = (\delta_k)$ denote the associated cost rate vector. Then, we define the profit rate function π as follows:

$$\pi(x) = x'(\Lambda^{-1}(x) - \delta), \quad x \in \mathcal{L}.$$

1. We assume that prices can be changed instantaneously at no cost.

For simplicity, we make the following regularity assumption on the profit rate function.

Assumption 6. *The profit rate function π is twice continuously differentiable on \mathcal{L} and has a negative definite Hessian matrix.*

An immediate consequence of Assumption 6 is that the profit rate function is strictly concave. Many common demand functions satisfy Assumptions 5-6; see Talluri and Van Ryzin [197, Section 7.3.3] for examples.

In what follows, we view the instantaneous demand rate process λ as the system manager's control, from which the price process p can be inferred using the inverse demand function Λ^{-1} , i.e., $p(t) = \Lambda^{-1}(\lambda(t))$ for $t \geq 0$. In order to avoid high backorder costs (due to long wait times), the system manager may outsource (or reject) customer orders when the backlog of work is judged excessive, thereby incurring outsourcing (or rejection) costs. We denote the number of class k orders outsourced (or rejected) up to time t by $O_k(t)$ for $k = 1, \dots, K$. Then, letting $O(t) = (O_k(t))$, we define $O = \{O(t) : t \geq 0\}$ to be K -dimensional outsourcing process. In addition to pricing and outsourcing decisions, the system manager makes dynamic scheduling decisions. We allow preemptive-resume scheduling and focus on head-of-line scheduling policies. Letting $T_k(t)$ denote the cumulative amount of time devoted to class k until time t , the system manager's dynamic scheduling policy can be described by a K -dimensional allocation process $T = (T_k)$. In summary, the dynamic control policy is denoted by (O, T, λ) , where O is the outsourcing process, T is the allocation process, and λ is the instantaneous demand process.

We let $Q_k(t)$ denote the (possibly negative) number of class k products in inventory at time t . Letting $Q(t) = (Q_k(t))$, the process $Q = \{Q(t), t \geq 0\}$ is called the inventory process. Assuming the system is empty initially, under policy (O, T, λ) , the inventory process Q evolves as follows: For $k = 1, \dots, K$ and $t \geq 0$,

$$Q_k(t) = S_k(T_k(t)) - N_k \left(\int_0^t \lambda_k(s) ds \right) + O_k(t). \quad (3.1)$$

The cumulative idleness process associated with scheduling policy T is defined as follows:

$$I(t) = t - \sum_{k=1}^K T_k(t), \quad t \geq 0. \quad (3.2)$$

The dynamic control policy (O, T, λ) is said to be feasible if it is non-anticipating and

$$I, O, T \text{ are non-decreasing with } I(0) = O(0) = T(0) = 0, \quad (3.3)$$

$$I, T \text{ are continuous,} \quad (3.4)$$

$$\lambda(t) \in \mathcal{L} \text{ for } t \geq 0. \quad (3.5)$$

We define the state cost function $q_k : \mathbb{R} \rightarrow \mathbb{R}_+$ that comprises of holding and backlogging costs as follows: For $k = 1, \dots, K$ and $x \in \mathbb{R}$,

$$q_k(x) = \begin{cases} \alpha_k x, & x \geq 0, \\ -\beta_k x, & x < 0, \end{cases}$$

where $\alpha_k > 0$ and $\beta_k > 0$ are the per unit holding and backorder cost of class k products, respectively. We denote the cost associated with outsourcing (or rejecting) a class k order by $\nu_k > 0$ for $k = 1, \dots, K$. Given π , q_k , and ν_k for $k = 1, \dots, K$, we define the cumulative profit process² associated with policy (O, T, λ) as follows:

$$V(t) = \int_0^t \pi(\lambda(s))ds - \sum_{k=1}^K \int_0^t q_k(Q_k(s))ds - \sum_{k=1}^K \nu_k O_k(t), \quad t \geq 0. \quad (3.6)$$

The first term on the right-hand side of (3.6) is a surrogate for the revenue (minus the variable costs) obtained from selling the products. This term can be interpreted as the expected revenue obtained from selling the products; see Plambeck et al. [175], Çelik and Maglaras [67], and Rubino and Ata [181] for similar treatments. The second term captures

2. Since the fixed cost of manufacturing does not depend on the policy, it is excluded from the profit calculation.

the holding cost associated with the unsold products and the backorder cost associated with the unfulfilled orders. The third term is the cost associated with the outsourced orders. Adopting the long-run average cost criterion, the system manager seeks to find the policy (O, T, λ) so as to

$$\text{maximize} \liminf_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}[V(t)] \text{ subject to (3.1) -- (3.5).}$$

Unfortunately, this formulation is not analytically tractable. Following Harrison [110, 114], Section 3.4 considers a sequence of closely related systems in the heavy-traffic regime and formulates the approximating Brownian control problem, which is tractable analytically.

To facilitate the analysis to follow, ignoring the randomness in the system, consider the following static planning problem: Find the instantaneous demand rate vector λ so as to

$$\text{maximize } \pi(\lambda) \text{ subject to } \lambda \in \mathcal{L}. \quad (3.7)$$

This formulation seeks to maximize the profit rate subject to the constraint $\lambda \in \mathcal{L}$ that ensures the instantaneous demand rate vector λ is achievable. The following assumption is crucial in deriving the approximating Brownian control problem and states that the profit maximizing demand rate puts the system in heavy traffic.

Assumption 7 (Balanced Loading Assumption). *The static planning problem (3.7) has a unique optimal solution $\lambda^* \in \text{interior}(\mathcal{L})$. Moreover, $\sum_{k=1}^K \lambda_k^* m_k = 1$.*

Since $\mathcal{L} \subset \mathbb{R}_+^K$, it follows from Assumption 7 that $\lambda_k^* > 0$ for $k = 1, \dots, K$. We refer to λ^* as the nominal instantaneous demand rate. This name reflects the fact that λ^* is derived from an idealized planning problem in which stochastic variability is suppressed. In the presence of such variability, it may be desirable for actual instantaneous demand rate to vary around the nominal rates λ^* depending on system status, which is the topic of the next section.

3.4 Approximating Brownian Control Problem

In Brownian approximations, one considers a sequence of closely related systems indexed by a parameter n , whose formal limit is the Brownian control problem. We attach a superscript n to various quantities of interest corresponding to the n -th system in this sequence. The asymptotic regime we focus on is the one where both the demand and system capacity grow with n . To be specific, we assume that for all $n \in \mathbb{N}$ and some $\eta \in \mathbb{R}_+^K$,

$$\Lambda^n(x) = n\Lambda(x), \quad x \in \mathcal{P}, \quad (3.8)$$

$$\mu_k^n = n\mu_k + \sqrt{n}\eta_k, \quad k = 1, \dots, K. \quad (3.9)$$

It is straightforward to show using (3.8) that for $n \in \mathbb{N}$ and $x \in \mathcal{L}$,

$$(\Lambda^n)^{-1}(nx) = \Lambda^{-1}(x) \quad \text{and} \quad \pi^n(nx) = n\pi(x). \quad (3.10)$$

It is evident from (3.7) and (3.10) that if the system manager were to ignore the randomness in the system, thus ignore all congestion-related and backlog-related costs, she would choose the instantaneous demand rate of $n\lambda^*$. However, in response to the system state, she may benefit from (dynamic) adjustments to the instantaneous demand rate. Moreover, the system inventory is of second order relative to the system size in the heavy-traffic asymptotic regime. Therefore, we focus our attention on instantaneous demand rate vectors of the following form: For all $n \in \mathbb{N}$ and some $\zeta : [0, \infty) \rightarrow \mathbb{R}^K$,

$$\lambda^n(t) = n\lambda^* + \sqrt{n}\zeta(t), \quad t \geq 0; \quad (3.11)$$

see Ata [19], Çelik and Maglaras [67], and Ata et al. [28] for similar treatments.

In other words, we study a large balanced-flow system for large n . For such systems, the inventory and outsourcing processes are expected to be of order \sqrt{n} . Thus, we scale them accordingly. To be specific, for $t \geq 0$, we define the scaled inventory and outsourcing

processes as follows:

$$Z^n(t) = \frac{Q^n(t)}{\sqrt{n}} \quad \text{and} \quad R^n(t) = \frac{O^n(t)}{\sqrt{n}}. \quad (3.12)$$

For $k = 1, \dots, K$, let $\rho_k = \lambda_k^*/\mu_k$ denote the proportion of time the system should devote to class k . We define the centered and scaled allocation and the scaled idleness processes for product $k = 1, \dots, K$ as follows:

$$Y_k^n(t) = \sqrt{n}(\rho_k t - T^n(t)) \quad \text{and} \quad U^n(t) = \sqrt{n}I^n(t), \quad t \geq 0. \quad (3.13)$$

We assume that the outsourcing cost ν_k^n , holding cost α_k^n , and backorder cost β_k^n vary with n as follows: For $k = 1, \dots, K$,

$$\nu_k^n = \frac{r_k}{\sqrt{n}}, \quad \alpha_k^n = \frac{h_k}{\sqrt{n}}, \quad \text{and} \quad \beta_k^n = \frac{b_k}{\sqrt{n}}, \quad (3.14)$$

where r_k , h_k , and b_k are given constants. It is then immediate from (3.14) that for $k = 1, \dots, K$ and $x \in \mathbb{R}$,

$$q_k^n(x) = \frac{g_k(x)}{\sqrt{n}}. \quad (3.15)$$

where

$$g_k(x) = \begin{cases} h_k x, & x \geq 0, \\ -b_k x, & x < 0. \end{cases}$$

Moreover, we define the cumulative cost process as the deviation of the cumulative profit process V^n from that in the corresponding deterministic system, $n\pi(\lambda^*)t$, which is an upper

bound on system performance. That is,

$$\xi^n(t) = n\pi(\lambda^\star)t - V^n(t), \quad t \geq 0. \quad (3.16)$$

In a similar fashion to Harrison [110], Appendix C.2.1 formally derives the approximating Brownian control problem as n gets large. In the approximating Brownian control problem, the performance-related processes ξ^n , Z^n , Y^n , U^n , and R^n are replaced with their formal limits ξ , Z , Y , U , and R , which jointly satisfy the following for $t \geq 0$:

$$\xi(t) = \int_0^t \zeta(s)' H \zeta(s) ds + \sum_{k=1}^K \int_0^t g_k(Z_k(s)) ds + \sum_{k=1}^K r_k R_k(t), \quad (3.17)$$

$$Z_k(t) = X_k(t) - \int_0^t \zeta_k(s) ds - \mu_k Y_k(t) + R_k(t), \quad k = 1, \dots, K, \quad (3.18)$$

$$U(t) = \sum_{k=1}^K Y_k(t), \quad (3.19)$$

$$U, R \text{ are nondecreasing with } U(0) = R(0) = 0, \quad (3.20)$$

where $H = -\nabla^2 \pi(\lambda^\star)/2$ and $X_k = \{X_k(t), t \geq 0\}$ are independent Brownian motions with infinitesimal drift $\rho_k \eta_k$ and infinitesimal variance $\sigma_k^2 = \lambda_k^\star(1 + \nu_{sk}^2)$.

Hereafter, we refer to (3.18)-(3.20) as the Brownian formulation and call the adapted control (R, Y, ζ) admissible if the corresponding state process Z satisfies

$$\limsup_{t \rightarrow \infty} \frac{\mathbb{E}[\|Z(t)\|]}{t} = 0, \quad (3.21)$$

where $\|\cdot\|$ denotes the Euclidean norm. The approximating Brownian Control Problem (BCP) can now be stated as follows: Choose the adapted control (R, Y, ζ) to

$$\text{minimize } \limsup_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}[\xi(t)] \text{ subject to (3.17) -- (3.21).} \quad (3.22)$$

As mentioned earlier, we consider a sequence of closely related systems indexed by n ,

whose formal limit is the Brownian control problem. The original control problem can be viewed as a specific element of this sequence of problems, which is determined by the particular choice of the system parameter n . Therefore, the final step in the Brownian approximation is to choose a system parameter n , which will be used to unscale the processes. The underlying assumption of the Brownian approximation is that the production and instantaneous demand rates (for the original problem of interest modeled in Section 3.3) are large enough so that various (scaled) processes of the original system can be approximated by the corresponding processes of the Brownian control problem. Eventually, one uses the same system parameter to interpret the solution to the Brownian control problem in the context of the original control problem.

We conclude this section by discussing the dynamic pricing policy prescribed by the BCP (3.22). Recall from (3.11) that the instantaneous demand rate $\lambda^n(t)$ is of the form $\lambda^n(t) = n\lambda^* + \sqrt{n}\zeta(t)$. It follows from Equation (3.8) and the one-to-one correspondence between the instantaneous demand rate process and the price process that the proposed price process in the n -th system is equal to the nominal price vector, $\Lambda^{-1}(\lambda^*)$, plus a term of order $1/\sqrt{n}$. To be specific, we have that

$$p^n(t) \simeq \Lambda^{-1}(\lambda^*) + \frac{\nabla\Lambda^{-1}(\lambda^*)\zeta(t)}{\sqrt{n}}, \quad t \geq 0, \quad (3.23)$$

where $\nabla\Lambda^{-1}(\lambda^*)$ denotes the Jacobian of Λ^{-1} evaluated at λ^* ; see Appendix C.2.2 for a derivation of (3.23).

3.5 Equivalent Workload Formulation

Although the Brownian control problem advanced in Section 3.4 is much simpler than the original control problem it approximates, its solution is not easy since it is a multidimensional control problem. This section develops a one-dimensional formulation that is equivalent to the Brownian control problem and is considerably more tractable. To this end, we define

the one-dimensional workload process W as follows:

$$W(t) = \sum_{k=1}^K m_k Z_k(t), \quad t \geq 0, \quad (3.24)$$

which represents the total (scaled) inventory (or backlog) in the system at time t measured in hours of total work for the server. Pre-multiplying Equation (3.18) by m' gives

$$W(t) = B(t) - \int_0^t \theta(s) ds - U(t) + L(t), \quad t \geq 0,$$

that describes the evolution of the workload process, where for $t \geq 0$,

$$B(t) = \sum_{k=1}^K m_k X_k(t), \quad \theta(t) = \sum_{k=1}^K m_k \zeta_k(t), \quad U(t) = \sum_{k=1}^K Y_k(t), \quad L(t) = \sum_{k=1}^K m_k R_k(t). \quad (3.25)$$

We refer to $L = \{L(t) : t \geq 0\}$ as the effective outsourcing process and refer to $\theta = \{\theta(t) : t \geq 0\}$ as the effective drift rate control process. One can interpret $U(t)$ as the cumulative (scaled) idleness by time t and interpret $L(t)$ as the cumulative (scaled) amount of work outsourced by time t . Defining the indices j^* and l^* as

$$j^* = \operatorname{argmin} \left\{ \frac{h_k}{m_k} : k = 1, \dots, K \right\} \quad \text{and} \quad l^* = \operatorname{argmin} \left\{ \frac{b_k}{m_k} : k = 1, \dots, K \right\},$$

we let $h^* = h_{j^*}/m_{j^*}$ and $b^* = b_{l^*}/m_{l^*}$ and define the effective holding cost function as follows: For $w \in \mathbb{R}$,

$$h(w) = \begin{cases} h^* w, & w \geq 0, \\ -b^* w, & w < 0. \end{cases} \quad (3.26)$$

The effective holding cost function h can be interpreted as the state cost associated with the following policy: If workload is positive, it is held as inventory of product j^* , the cheapest product to hold inventory per unit of work. If workload is negative, it is held as backordered

requests of product l^* , the cheapest product to have backordered requests per unit of work; see Wein [213] for a similar interpretation. Also, define the effective outsourcing cost and the corresponding class as follows:

$$i^* = \operatorname{argmin} \left\{ \frac{r_k}{m_k} : k = 1, \dots, K \right\} \quad \text{and} \quad \kappa = \frac{r_{i^*}}{m_{i^*}}. \quad (3.27)$$

The effective outsourcing cost κ is the outsourcing cost associated with the greedy outsourcing policy that only outsources product i^* , which is the cheapest product to outsource in the sense that it has the minimal outsourcing cost per unit of work outsourced. Moreover, we define the cost function $c(x)$ associated with the effective drift rate x and the corresponding optimal drift rate vector $\zeta^*(x)$ as follows: For $x \in \mathbb{R}$,

$$c(x) = \min \left\{ \zeta' H \zeta : m' \zeta = x, \zeta \in \mathbb{R}^K \right\} \quad \text{and} \quad \zeta^*(x) = \operatorname{argmin} \left\{ \zeta' H \zeta : m' \zeta = x, \zeta \in \mathbb{R}^K \right\}.$$

The next lemma provides closed-form expressions for c and ζ^* .

Lemma 7. *For $x \in \mathbb{R}$, we have that $c(x) = \frac{1}{m' H^{-1} m} x^2$ and $\zeta^*(x) = \frac{H^{-1} m}{m' H^{-1} m} x$.*

The Equivalent Workload Formulation (EWF) is then formulated as follows: Choose the adapted control (L, U, θ) to

$$\text{minimize} \quad \limsup_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\int_0^t c(\theta(s)) ds + \int_0^t h(W(s)) ds + \kappa L(t) \right] \quad (3.28)$$

subject to

$$W(t) = B(t) - \int_0^t \theta(s) ds - U(t) + L(t), \quad (3.29)$$

$$L, U \text{ are nondecreasing with } L(0) = U(0) = 0, \quad (3.30)$$

where B is a Brownian motion with infinitesimal drift $\mu = \sum_{k=1}^K m_k \rho_k \eta_k$ and infinitesimal variance $\sigma^2 = \sum_{k=1}^K m_k^2 \sigma_k^2$. We call the adapted control (L, U, θ) admissible if it satisfies

(3.29)-(3.30) and

$$\limsup_{t \rightarrow \infty} \frac{\mathbb{E}[|W(t)|]}{t} = 0. \quad (3.31)$$

We refer to W as the workload process associated with control (L, U, θ) . Proposition 5 establishes the equivalence of the EWF (3.28)-(3.30) and the BCP (3.22). As a preliminary to Proposition 5, we define the optimal workload configuration (or, the lifting map) $\Delta : \mathbb{R} \rightarrow \mathbb{R}^K$ by

$$\Delta_k(w) = \begin{cases} \frac{w}{m_k}, & \text{if } k = j^* \text{ and } w \geq 0, \\ \frac{w}{m_k}, & \text{if } k = l^* \text{ and } w < 0, \\ 0, & \text{otherwise,} \end{cases} \quad (3.32)$$

for $k = 1, \dots, K$. The optimal workload configuration distributes the workload as follows: If workload is positive, it is held as inventory of product j^* , the cheapest product to hold inventory per unit of work. If workload is negative, it is held as backordered requests of product l^* , the cheapest product to have backordered requests per unit of work.

Proposition 5. *The EWF (3.28)-(3.30) is equivalent to the BCP (3.22) in the following sense:*

- (i) *Suppose that (R, Y, ζ) is an admissible control for the BCP (3.22) with Brownian motion X and state process Z . Then (L, U, θ) defined by (3.19) and (3.25) is an admissible control for the EWF (3.28)-(3.30) with Brownian motion B given by (3.25) and state process W given by (3.24). Moreover, the cost of control (L, U, θ) for the EWF (3.28)-(3.30) is less than or equal to the cost of control (R, Y, ζ) for the BCP (3.22).*
- (ii) *Suppose that (L, U, θ) is an admissible control for the EWF (3.28)-(3.30) with Brownian motion B and state process W . Then, there exists a Brownian motion X that satisfies (3.25). Furthermore, there exists (R, Y, ζ) that is an admissible control for the BCP (3.22) with Brownian motion X and state process $\{\Delta(W(t)) : t \geq 0\}$, where Δ is given*

by (3.32). These two policies have the same cost.

3.6 Solution to the Equivalent Workload Formulation

In this section, we propose an optimal solution to the EWF (3.28)-(3.30). To minimize technical complexity, we restrict our attention to stationary Markov control policies. That is, the (effective) drift rate $\theta(t)$ chosen at any time t is assumed to depend on past history only through the workload $W(t)$ at time t . To reflect this assumption, in what follows, we write $\theta(W(t))$ as opposed to $\theta(t)$. The optimal outsourcing and idling policies we derive can be viewed as a barrier policy, which we define in Section 3.6.1.

3.6.1 Barrier Policies

Following Harrison [115, Section 7.7], we define a barrier policy as follows.

Definition 4 (Barrier Policy). *Given policy parameters $l, u \in \mathbb{R}$ with $l < u$, we call (L, U, θ) a barrier policy if it is an admissible control for the EWF (3.28)-(3.30), and it satisfies $W(t) \in [l, u]$,*

$$\int_0^t \mathbf{1}_{\{W(s)>l\}} dL(s) = 0 \quad \text{and} \quad \int_0^t \mathbf{1}_{\{W(s)<u\}} dU(s) = 0, \quad t \geq 0. \quad (3.33)$$

Processes L and U enforce the lower and upper reflecting barriers at l and u , respectively. Let $C^2[l, u]$ denote the space of functions $f : [l, u] \rightarrow \mathbb{R}$ that are twice continuously differentiable up to the boundary (i.e., f is twice continuously differentiable on the interior of the interval, and its first and second derivatives both approach finite limits at the end points). For a given real-valued function $\theta(\cdot)$, define the differential operator Γ_θ as follows:

$$\Gamma_\theta f(w) = \frac{1}{2}\sigma^2 f''(w) + \mu f'(w) - \theta(w)f'(w), \quad w \in (l, u).$$

Consider $\gamma \in \mathbb{R}$ and $f \in C^2([l, u])$ and assume that they jointly satisfy

$$\Gamma_\theta f(w) + c(\theta(w)) + h(w) = \gamma, \quad w \in (l, u), \quad (3.34)$$

subject to the boundary conditions

$$f'(l) = -\kappa \quad \text{and} \quad f'(u) = 0. \quad (3.35)$$

Proposition 6 provides a useful identity and will help motivate the Bellman equation; see Appendix C.3 for its proof.

Proposition 6. *Consider the barrier policy (L, U, θ) with a lower barrier at l and an upper barrier at u . If $\gamma \in \mathbb{R}$ and $f \in C^2([l, u])$ jointly satisfy (3.34)-(3.35), then*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\int_0^t c(\theta(W(s))) ds + \int_0^t h(W(s)) ds + \kappa L(t) \right] = \gamma.$$

3.6.2 Bellman Equation

Proposition 6 and the smooth pasting arguments (see e.g., Harrison [115, Section 7.7]) motivate the following Bellman equation: find $l, u, \gamma \in \mathbb{R}$ and $f \in C^2([l, u])$ that satisfy

$$\min_{x \in \mathbb{R}} \left\{ \frac{1}{2} \sigma^2 f''(w) + \mu f'(w) - x f'(w) + c(x) + h(w) \right\} = \gamma, \quad w \in (l, u), \quad (3.36)$$

subject to the boundary conditions

$$f'(l) = -\kappa \quad \text{and} \quad f'(u) = 0, \quad (3.37)$$

and the smooth pasting conditions

$$f''(l) = 0 \quad \text{and} \quad f''(u) = 0. \quad (3.38)$$

Motivated by Proposition 6, we interpret γ as a guess at the minimum average cost and interpret l and u as the lower and upper reflecting barriers to be imposed on the workload process. The unknown function f is often called the relative value function in average cost dynamic programming.

The Bellman equation is introduced primarily to motivate our solution approach; the properties of the Bellman equation that we require will be proved from first principles. We shall develop an explicit solution (l, u, γ, f) for this Bellman equation, and define the candidate policy as the one that chooses in each state w the (effective) drift rate $\theta(w)$ equal to the minimizer x in (3.36). Then, we will prove that this candidate policy is optimal.

As a preliminary to analyzing (3.36)-(3.38), following Ata et al. [26, Section 2], define the convex conjugate ϕ of c , and its derivative as follows:

$$\phi(y) = \sup_{x \in \mathbb{R}} \{yx - c(x)\} \quad \text{and} \quad \psi(y) = \operatorname{argmax}_{x \in \mathbb{R}} \{yx - c(x)\}, \quad y \in \mathbb{R}. \quad (3.39)$$

The following lemma provides closed-form expressions for ϕ and ψ .

Lemma 8. *For $y \in \mathbb{R}$, we have that*

$$\phi(y) = \frac{m' H^{-1} m}{4} y^2 \quad \text{and} \quad \psi(y) = \frac{m' H^{-1} m}{2} y. \quad (3.40)$$

Since Equation (3.36) does not involve the unknown function f itself, it is really a first-order equation. Setting $v(w) = f'(w)$ and recalling the definition (3.39), we can rewrite (3.36) as

$$\frac{1}{2}\sigma^2 v'(w) + \mu v(w) + h(w) - \phi(v(w)) = \gamma, \quad w \in (l, u). \quad (3.41)$$

Then, using Lemma 8 and rearranging the terms in (3.41), we rewrite the Bellman equation

as follows: find $l, u, \gamma \in \mathbb{R}$ and $v \in C^1[l, u]$ that satisfy³

$$v'(w) = \frac{m'H^{-1}m}{2\sigma^2} v^2(w) - \frac{2\mu}{\sigma^2} v(w) + \frac{2(\gamma - h(w))}{\sigma^2}, \quad w \in (l, u), \quad (3.42)$$

subject to the boundary conditions

$$v(l) = -\kappa \quad \text{and} \quad v(u) = 0, \quad (3.43)$$

and the smooth pasting conditions

$$v'(l) = 0 \quad \text{and} \quad v'(u) = 0. \quad (3.44)$$

Equation (3.42) is a Riccati equation; see Appendix C.1 for a brief overview of Riccati equations.

3.6.3 Solution to the Bellman Equation

To solve the Bellman equation, we proceed in three steps. First, we use the boundary and smooth pasting conditions (3.43)-(3.44) to express l, u in terms of γ and show that $l < 0 < u$. Second, for each γ , we solve Equation (3.42) on $[l, 0]$ and $[0, u]$, separately. This step exploits the special structure of the Bellman equation, i.e., the solution of the Riccati equation. Finally, we use the desired continuity of $v(w)$ at $w = 0$ to pin down γ , which completes the solution of the Bellman equation.

To this end, we substitute $v(l) = -\kappa$ and $v'(l) = 0$ into Equation (3.42) for $w = l$ to obtain

$$\gamma = h(l) - \mu\kappa - \frac{m'H^{-1}m}{4} \kappa^2. \quad (3.45)$$

3. Note that $C^1[l, u]$ denotes the space of functions $v : [l, u] \rightarrow \mathbb{R}$ that are continuously differentiable up to the boundary (i.e., v is continuously differentiable on the interior of the interval, and its first derivative approaches finite limits at the end points).

Similarly, substituting $v(u) = v'(u) = 0$ into Equation (3.42) for $w = u$ gives

$$\gamma = h(u). \quad (3.46)$$

Combining (3.45) and (3.46) gives

$$h(u) = h(l) - \mu\kappa - \frac{m'H^{-1}m}{4}\kappa^2. \quad (3.47)$$

The next result uses Equations (3.45)-(3.47) to obtain useful structural insights.

Lemma 9. *Assume $l, u, \gamma \in \mathbb{R}$ and $v \in C^1([l, u])$ jointly satisfy (3.42)-(3.44). Then, we have that $\gamma > 0$ and*

$$l = \frac{-1}{b^*}(\gamma + \mu\kappa + \frac{m'H^{-1}m}{4}\kappa^2) < 0 \quad \text{and} \quad u = \frac{\gamma}{h^*} > 0. \quad (3.48)$$

Next, we split the analysis of the Bellman equation (3.42)-(3.44) into two sub-intervals: $[l, 0]$ and $[0, u]$. We fix γ and solve the Bellman equation (3.42)-(3.44) on each sub-interval, separately. Then, we use the continuity of v at zero to pin down γ . To do so, fix $\gamma \geq 0$ and define⁴

$$l_\gamma = \frac{-1}{b^*}(\gamma + \mu\kappa + \frac{m'H^{-1}m}{4}\kappa^2) < 0 \quad \text{and} \quad u_\gamma = \frac{\gamma}{h^*} \geq 0.$$

Substituting (3.45) into (3.42) and focusing on the interval $(l_\gamma, 0]$ yields the following:

$$v'(w) = \frac{m'H^{-1}m}{2\sigma^2}(v^2(w) - \kappa^2) - \frac{2\mu}{\sigma^2}(v(w) + \kappa) + \frac{2(h(l_\gamma) - h(w))}{\sigma^2}, \quad w \in (l_\gamma, 0], \quad (3.49)$$

4. Although Lemma 9 proves that the average cost parameter γ that solves the Bellman equation (3.42)-(3.44) is strictly positive, we do not assume this apriori and allow $\gamma \geq 0$ for mathematical convenience.

subject to the boundary condition

$$v(l_\gamma) = -\kappa. \quad (3.50)$$

It is straightforward to verify by substituting (3.50) into (3.49) for $w = l_\gamma$ that the solution to (3.49)-(3.50) satisfies the smooth pasting condition $v'(l_\gamma) = 0$.

Similarly, substituting (3.46) into (3.42) and focusing on the interval $[0, u_\gamma]$ yields the following:

$$v'(w) = \frac{m'H^{-1}m}{2\sigma^2} v^2(w) - \frac{2\mu}{\sigma^2} v(w) + \frac{2(h(u_\gamma) - h(w))}{\sigma^2}, \quad w \in [0, u_\gamma] \quad (3.51)$$

subject to the boundary condition

$$v(u_\gamma) = 0. \quad (3.52)$$

It is again straightforward to verify by substituting (3.52) into (3.51) for $w = u_\gamma$ that the solution to (3.51)-(3.52) satisfies the smooth pasting condition $v'(u_\gamma) = 0$.

It is straightforward to show that (3.49)-(3.50) has a unique continuously differentiable solution, denoted by $v_\gamma^- : [l_\gamma, 0] \rightarrow \mathbb{R}$ for each $\gamma \geq 0$. Similarly, for $\gamma \geq 0$, (3.51)-(3.52) has a unique continuously differentiable solution, denoted by $v_\gamma^+ : [0, u_\gamma] \rightarrow \mathbb{R}$.⁵ Lemma 10 shows that $v_\gamma^-(0)$ is strictly increasing in γ and $v_\gamma^+(0)$ is strictly decreasing in γ ; see Appendix C.3 for its proof.

Lemma 10. *We have the following:*

(i) $v_\gamma^-(0)$ is continuous and strictly increasing in γ with $v_0^-(0) < 0$ and $\lim_{\gamma \rightarrow \infty} v_\gamma^-(0) = \infty$.

(ii) $v_\gamma^+(0)$ is continuous and strictly decreasing in γ with $v_0^+(0) = 0$ and $\lim_{\gamma \rightarrow \infty} v_\gamma^+(0) = -\infty$.

5. Equations (3.49)-(3.50) and (3.51)-(3.52) belong to a class of Riccati differential equations that admits unique C^1 solutions; see e.g., Zaitsev and Polyanin [221, Section 1.2.2].

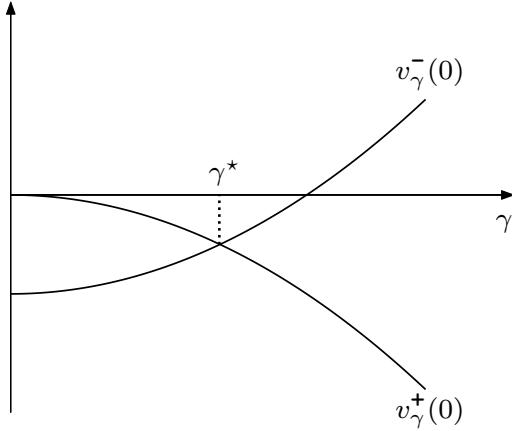


Figure 3.1: $v_\gamma^-(0)$ and $v_\gamma^+(0)$ as a function of the average cost γ .

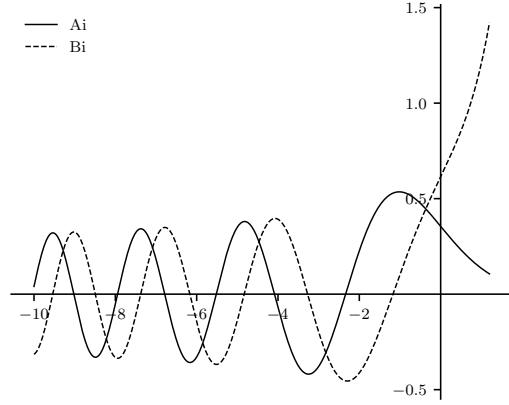


Figure 3.2: Airy functions of the first and second kind.

The following corollary is immediate from Lemma 10 and the unique γ characterized in the corollary is illustrated in Figure 3.1. Corollary 3 is crucially used in solving the Bellman equation (3.42)-(3.44).

Corollary 3. *There exists a unique γ^* such that $v_{\gamma^*}^-(0) = v_{\gamma^*}^+(0)$.*

As mentioned earlier, Equations (3.49)-(3.50) and (3.51)-(3.52) both fall in the class of Riccati differential equations; e.g., see Zaitsev and Polyanin [221, Section 1.2]. Next, we use the structure of the Riccati differential equations to obtain closed-form expressions for v_γ^- and v_γ^+ . To facilitate our analysis, define $\text{Ai} : \mathbb{R} \rightarrow \mathbb{R}$ and $\text{Bi} : \mathbb{R} \rightarrow \mathbb{R}$ as

$$\begin{aligned}\text{Ai}(w) &= \frac{1}{\pi} \int_0^\infty \cos\left(\frac{1}{3}x^3 + xw\right) dx, \quad w \in \mathbb{R}. \\ \text{Bi}(w) &= \frac{1}{\pi} \int_0^\infty \left(\cos\left(-\frac{1}{3}x^3 + xw\right) + \sin\left(\frac{1}{3}x^3 + xw\right) \right) dx, \quad w \in \mathbb{R}.\end{aligned}$$

Functions Ai and Bi are referred to as the Airy functions of the first and second kind, respectively. The Airy functions of the first and second kind are depicted in Figure 3.2. The Airy functions oscillate on $(-\infty, 0)$ and are monotone on $(0, \infty)$.⁶ As a preliminary to

6. The Airy functions can be expressed in terms of the Bessel functions and the modified Bessel functions of order 1/3; see e.g., Abramowitz and Stegun [4, Section 10.4].

discussing closed-form expressions for v_γ^- and v_γ^+ , define the functions

$$\begin{aligned}\bar{z}(w) &= (m'H^{-1}mb^*\sigma^2)^{-2/3} (\mu^2 - m'H^{-1}m(b^*w + b^*l_\gamma + \gamma)), \quad w \in [0, -l_\gamma] \\ \hat{z}(w) &= (m'H^{-1}mh^*\sigma^2)^{-2/3} (\mu^2 - m'H^{-1}mh^*w), \quad w \in [0, u_\gamma].\end{aligned}$$

Similarly, define the constants

$$\begin{aligned}\bar{C}_1 &= \frac{\text{Bi}'(\bar{z}(0)) + \text{Bi}(\bar{z}(0))(m'H^{-1}mb^*\sigma^2)^{-1/3}(\frac{1}{2}m'H^{-1}m\kappa + \mu)}{\text{Ai}(\bar{z}(0))\text{Bi}'(\bar{z}(0)) - \text{Ai}'(\bar{z}(0))\text{Bi}(\bar{z}(0))} \\ \bar{C}_2 &= \frac{-\text{Ai}'(\bar{z}(0)) - \text{Ai}(\bar{z}(0))(m'H^{-1}mb^*\sigma^2)^{-1/3}(\frac{1}{2}m'H^{-1}m\kappa + \mu)}{\text{Ai}(\bar{z}(0))\text{Bi}'(\bar{z}(0)) - \text{Ai}'(\bar{z}(0))\text{Bi}(\bar{z}(0))} \\ \hat{C}_1 &= \frac{\text{Bi}'(\hat{z}(0)) - \text{Bi}(\hat{z}(0))(m'H^{-1}mh^*\sigma^2)^{-1/3}\mu}{\text{Ai}(\hat{z}(0))\text{Bi}'(\hat{z}(0)) - \text{Ai}'(\hat{z}(0))\text{Bi}(\hat{z}(0))} \\ \hat{C}_2 &= \frac{-\text{Ai}'(\hat{z}(0)) + \text{Ai}(\hat{z}(0))(m'H^{-1}mh^*\sigma^2)^{-1/3}\mu}{\text{Ai}(\hat{z}(0))\text{Bi}'(\hat{z}(0)) - \text{Ai}'(\hat{z}(0))\text{Bi}(\hat{z}(0))}\end{aligned}$$

and define the functions

$$\bar{y}(w) = \bar{C}_1 \exp(-\mu w/\sigma^2) \text{Ai}(\bar{z}(w)) + \bar{C}_2 \exp(-\mu w/\sigma^2) \text{Bi}(\bar{z}(w)), \quad w \in [0, -l_\gamma]$$

$$\hat{y}(w) = \hat{C}_1 \exp(\mu w/\sigma^2) \text{Ai}(\hat{z}(w)) + \hat{C}_2 \exp(\mu w/\sigma^2) \text{Bi}(\hat{z}(w)), \quad w \in [0, u_\gamma].$$

Lemma 11 provides closed-form expressions for v_γ^- and v_γ^+ in terms of the Airy functions; see Appendix C.3 for its proof.

Lemma 11. *We have the following closed form solutions for v_γ^- and v_γ^+ :*

$$\begin{aligned}v_\gamma^-(w) &= -\frac{2\sigma^2}{m'H^{-1}m} \frac{\bar{y}'(w - l_\gamma)}{\bar{y}(w - l_\gamma)} w \in [l_\gamma, 0], \\ v_\gamma^+(w) &= \frac{2\sigma^2}{m'H^{-1}m} \frac{\hat{y}'(u_\gamma - w)}{\hat{y}(b_\gamma - w)} w \in [0, u_\gamma].\end{aligned}$$

The parameter γ^* , characterized in Corollary 3, can be computed using the closed-form expressions for v_γ^- and v_γ^+ provided in Lemma 11. By Lemma 10, this computation can be

done using a simple line search. Given γ^* , we define

$$v(w) = \begin{cases} v_{\gamma^*}^-(w), & w \in [l_{\gamma^*}, 0], \\ v_{\gamma^*}^+(w), & w \in (0, u_{\gamma^*}]. \end{cases} \quad (3.53)$$

Proposition 7 shows that $(l_{\gamma^*}, u_{\gamma^*}, \gamma^*, v)$ is the unique solution to the Bellman equation (3.42)-(3.44); see Appendix C.3 for its proof.

Proposition 7. *The function v is non-positive, strictly increasing, and continuously differentiable on $[l_{\gamma^*}, u_{\gamma^*}]$. Moreover, $(l_{\gamma^*}, u_{\gamma^*}, \gamma^*, v)$ is the unique solution to (3.42)-(3.44).*

To solve the Bellman equation (3.36)-(3.38), define

$$f(x) = \int_{l_{\gamma^*}}^w v(x)dx, \quad w \in [l_{\gamma^*}, u_{\gamma^*}], \quad (3.54)$$

where v is given by (3.53). The following corollary provides the unique solution to the Bellman equation (3.36)-(3.38).

Corollary 4. *The function f is strictly decreasing, strictly convex, and twice continuously differentiable. Moreover, $(l_{\gamma^*}, u_{\gamma^*}, \gamma^*, f)$ solves the Bellman equation (3.36) subject to the boundary condition (3.37) and the smooth pasting condition (3.38). Furthermore, f is unique up to an additive constant.*

3.6.4 Candidate Policy and Its Optimality for the Equivalent Workload Formulation

Our candidate policy is the barrier policy with the lower barrier at l_{γ^*} , the upper barrier at u_{γ^*} , and the (effective) drift rate that is the minimizer of the left-hand side of Equation

(3.36) at every $w \in [l_{\gamma^*}, u_{\gamma^*}]$, i.e.,

$$\theta(w) = \frac{m' H^{-1} m}{2} v(w), \quad w \in [l_{\gamma^*}, u_{\gamma^*}], \quad (3.55)$$

where v is given by (3.53). The candidate policy is stationary and its (effective) drift rate is non-positive, continuous, and strictly increasing in the workload. Under the candidate policy, the workload process W evolves as a diffusion process with a strictly increasing state-dependent drift rate and reflecting barriers at l_{γ^*} and u_{γ^*} . Theorem 2 shows that the candidate policy is optimal for the EWF (3.28)-(3.30); see Appendix C.3 for its proof.

Theorem 2. *The barrier policy with the lower barrier at l_{γ^*} , the upper barrier at u_{γ^*} , and the (effective) drift rate control $\theta(\cdot)$ given by (3.55) is optimal for the EWF (3.28)-(3.30) and it has a long-run average cost of γ^* .*

3.7 Proposed Policy and Its Performance

This section proposes a dynamic control policy for the problem introduced in Section 3.3 by interpreting the solution to the EWF (3.28)-(3.30) in the context of the original control problem. The proposed policy has three components: pricing, outsourcing, and scheduling. To describe the policy, recall that we consider a sequence of systems indexed by n , whose formal limit is the BCP (3.22). The original control problem introduced in Section 3.3 is a specific element of this sequence of problems, which is determined by the particular choice of n . Therefore, in order to describe the proposed policy, first, we need to choose the system parameter n that will be used to unscale the processes of interest. Given n , define $\mathcal{W} = \{\mathcal{W}(t) : t \geq 0\}$ as the unscaled, or nominal, workload process associated with the control problem introduced in Section 3.3. To be specific, $\mathcal{W}(t) = \sum_{k=1}^K m_k Q_k^n(t)$ for $t \geq 0$.

Pricing Policy. Given the nominal workload process \mathcal{W} , we choose the price vector

$$p(t) = \Lambda^{-1}(\lambda^*) + \frac{\nabla \Lambda^{-1}(\lambda^*)}{\sqrt{n}} \frac{H^{-1}m}{2} v(l_{\gamma^*} \vee (\frac{\mathcal{W}(t)}{\sqrt{n}} \wedge u_{\gamma^*})), \quad t \geq 0, \quad (3.56)$$

where v is given by (3.53) and $\nabla \Lambda^{-1}(\lambda^*)$ denotes the Jacobian of Λ^{-1} evaluated at λ^* .

This pricing policy follows from Theorem 2, Lemma 7, and Equation (3.23). The term $l_{\gamma^*} \vee ((\mathcal{W}(t)/\sqrt{n}) \wedge u_{\gamma^*})$, which projects $\mathcal{W}(t)/\sqrt{n}$ onto $[l_{\gamma^*}, u_{\gamma^*}]$, enables us to prescribe a price for any value of the nominal workload process. Note that although the workload process W in the equivalent workload formulation lives in $[l_{\gamma^*}, u_{\gamma^*}]$, the scaled nominal workload process $\mathcal{W}(t)/\sqrt{n}$ in the original control problem may leave the interval $[l_{\gamma^*}, u_{\gamma^*}]$, albeit with vanishing likelihood as the system gets large. Recall that the optimal effective drift rate control θ is non-positive and strictly increasing in the workload. Thus, our proposed pricing policy induces an effective demand rate $\sum_{k=1}^K \lambda_k(t)m_k$ that is strictly increasing in the workload and less than or equal to $\sum_{k=1}^K \lambda_k^* m_k = 1$.

Outsourcing Policy. Outsource orders of product i^* when the nominal workload \mathcal{W} is less than or equal to the outsourcing threshold $l^n = \sqrt{n}l_{\gamma^*} < 0$.

The optimal solution to the EWF (3.28)-(3.30) imposes a lower barrier on the workload process W at l_{γ^*} using the effective outsourcing process L . We interpret this in the context of the original control problem as stated above. The intuition behind this policy is that when the backlog is so large that $\mathcal{W} < l^n$, it is worthwhile to outsource orders as opposed to letting the backlog grow further. Furthermore, the system manager only outsources the "cheapest" product, i.e., product i^* ; see Equation (3.27).

The approximating Brownian control problem sets all but one of the inventory levels to zero (whenever the total workload is positive). As articulated in Harrison [111], the zeros in the Brownian control problem correspond to positive but small inventory levels, e.g., safety stocks, for further details see Harrison [111, 112], Harrison and López [116], Maglaras [160], and Ata and Kumar [27]. Therefore, as is customary in the heavy traffic literature, we will

put small safety stocks in various buffers. The safety stock for buffer k is denoted by s_k for $k = 1, \dots, K$; the values s_k can be calibrated via simulation as done in Wein [213].

Scheduling Policy. Interpreting the solution to the equivalent workload formulation in the context of the original control problem, we propose idling the server whenever the nominal workload \mathcal{W} exceeds the idling threshold $u^n = \sqrt{n}u_{\gamma^*}$ and $Q_k^n \geq s_k$ for $k = 1, \dots, K$. Whenever the server is working, it first prioritizes those products for which $Q_k^n < s_k$, i.e., the inventory is below the safety stock. Among such products, the system manager prioritizes them in the descending order of b_k/m_k . Finally, if $Q_k^n \geq s_k$ for all $k = 1, \dots, K$ and $\mathcal{W} < u^n$, i.e., the server is working, then the server focuses solely on product j^* , i.e., the cheapest product to hold in inventory.

Next, we consider an example that illustrates the effectiveness of the proposed policy. Following Wein [213, Section 10], we consider a manufacturing system with $K = 3$ products. We assume the production times are exponentially distributed with production rates $\mu_1 = 100$, $\mu_2 = 100/2$, and $\mu_3 = 100/3$; the production rates in our example are hundred times the corresponding values in Wein [213]. Following Wein [213], we take the holding and backorder costs to be $\alpha_1 = \alpha_2 = \alpha_3 = 1$ and $\beta_1 = \beta_2 = \beta_3 = 2$, respectively. We take the outsourcing costs to be higher but in the same order as the holding and backorder costs. To be specific, we take $\nu_1 = \nu_2 = \nu_3 = 3$. Following Çelik and Maglaras [67], we use a multinomial logit demand model. We assume customers arrive according to a Poisson process with rate 100 and choose to order product k with probability

$$\mathbb{P}(\text{Ordering Product } k) = \frac{\exp(a_k - d_k p_k(t))}{1 + \sum_{i=1}^3 \exp(a_i - d_i p_i(t))},$$

where $a_1 = 1.87$, $a_2 = 2.46$, $a_3 = 0.13$, $d_1 = 1$, $d_2 = 2$, and $d_3 = 1/2$. The nominal instantaneous demand rate under this demand model is $\lambda^* = (30, 15, 10)$ and the corresponding average server utilization is $\sum_{k=1}^3 \lambda_k^*/\mu_k = 0.9$. Demand model primitives a_k and d_k are chosen such that the nominal instantaneous demand rate λ^* is 100 times the demand rate

in Wein [213].⁷

We take the system parameter to be $n = 100$. As done by Wein [213], for simplicity, we focus on safety stock parameters of the form $s_1 = s_2 = s_3$. Throughout the simulation study, we observe that the optimal (common) safety stock parameter is $s_1 = s_2 = s_3 = 1$. Table 3.1 reports the simulation results. Our proposed policy is compared against two other policies. The first policy is the optimal joint dynamic pricing, outsourcing, and scheduling policy that is computed using the dynamic programming approach numerically. The second policy, indicated by "Static Policy", uses the static price $\Lambda^{-1}(\lambda^*)$ alongside the outsourcing and scheduling policies proposed in Section 3.7. Each policy is simulated 100 times for 10^6 time units starting from an initial inventory of zero. In order to assess the steady state performance of the system, the initial 20% of each run is discarded. The values inside the parenthesis in Table 3.1 indicate the 95% confidence interval. Table 3.1 shows that the dynamic pricing policy is effective and has a small optimality gap. Moreover, the gap between the static policy and the proposed policy (that uses dynamic prices) is substantial, which highlights the value of dynamic pricing.

Table 3.1: Simulation output.

Utilization	Profit (per unit of time)			Optimality Gap	
	Static Policy	Proposed Policy	Optimal Policy	Static Policy	Proposed Policy
0.90	116.75 (± 0.03)	120.66 (± 0.03)	122.72	4.86%	1.68%
0.95	114.04 (± 0.03)	120.03 (± 0.03)	122.09	6.82%	1.69%
0.97	113.00 (± 0.03)	119.74 (± 0.03)	121.82	7.24%	1.70%
0.99	112.14 (± 0.03)	119.25 (± 0.03)	121.54	7.73%	1.88%

Next, we decrease the production rates with the same factor to achieve server utilization values of 0.95, 0.97, and 0.99. Table 3.1 shows that the gap between the static and proposed policies grows with utilization, which shows that the value of dynamic pricing increases with the utilization of the system.

7. These values of λ_k^* and μ_k for $k = 1, 2, 3$ are consistent with Wein [213] in the heavy traffic limit, because Wein [213] scales time by $n = 100$ to arrive at the approximating Brownian control problem whereas we scale the instantaneous demand and service rates by $n = 100$.

Finally, to gain a better understanding of the parameters impacting the optimality gap of the static and proposed policies, we take the demand parameters $d_1 = \bar{d}$, $d_2 = 2\bar{d}$, and $d_3 = \bar{d}/2$, where \bar{d} is a tunable parameter. Tables 3.2 and 3.3 depict the impact of \bar{d} and the (common) outsourcing cost on the optimality gap of the static and proposed policies, respectively. In these simulations, server utilization is maintained at 0.90. We observe that as the price sensitivity of demand increases the proposed policy significantly outperforms the static one, although both their performances degrade relative to the optimal policy computed numerically by dynamic programming. Moreover, the gap between the static and proposed policies grows with the outsourcing cost, which shows that the dynamic pricing policy is more effective when the outsourcing cost is higher since dynamic pricing can decrease the chances of an outsourcing event by reducing the effective demand rate when backlog is high.

Table 3.2: Opt. gap of the static policy.

\bar{d}	Outsourcing Cost			
	1	2	3	4
1	1.69%	3.32%	4.86%	6.01%
2	4.48%	10.19%	14.45%	16.31%
3	8.12%	18.19%	24.49%	24.98%

Table 3.3: Opt. gap of the proposed policy.

\bar{d}	Outsourcing Cost			
	1	2	3	4
1	1.27%	1.64%	1.68%	1.69%
2	2.50%	3.12%	3.12%	3.13%
3	3.81%	4.30%	4.30%	4.30%

Table 3.4: Optimality gaps.

3.8 Concluding Remarks

This chapter studies the joint dynamic pricing, outsourcing, and scheduling of a multiclass make-to-stock manufacturing system. By considering the manufacturing system under the heavy traffic regime, an approximating Brownian control problem is obtained. It is shown that the Brownian control problem is equivalent to a one-dimensional workload formulation, whose state process, the workload process, represents the total (scaled) inventory (or backlog) in the system measured in hours of total work for the manufacturing system. The optimal

solution to the equivalent workload formulation is a two-sided barrier policy with state-dependent drift rate. Outsourcing and idling processes are used to keep the workload process between the lower and upper reflecting barriers. Between the two barriers, a state-dependent drift rate is used to control the workload process. By interpreting this solution in the context of the original problem, a dynamic pricing, outsourcing, and scheduling policy is proposed for the manufacturing system. It is demonstrated through a simulation study that the proposed policy performs well compared to the optimal control policy.

Although this chapter focuses on a make-to-stock manufacturing system with backlogged orders, its solution technique can be easily adapted to the lost sales case by imposing an exogenous lower reflecting barrier at zero on the workload process. Letting κ denote the (effective) lost sales cost in the workload formulation, the Bellman equation would be: find $u, \gamma \in \mathbb{R}$ and $f \in C^2([0, u])$ that satisfy Equation (3.36) for $w \in (0, u)$ subject to the boundary conditions $f'(0) = -\kappa$ and $f'(u) = 0$ and the smooth pasting condition $f''(u) = 0$.⁸ As done in Section 3.6.2, the Bellman equation can be equivalently written in terms of $u, \gamma \in \mathbb{R}$ and $v \in C^1[0, u]$. Then, u_γ and $v_\gamma : [0, u_\gamma] \rightarrow \mathbb{R}$ can be defined for each $\gamma \geq 0$ as done in Section 3.6.3. It is straightforward to characterize v_γ as done in Lemma 11, argue that γ^* is the unique solution to $v_{\gamma^*}(0) = -\kappa$, and show that the optimal policy is characterized by $(u_{\gamma^*}, \gamma^*, v_{\gamma^*})$. This two-sided barrier policy with drift rate control can then be interpreted in the context of the original control problem to obtain an effective dynamic control policy for the manufacturing system with lost sales as done above for the backlog case.

8. Note that since the lower reflecting barrier is exogenous, no smooth pasting condition is needed for it.

APPENDIX A

APPENDIX OF CHAPTER 1

A.1 Node Definitions

This section discusses the procedure used to draw the boundaries of the nodes in Figure A.3. We use the approach commonly referred to as regionalization/districting (e.g., see Duque et al. [84]). In this approach, one starts with a set of small districts and aggregates them to create nodes of appropriate size. We use 2010 census tracts¹ as districts in our analysis and use a MIP (Mixed Integer Program) to aggregate the census tracts into nodes. Census tracts are permanent statistical subdivisions of a county defined by the United States Census Bureau and used in many statistical spatial analyses that require this level of granularity. NYC consists of 2166 census tracts of various shapes and sizes.² We limit our analysis to the 488 census tracts that have a minimum of one customer per street mile per day and are not isolated from the rest of the cluster (except for JFK airport). These census tracts cover 99.42% of the pick-ups in NYC between January 2010 and December 2013.

Objective. One of the assumptions used in the matching model introduced in Section 1.3 is that customers are uniformly distributed within the node. Therefore, we would like the census tracts that are aggregated to create a node to share similar demand densities (we use the number of pick-ups as a proxy for demand). We would also like the nodes to satisfy certain spatial contiguity, shape, and demand constraints. Specifically, we would like to avoid isolated census tracts in a node, avoid long and narrow nodes, and ensure a minimum number of pick-ups in each node. Moreover, we would like the number of regions (the only notion of node size in the matching model; see Section 1.3) in the nodes to show little variation across

1. Census tracts should not be confused with the regions discussed in Section 1.3. Regions are defined in the matching section of this dissertation to model the search of taxi drivers for customers within a node. Census tracts, however, are permanent subdivisions of a county defined by the US Census Bureau that provide a stable set of geographic units for the presentation of statistical data.

2. See <https://data.cityofnewyork.us/City-Government/2010-Census-Tracts/fxpq-c8ku/data>.

the city. This prevents the size of nodes from being a factor in the relocation decisions of the taxi drivers. We aim to achieve this goal by ensuring that the street miles in a node are such that it takes a taxi between t and \bar{t} hours to search all the streets in the node.³ Due to the significant spatial variation in the average traffic speed and density of pick-ups in NYC, we divide the city into six zones; see Figure A.1 for the boundaries of the zones. These zones correspond to i) downtown/lower Manhattan (south of 14th street), ii) midtown Manhattan (14th street to 59th street), iii) central park area (59th street to about 110th street excluding central park), iv) upper Manhattan (north of 110th street to Inwood),⁴ v) Brooklyn, and vi) Queens (excluding La Guardia and JFK airports). The average traffic speed and the density of pick-ups in each zone are presented in Table A.1.

Formulation. Next, we discuss the formulation we use to aggregate the census tracts in a fixed zone into nodes whilst satisfying the aforementioned conditions. We start by introducing the notation. Then, we proceed to the formulation and discuss each constraint in detail. This formulation is conceptually similar to the formulation used in Gentry et al. [96]. Let $(\mathcal{V}^c, \mathcal{E}^c)$ denote the graph of the census tracts that we would like to aggregate into nodes, where $(i, i) \notin \mathcal{E}^c$ and for $i \neq j$, $(i, j) \in \mathcal{E}^c$ if census tracts i and j share a border (of positive length). Let A^c denote the adjacency matrix of the graph, i.e., $A_{ij}^c = 1$ if $(i, j) \in \mathcal{E}^c$ and zero otherwise.

We would like the census tracts that will be aggregated into a node to be physically close to each other, i.e., we would like to avoid long and narrow nodes if possible. Therefore, we use an approach similar to Gentry et al. [96] and choose a census tract as the center of each node and assign each census tract to the node whose center is physically closer to it. To that end, let z_i be a binary decision variable such that z_i is equal to one if census tract i

3. The lower bound ensures that nodes are sufficiently large for computational tractability of the equilibrium. The upper bound helps avoid long travel times between adjacent nodes. Long travel times on adjacent nodes create an unintended incentive for taxi drivers to stay in their current node and avoid relocation due to the opportunity cost of long travels.

4. For further information on Manhattan zones, see <https://www.nybits.com/manhattan>.

is chosen as the center of a node and zero otherwise. Let x_{ij} be a binary variable that is equal to one if census tract i is assigned to the node whose center is census tract j and let y_j be a non-negative decision variable that denotes the maximum variation in the density of pick-ups (pick-ups per mile) across the census tracts assigned to the node whose center is census tract j . If census tract j is not the center of a node, $y_j = 0$.

Denote the average number of pick-ups per period in census tract i by D_i and the street miles of this census tract by S_i . Similarly, denote the lower bound on the number of pick-ups per period in a node by \underline{D} and the lower and upper bounds on the street miles in a node by \underline{S} and \bar{S} , respectively. Let $\hat{\delta}_{ij}$ denote the distance between the geometric centers of census tracts i and j and $\tilde{\delta}_{ijk} = 1$ if $\hat{\delta}_{ij} < \hat{\delta}_{ik}$ and zero otherwise. Also, let $\bar{\delta}$ denote the maximum allowable distance between the census tracts assigned to a node. We use a $\bar{\delta}$ that is equivalent to 10 minutes of travel in the zone.

The optimal node definition is the solution to

$$\underset{x,y,z}{\text{minimize}} \sum_{j \in \mathcal{V}^c} y_j \quad (\text{P0})$$

$$\text{subject to } \sum_{i \in \mathcal{V}^c} D_i x_{ij} \geq \underline{D} z_j \quad \text{for all } j, \quad (\text{P0a})$$

$$\sum_{i \in \mathcal{V}^c} S_i x_{ij} \leq \bar{S} z_j \quad \text{for all } j, \quad (\text{P0b})$$

$$\sum_{i \in \mathcal{V}^c} S_i x_{ij} \geq \underline{S} z_j \quad \text{for all } j, \quad (\text{P0c})$$

$$\bar{\delta} \geq \hat{\delta}_{ik} (x_{ij} + x_{kj} - 1) \quad \text{for all } i, j, k, \quad (\text{P0d})$$

$$\sum_{k \in \mathcal{V}^c} A_{ki}^c x_{kj} \geq x_{ij} \quad \text{for all } i \neq j, \quad (\text{P0e})$$

$$\sum_{k \in \mathcal{V}^c} \tilde{\delta}_{ijk} x_{ik} \leq 1 - z_j \quad \text{for all } i, j, \quad (\text{P0f})$$

$$\sum_{j \in \mathcal{V}^c} x_{ij} = 1 \quad \text{for all } i, \quad (\text{P0g})$$

$$x_{ij} \leq z_j \quad \text{for all } i, j, \quad (\text{P0h})$$

$$y_j \geq \left| \frac{D_i}{S_i} - \frac{D_k}{S_k} \right| (x_{ij} + x_{kj} - 1) \quad \text{for all } i, j, k \quad (\text{P0i})$$

$$x_{ij}, z_i \in \{0, 1\}, y_j \geq 0 \quad \text{for all } i, j. \quad (\text{P0j})$$

Problem (P0) minimizes the sum of the maximum variations in the density of pick-ups in each node (a measure of the heterogeneity in the density of demand within the nodes). We will show that in the optimal solution, if census tract j is not the center of a node, $y_j = 0$.

The constraints of Problem (P0) can be categorized into three groups: constraints concerning the demand and size of the nodes, constraints concerning the assignment of the census tracts, and constraints concerning the definition of the decision variables. We start with Constraints (P0a)-(P0d) that focus on the demand and size of the nodes. Constraint (P0a) ensures that all nodes have a minimum of \underline{D} pick-ups per period. Similarly, Constraints (P0b)-(P0c) ensure that the street miles in each nodes fall between the lower bound

of \underline{S} and the upper bound of \bar{S} . Constraint (P0d) ensures that the census tracts in a node are not more than $\bar{\delta}$ miles apart. This constraint helps avoid long and narrow nodes. Note that the right hand-side of (P0d) is less than or equal to zero unless both census tracts i and k are assigned to the node whose center is census tract j , in which case the right-hand side of (P0d) is equal to $\hat{\delta}_{ik}$.

Constraints (P0e)-(P0g) focus on the assignment of the census tracts. Constraint (P0e) ensures that a census tract that is not the center of its node is connected to at least one other census tract in that node. Because, if census tract i is assigned to the node whose center is census tract j , the right-hand side of (P0e) is equal to one. Therefore, there must be at least one census tract connected to census tract i who is also assigned to the node whose center is census tract j .⁵ Constraint (P0f) ensures that census tracts are assigned to the node whose center is closest to them. Because, if census tract j is the center of a node, the right hand side of (P0f) is equal to zero. This ensures that $x_{kj} = 0$ for all nodes k in which $\tilde{\delta}_{ijk} = 1$. In other words, if census tract j is the center of a node, census tract i can not be assigned to a node whose center is further from i than census tract j ; see Gentry et al. [96] for a similar treatment. Constraints (P0g) ensures that all census tracts are assigned to a node.

Constraints (P0h)-(P0j) focus on the consistency in the definition of the decision variables. Constraint (P0h) ensures the consistency in the definitions of x_{ij} and z_i , i.e., if census tract i is assigned to the node whose center is census tract j , then $z_j = 1$. Constraint (P0i), alongside the objective function, ensures that y_j (at the optimal solution) is equal to the maximum variation in the density of pick-ups among the census tracts assigned to the node whose center is census tract j . Note that if census tract j is not the center of a node, the right-hand side of (P0i) is less than or equal to zero and $y_j = 0$ at the optimal solution. Finally, Constraint (P0j) enforces that x_{ij} and z_i are binary decision variables and y_j are

5. Note that Constraint (P0e) does not enforce spatial contiguity, rather it ensures that no census tract is isolated (unless it is the center of a node). One can enforce spatial contiguity by using the notion of contiguity order (e.g., see Duque et al. [84]). Such constraints increase the size of the MIP significantly, which has motivated Duque et al. [84] to propose an approximation algorithm to their problem. Motivated by this observation, we do not introduce such constraints in Problem (P0).

non-negative.

Inputs. In what follows, we attach a subscript l to various quantities to indicate they correspond to zone l . Note that $\underline{S}_l = \text{Average Traffic Speed in Zone } l \times \underline{t}_l$, where \underline{S}_l denotes the lower bound on the total number of miles in nodes in zone l and \underline{t}_l is the minimum time a taxi has to spend to travel all streets in a node in zone l . Similarly, $\bar{S}_l = \text{Average Traffic Speed in zone } l \times \bar{t}_l$, where \bar{S}_l and \bar{t}_l are the analogous upper bounds. Recall that D_i/S_i denotes the pick-up density of census tract i and \mathcal{V}_l^c denotes the set of census tracts in zone l for $l = 1, \dots, L$, where L denotes the number of zones. We observe that $\{D_i/S_i : i \in \mathcal{V}_l^c\}$ exhibits significant variation for all zones; see the third column of Table A.1 for their coefficient of variation. Define

$$\text{PD}_l = \frac{\sum_{i \in \mathcal{V}_l^c} D_i}{\sum_{i \in \mathcal{V}_l^c} S_i} \quad \text{for } l = 1, \dots, L$$

as the average pick-up density for zone l . There is also substantial variation in $\{\text{PD}_l : l = 1, \dots, L\}$; see the second column of Table A.1. For example, the pick-up density in midtown Manhattan is two orders of magnitude higher than pick-up density in Brooklyn. Therefore, to avoid vacuous constraints in nodes with high pick-up density, we use a lower bound \underline{D}_l (for zone l) on demand that is proportional to the average pick-up density of zone l . That is, defining the overall pick-up density $\text{PD} = \sum_i D_i / \sum_i S_i$, we set

$$\underline{D}_l = \frac{\text{PD}_l}{\text{PD}} \hat{D},$$

where \hat{D} is a tuning parameter that relates the lower bounds across zones. We use $\underline{t}_l = (1 - 0.25 \text{CV}_l) \hat{t}$ and $\bar{t}_l = (1 + 0.25 \text{CV}_l) \hat{t}$, where \hat{t} is a tuning parameter common across zones and CV_l is the coefficient of variation of pick-up density in zone l .⁶ Note that the term 0.25CV_l is used to allow variation in the pick-up densities of various tracts in a node in

6. Sensitivity analysis displays that node definitions show little sensitivity to the choice of the coefficient of CV_l .

zone l in Problem (P0); and it is equivalent to controlling the variation in pick-up densities of the various tracts in zone l .

Let $\text{LT}(\hat{t}, \hat{D})$ denote the fraction of arcs in \mathcal{E} (see Section 1.3) which take longer than one period to traverse. That is,

$$\text{LT}(\hat{t}, \hat{D}) = \frac{|\{(i, j) \in \mathcal{E} : \tau_{ij} > 1\}|}{|\mathcal{E}|}.$$

We would like to choose \hat{t} and \hat{D} such that this fraction is sufficiently small and that \hat{D} is sufficiently large to ensure considerable demand in all nodes. Next, for each fixed \hat{D} , we consider $\min_{\hat{t}} \text{LT}(\hat{t}, \hat{D})$. As depicted in Figure A.2, this shows a steady increase for $\hat{D} \leq 5.9$ and a sharp increase beyond $\hat{D} = 5.9$. Therefore, to accommodate both goals, we choose $\hat{D} = 5.9$ and $\hat{t} = \arg\min_{\hat{t}} \text{LT}(\hat{t}, 5.9) = 1.2$.

Results. The solution to Problem (P0) for $(\hat{t}, \hat{D}) = (1.2, 5.9)$ is depicted in Figure A.3. In particular, this node definition provides sufficient separation between the major hubs of the city. In Figure A.3, the nodes that accommodate the major hubs of the city are highlighted. JFK and La Guardia airports are colored in red and blue, respectively. The nodes accommodating the major railroad terminals, Penn Station and Grand Central Terminal, are colored in yellow and purple, respectively. The node accommodating the major bus terminal, Port Authority Bus Terminal, is colored in green.

Table A.1: Parameters used in Problem (P0).

	Traffic speed ¹	Average density of pick-ups ²	CV of density of pick-ups	Maximum distance within a node ($\bar{\delta}$) ³
Lower Manhattan	11.8	2.36	0.8	2.0
Midtown Manhattan	8.6	6.92	0.8	1.4
Central Park Area	10.2	4.75	0.7	1.7
Upper Manhattan	11.8	0.16	1.7	2.0
Brooklyn	11.4	0.08	1.8	1.9
Queens	12.8	0.08	1.6	2.1

¹ Miles per hour. ² Pick-ups per period per mile. ³ Miles.

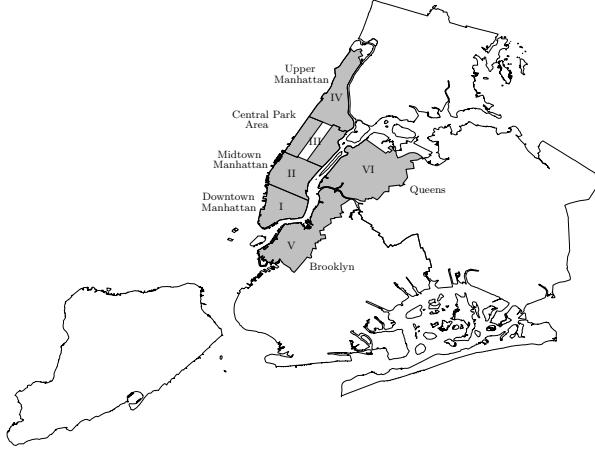


Figure A.1: Boundaries of the zones used in node generation.

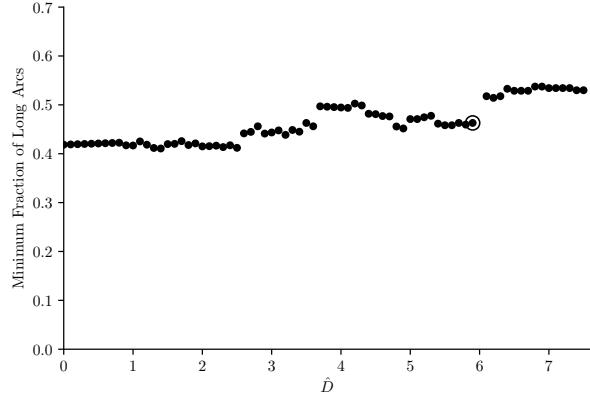


Figure A.2: Minimum achievable fraction of arcs between adjacent nodes with travel time longer than one period for given \hat{D} .

A.2 Further Discussion on the Matching Model

This section compares our matching model with that of Buchholz [56]. Buchholz uses the friction model introduced in Burdett et al. [59].⁷ Burdett et al. [59] provides a formula for friction in the following setting: Suppose that n buyers each want one unit of a good and m sellers each have one unit. First, sellers are revealed to buyers, and then each buyer chooses a seller. If more than one buyer chooses a seller, the seller can serve only one of the buyers. Burdett et al. [59] refers to this phenomena as friction and calculates the expected number of sales (matches between buyers and sellers) given n and m . Burdett et al. [59] does not allow the buyers to search for sellers if they fail in their first attempt. It also assumes that buyers know the number of sellers and their exact location (sellers are revealed to buyers).

Buchholz [56] tailors the imperfect matching model proposed in Burdett et al. [59] to the taxi market. This imperfect matching model is equivalent to the following setting: First the number of drivers and their locations are revealed to the customers. Then, each customer chooses a taxi. When more than one customer chooses a taxi, only one of the customers is served and the others leave the system unfulfilled. Therefore, given m taxis and Λ customers

7. This was initially formulated as an urn-ball problem in Hall [107], where n balls are randomly placed in m urns. In this problem, a match (only) occurs for the first ball placed in any urn.

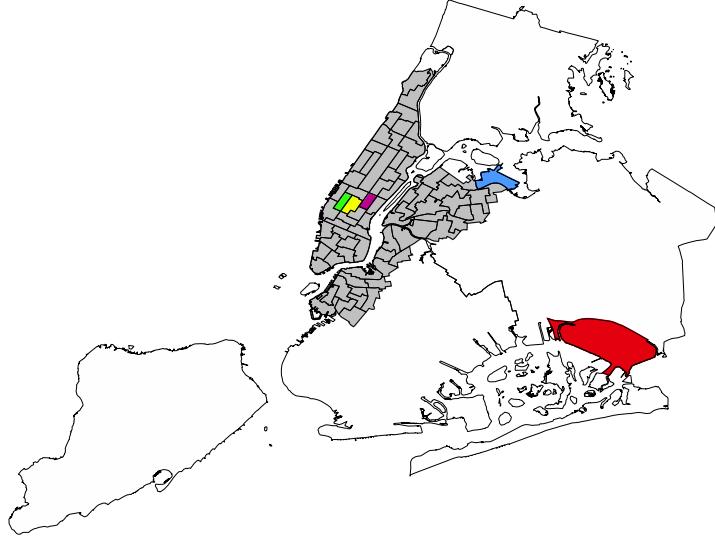


Figure A.3: The boundaries of the seventy five nodes used in the analysis and the area under consideration.

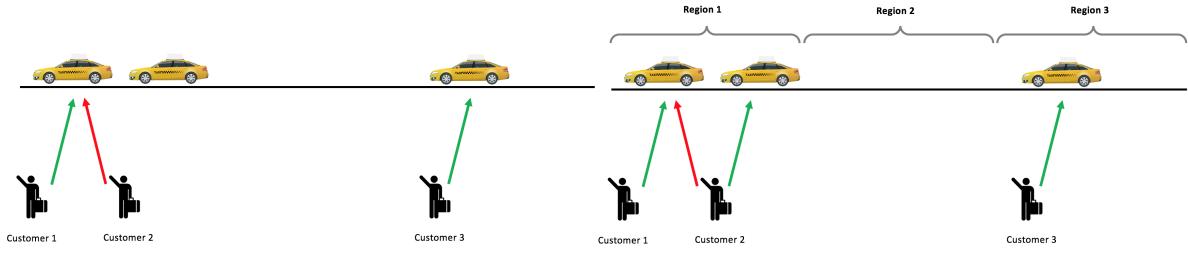
in a node, the expected number of matches (served customers) is $m(1 - [1 - 1/m]^\Lambda)$; see Burdett et al. [59] for the derivation. Buchholz introduces an efficiency parameter γ and models the expected number of matches as follows:

$$\text{Expected Number of Matches} = m\left(1 - \left[1 - \frac{1}{\gamma m}\right]^\Lambda\right).$$

Lower values of γ correspond to more efficient matching.

Figure A.4a depicts an example of matching under the model of Buchholz [56]. In this example, Customers 1 and 2 choose Taxi 1 and Customer 3 chooses Taxi 3. Taxi 1 chooses Customer 1 and Customer 2 is left without a taxi. He can not choose Taxi 2 although they are very close to each other. Figure A.4b depicts the matches in the same example under our matching model with three regions. In this case, after failure to obtain a taxi in the first attempt, Customer 2 is allowed to re-choose among the remaining empty taxis in its region. Customer 2 chooses Taxi 2 and there are no unfulfilled customers left.

Figure A.5 depicts the percentage increase in the number of matches under our matching model with $N = 10$ and $N = 20$, respectively, compared to the matching model of Buchholz



(a) Matching in the model of Buchholz [56]. (b) Matching in our model with three regions.

Figure A.4: An illustration of matching in our imperfect matching model and the model of Buchholz [56].

[56] with $\gamma = 0.84$ (as estimated by Buchholz for midtown Manhattan).⁸ We observe that for small number of taxis and customers, which is common in areas such as Brooklyn and Queens, our matching model results in fewer matches. However, for large number of taxis and customers, which is common in areas such as midtown Manhattan, our matching model results in more matches.

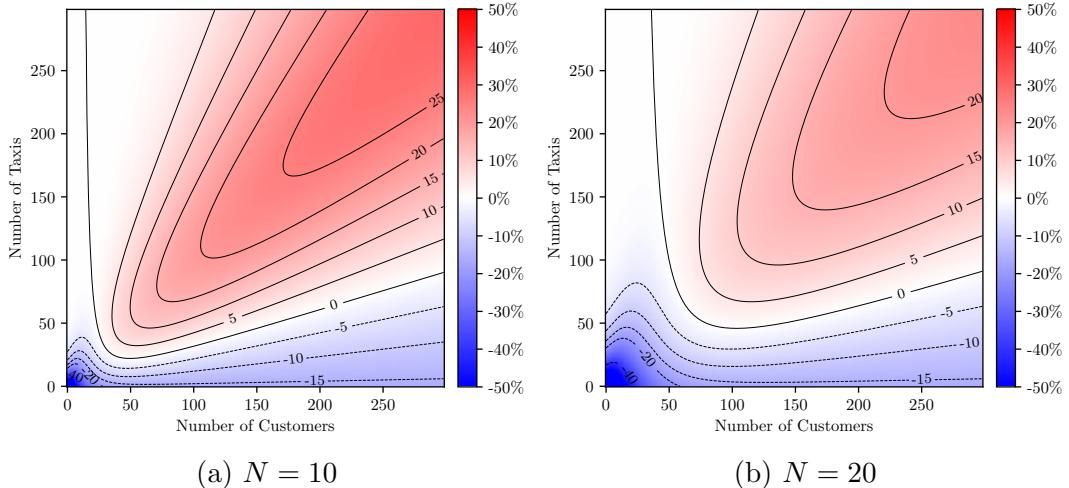


Figure A.5: Percentage change in the number of matches in our imperfect matching model compared to the imperfect matching model of Buchholz [56].

8. We use $N = 10$ and 20 as the number of regions in the overwhelming majority of the nodes in Figure A.3 fall between these numbers.

A.3 Supplemental Material for Section 1.4

The histograms of the average trip duration and distance (for rides originating at each node) are depicted in Figure A.6. Recall that a trip on average is 2.7 miles and 13.2 minutes long. There is considerable spatial variation in the average trip duration and distance across the city. Figure A.7 depicts the average trip duration and distance for rides originating at each node.

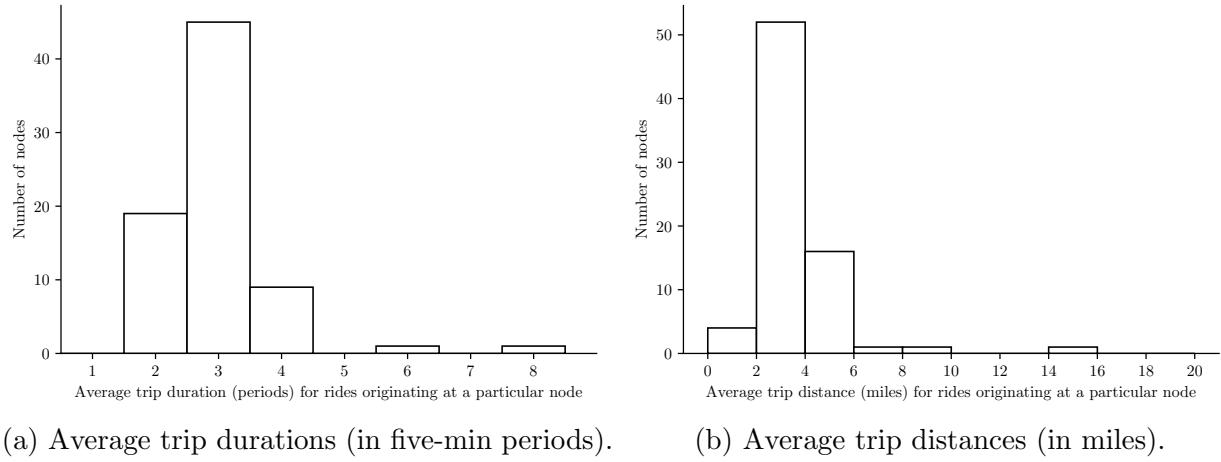


Figure A.6: Histograms of the average trip duration and distances across nodes.

A.4 Supplemental Material for Section 1.5

Table A.2 provides a summary of the model primitives. The estimation of these primitives is carried out in three steps; see Section 1.5 for further detail.

Next, we discuss the calculation of the likelihood of taxi m observing the sequence of drop-off and subsequent pick-ups on day d of month k that are no longer than two hours apart. We start by introducing some notation. Then, we calculate the likelihood of taxi m observing the entire sequence of drop-off and subsequent pick-ups on day d of month k . We conclude this section by removing the drop-off and subsequent pick-up tuples that are longer than two hours apart.

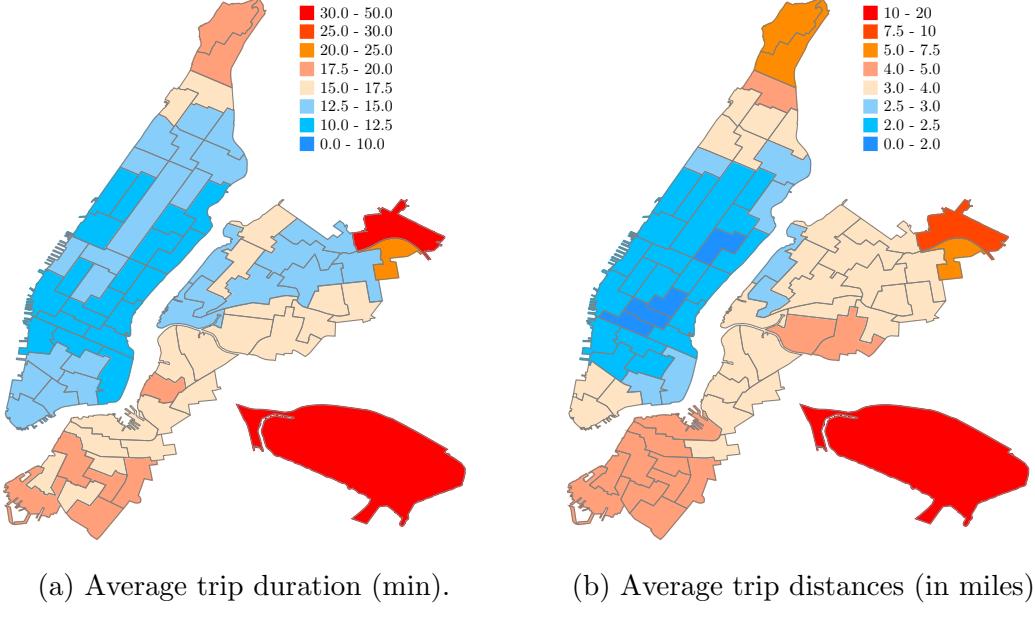


Figure A.7: Average trip duration and distance for rides originating at each node.

Let $N_{m,d,k}$ denote the number of riders that taxi m served on day d of month k . We denote the pick-up periods of taxi m on day d of month k by $t_{m,d,k}^{(1)}, \dots, t_{m,d,k}^{(l)}, \dots, t_{m,d,k}^{(N_{m,d,k})}$, where superscript l denotes the number of the pick-up. Similarly, we denote the pick-up and drop-off nodes of taxi m on day d of month k by $i_{m,d,k}^{(1)}, \dots, i_{m,d,k}^{(l)}, \dots, i_{m,d,k}^{(N_{m,d,k})}$ and $j_{m,d,k}^{(1)}, \dots, j_{m,d,k}^{(l)}, \dots, j_{m,d,k}^{(N_{m,d,k})}$, respectively. Then, the likelihood of taxi m observing the entire sequence of pick-ups on day d of month k is

$$\begin{aligned} & \mathbb{P} \left\{ (t_{m,d,k}^{(1)}, i_{m,d,k}^{(1)}, j_{m,d,k}^{(1)}), \dots, (t_{m,d,k}^{(N_{m,d,k})}, i_{m,d,k}^{(N_{m,d,k})}, j_{m,d,k}^{(N_{m,d,k})}) \right\} \\ &= \left[\prod_{l=2}^{N_{m,d,k}} \mathbb{P} \left\{ t_{m,d,k}^{(l)}, i_{m,d,k}^{(l)}, j_{m,d,k}^{(l)} \mid t_{m,d,k}^{(l-1)}, i_{m,d,k}^{(l-1)}, j_{m,d,k}^{(l-1)} \right\} \right] \\ & \quad \times \mathbb{P}_I \{ t_{m,d,k}^{(1)}, i_{m,d,k}^{(1)}, j_{m,d,k}^{(1)} \} \times \mathbb{P}_F \{ t_{m,d,k}^{(l)}, i_{m,d,k}^{(l)}, j_{m,d,k}^{(l)} \}, \end{aligned} \quad (\text{A.1})$$

where the term $\mathbb{P}_I \{ t_{m,d,k}^{(1)}, i_{m,d,k}^{(1)}, j_{m,d,k}^{(1)} \}$ is the probability that the customer picked up at period $t_{m,d,k}^{(1)}$ at node $i_{m,d,k}^{(1)}$ for destination $j_{m,d,k}^{(1)}$ is the first pick-up of the driver, the term $\mathbb{P}_F \{ t_{m,d,k}^{(l)}, i_{m,d,k}^{(l)}, j_{m,d,k}^{(l)} \}$ is the probability that the customer picked up at pe-

Table A.2: Summary of the primitives.

Parameter	Number of elements	Estimation type
α	1	Endogenous
β	1	Endogenous
σ	1	Endogenous
A_{ij}^t	$120 \times 75 \times 75$	Endogenous
S_{ij}	75×75	Exogenous
τ_{ij}	75×75	Exogenous
d_{ij}	75×75	Exogenous
m_i^1	75×48	Exogenous
N_i	75	Exogenous
M	48	Exogenous
c	48	Exogenous

riod $t_{m,d,k}^{(l)}$ at node $i_{m,d,k}^{(l)}$ for destination $j_{m,d,k}^{(l)}$ is the final pick-up of the driver, and $\mathbb{P}\{t_{m,d,k}^{(l)}, i_{m,d,k}^{(l)}, j_{m,d,k}^{(l)} | t_{m,d,k}^{(l-1)}, i_{m,d,k}^{(l-1)}, j_{m,d,k}^{(l-1)}\}$ is the probability that the driver picks up his l -th customer in period $t_{m,d,k}^{(l)}$ at node $i_{m,d,k}^{(l)}$ for destination $j_{m,d,k}^{(l)}$ given that his previous customer was picked up in period $t_{m,d,k}^{(l-1)}$ at node $i_{m,d,k}^{(l-1)}$ for destination $j_{m,d,k}^{(l-1)}$. Note that the distribution of $(t_{m,d,k}^{(l)}, i_{m,d,k}^{(l)}, j_{m,d,k}^{(l)})$ depends on all elements of $(t_{m,d,k}^{(l-1)}, i_{m,d,k}^{(l-1)}, j_{m,d,k}^{(l-1)})$. Because, all three elements of $(t_{m,d,k}^{(l-1)}, i_{m,d,k}^{(l-1)}, j_{m,d,k}^{(l-1)})$ impact the period and node at which the taxi drops its $(l-1)$ -th customer and starts the search for the l -th customer. Since it takes $\tau_{i,j}$ periods to transport a customer from node i to node j , the term

$$\mathbb{P}\{t_{m,d,k}^{(l)}, i_{m,d,k}^{(l)}, j_{m,d,k}^{(l)} | t_{m,d,k}^{(l-1)}, i_{m,d,k}^{(l-1)}, j_{m,d,k}^{(l-1)}\} \quad (\text{A.2})$$

is equal to the probability that driver m picks up his next customer in period $t_{m,d,k}^{(l)}$ at node $i_{m,d,k}^{(l)}$ for destination $j_{m,d,k}^{(l)}$ if he starts empty in period $t_{m,d,k}^{(l-1)} + \tau_{i_{m,d,k}^{(l-1)}, j_{m,d,k}^{(l-1)}}$ at node $j_{m,d,k}^{(l-1)}$.

We are interested in the likelihood of drop-off and subsequent pick-ups that are no longer than two hours apart. Therefore, we must omit drop-off and subsequent pick-up tuples that take longer than two hours. To do so, let $J(t^2, i^2, j^2 | t^1, i^1, j^1) = \mathbb{P}\{t^2, i^2, j^2 | t^1, i^1, j^1\}$ if $t^2 - t^1 - \tau_{j^1, i^2} > 24$ and one otherwise. Define J_I and J_F in a similar manner for the initial

and final pick-up. Then, the likelihood of taxi m observing the sequence of drop-off and subsequent pick-ups on day d of month k that are no longer than two hours apart is

$$\begin{aligned} \mathcal{L}_{m,d,k}(\sigma) = & \left[\prod_{l=2}^{N_{m,d,k}} J\left(t_{m,d,k}^{(l)}, i_{m,d,k}^{(l)}, j_{m,d,k}^{(l)} \mid t_{m,d,k}^{(l-1)}, i_{m,d,k}^{(l-1)}, j_{m,d,k}^{(l-1)}\right) \right] \\ & \times J_I\left(t_{m,d,k}^{(1)}, i_{m,d,k}^{(1)}, j_{m,d,k}^{(1)}\right) \times J_F\left(t_{m,d,k}^{(l)}, i_{m,d,k}^{(l)}, j_{m,d,k}^{(l)}\right). \end{aligned} \quad (\text{A.3})$$

To compute (A.2), we transform the graph of the problem to an augmented graph, in which all arcs have length one. This can be achieved by adding $n - 1$ auxiliary equi-spaced nodes on each arc of length n (for all n). Movements of an empty taxi on the augmented graph follow a non-homogeneous discrete time Markov Chain and (A.2) can be computed from the transition matrix of this Markov chain. As discussed in Section 1.5, solving (P1) using (A.3) gives an estimate of $\sigma = 1.4048$; see Figure A.8 for log likelihood values in the vicinity of the maximum likelihood estimate $\sigma = 1.4048$.

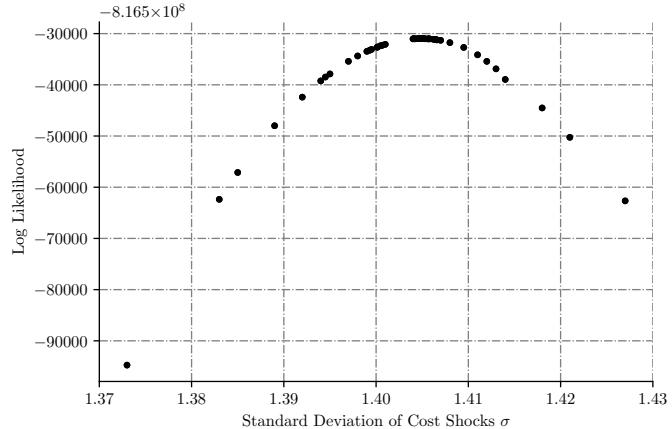


Figure A.8: Log likelihood values in the vicinity of the maximum likelihood estimate $\sigma = 1.4048$.

Next, let us illustrates through an example how an augmented graph is created and used to compute (A.2). Consider the graph in Figure A.9, where $\tau_{12} = \tau_{21} = 1$ and $\tau_{23} = \tau_{32} = 2$. The augmented graph is generated by placing node 4 on arc $(2, 3)$ and node 5 on arc $(3, 2)$, such that $\tau_{24} = \tau_{43} = \tau_{35} = \tau_{52} = 1$ and nodes 4 and 5 do not have a loop. The

augmented graph is depicted in Figure A.10. The relocation decisions of an empty taxi on

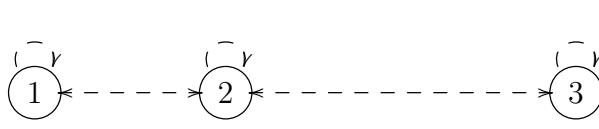


Figure A.9: Original graph.

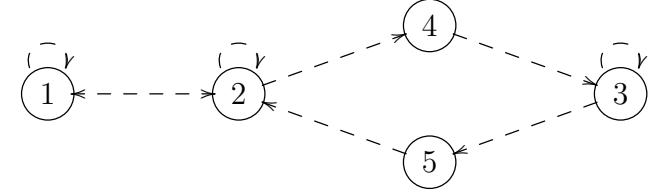


Figure A.10: Augmented graph.

the augmented graph correspond to the relocation decisions of the same taxi on the original graph. Consequently, the (path-wise) relocations of an empty taxi on the augmented graph are defined as follows:

1. If the taxi decides to relocate on an arc of unit length in the original graph, it relocates on the same arc in the augmented graph. For example, a taxi relocating from node 1 to node 2 in the original graph relocates from node 1 to node 2 in the augmented graph.
2. If the taxi decides to relocate on a long arc (arc of length longer than one period) in the original graph, it relocates to the first node created as a replacement of the long arc in the augmented graph. For example, a taxi relocating from node 2 to node 3 in the original graph relocates from node 2 to node 4 in the augmented graph.
3. If the taxi is currently on an auxiliary node in the augmented graph (which corresponds to a taxi relocating on a long arc in the original graph), it relocates on the only arc originating from the auxiliary node. For example, a taxi on node 4 of the augmented graph relocates to node 3 and a taxi on node 5 of the augmented graph relocates to node 2.

Given the aforementioned relocation policy on the augmented graph, the transition matrix

of the augmented graph in period t is

$$\tilde{Q}^t = \begin{bmatrix} q_{11}^t & q_{12}^t & 0 & 0 & 0 \\ q_{21}^t & q_{22}^t & 0 & q_{23}^t & 0 \\ 0 & 0 & q_{33}^t & 0 & q_{32}^t \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{bmatrix}, \quad (\text{A.4})$$

where q_{ij}^t are relocation probabilities on the original graph. Let B^t be a matrix whose (i, j) -th element is given by

$$B_{ij}^t = \begin{cases} (1 - \frac{\lambda_i^t}{m_i^t})\tilde{Q}_{ij}^t & \text{for } i \leq n, \\ \tilde{Q}_{ij}^t & \text{for } i > n, \end{cases}$$

and denotes the probability that an empty taxi at node i does not pick up a customer in period t and relocates to node j . Following the numbering in Figure A.10, since n is the number of nodes on the original graph, $i < n$ corresponds to a node common to the original and the augmented graph and $i > n$ corresponds to an auxiliary node on the augmented graph.

Next, let us compute the probability that the taxi driver picks up his next customer in period t_1 at node i_1 for destination j_1 if he starts empty in period t_0 at node j_0 , denoted by $\mathbb{P}\{t_1, i_1, j_1 | t_0, j_0\}$. Computing this probability is equivalent to computing (A.2).⁹ By the definition of B^t , we have

$$\mathbb{P}\{t_1, i_1, j_1 | t_0, j_0\} = \left[\prod_{t=t_0}^{t_1-1} B^t \right] \Big|_{i=j_0, j=i_1} \frac{\lambda_{i_1}^{t_1}}{m_{i_1}^{t_1}} \pi_{i_1, j_1}^{t_1}(F, P). \quad (\text{A.5})$$

The first term on the right hand side is the probability that a taxi driver who starts empty

9. Note that (A.2) is equal to the probability that driver m picks up his next customer in period $t_{m,d,k}^{(l)}$ at node $i_{m,d,k}^{(l)}$ for destination $j_{m,d,k}^{(l)}$ if he starts empty in period $t_{m,d,k}^{(l-1)} + \tau_{i_{m,d,k}^{(l-1)}, j_{m,d,k}^{(l-1)}}$ at node $j_{m,d,k}^{(l-1)}$.

at period t_0 in node j_0 unsuccessfully searches for customers until he arrives in period t_1 at node i_1 . The second term on the right hand side is the probability that the driver picks up a customer in period t_1 at node i_1 , and the third term is the probability that the destination of the customer is node j_1 .

We conclude this section by discussing the details of the equilibrium calculations. Given \bar{M} , M^k , N_i , c^k , F_{ij}^k , P_{ij}^k , $\pi_{ij}^{tk}(F^k, P^k)$, and λ_i^{tk} , for each σ , we would like to solve for the mean field equilibrium that satisfies Equations (1.4)-(1.9) and $\lambda_i^{tk} \leq m_i^{tk}$, with m_i^1 having the same distribution as the first pick-up of the taxis in the first fifteen minutes of the day shift. However, due to the potential misspecification of the initial distribution of taxis m_i^1 , such an equilibrium may not exist (in some months). Consider the following formulation:

$$\underset{m_i^1}{\text{minimize}} \quad \sum_{i=1}^n |m_i^1 - \hat{m}_i^1|^2 \quad (\text{A.6})$$

$$\text{subject to (1.4) - (1.9) and } \lambda_i^t \leq m_i^t \quad (\text{A.7})$$

In this formulation, m_i^1 denotes the initial distribution of empty taxis used in equilibrium calculations and \hat{m}_i^1 represents the desired (empirical) distribution (distribution of the first pick-up of the taxis in the first fifteen minutes of the day shift). Problem (A.6) searches among all mean-field equilibria (that satisfy constraints (A.7) given the estimated primitives) for the equilibrium whose initial distribution is closest (in the L^2 sense) to the desired (empirical) initial distribution. Numerical experiments show that the optimal objective function value of Problem (A.6) is non-zero in less than a handful of months. In other words, in the overwhelming majority of the months, there exists an equilibrium whose m_i^1 has the desired distribution.

A.5 Monte Carlo Experiments

This section uses Monte Carlo experiments to evaluate the capability of the maximum likelihood estimator described in Section 1.5 in identifying the true σ . We generate 100 simulated datasets (of pick-ups and drop-offs) assuming a true σ and estimate new σ 's from the simulated datasets. Then, we construct the 95% confidence interval and check whether the assumed true value falls within the confidence interval.

Assuming the true value of $\sigma = 0.7$ and using the estimated values of \bar{M} , M^k , N_i , m_i^{1k} , c^k , F_{ij}^k , P_{ij}^k , $\pi_{ij}^{tk}(F^k; P^k)$, and λ_i^{tk} from Section 1.5, we calculate the mean field equilibrium for each month. For simplicity, the month index k is suppressed in the remainder of this section. Using the mean field equilibria (for all months), we generate 100 simulated datasets with the same size as the original dataset (same number of months/days/active taxi drivers). To simulate the movements of the drivers, we use the following procedure for each active taxi driver on each weekday of each month.

1. Set $t = 1$ and generate a discrete r.v. with distribution $m^1 = \{m_j^1 : 1 \leq j \leq n\}$. This r.v. corresponds to the initial location of the taxi. Set i equal to this location.
2. For an empty taxi in period t at node i , generate a Bernoulli r.v. with success probability λ_i^t/m_i^t , to simulate the search of the taxi for customers.
 - (a) If the Bernoulli r.v. is equal to one (driver picking up a customer), generate r.v. X with distribution $\pi_{i,\cdot}^t = \{\pi_{i,j}^t : 1 \leq j \leq n\}$.
 - (b) If the Bernoulli r.v. is equal to zero (driver failing to pick up a customer), generate r.v. X with distribution $q_{i,\cdot}^t = \{q_{i,j}^t : j \in \mathcal{A}(i)\}$.
3. Update $i \leftarrow X$ and add $t \leftarrow t + \tau_{ij}$.
4. If $t < T$, return to Step 2.

Estimating σ from the simulated datasets results in the 95% confidence interval of (0.699, 0.702).

We observe that the (assumed) true σ falls within the confidence interval. This shows that our estimator can recover the true structural parameter σ from the data.

A.6 Cross-Validation

This section uses five-fold cross-validation to examine the ability of our model in predicting the relocation decisions of the empty taxi drivers (Kohavi et al. [139]). We randomly split the non-holiday weekdays in the dataset into five groups of roughly equal sizes. Keeping the node definitions in Figure A.3, we hold out one of the groups and use the other groups to estimate the primitives of the model (see Appendix A.4 for a summary of the primitives and Section 1.5 for the estimation procedure). The goal is to use the estimated primitives from the other groups to make predictions about the hold-out group and assess the accuracy of the prediction. The relocation probabilities would be an ideal performance metric. However, in the data, we do not observe the location of the empty taxis at all times. Therefore, we use the distribution of the next pick-up location as a proxy.

Let μ_{it}^{kl} denote the ex ante distribution of the next pick-up location for a taxi driver who dropped off his customer at node i in period t of a weekday in month k and group l . Let ν_{it}^{kl} denote the corresponding empirical quantity and w_{it}^{kl} denote the number of drop-offs at node i in period t of a weekday in month k and group l . Furthermore, let $D(\cdot, \cdot)$ be a measure of similarity between two (spatial) distributions. We take

$$\frac{\sum_{l=1}^5 \sum_{k=1}^K \sum_{i=1}^n \sum_{t=1}^T w_{it}^{kl} D(\mu_{it}^{kl}, \nu_{it}^{kl})}{\sum_{l=1}^5 \sum_{k=1}^K \sum_{i=1}^n \sum_{t=1}^T w_{it}^{kl}}, \quad (\text{A.8})$$

the weighted average¹⁰ of the similarity between the ex ante and empirical distributions of next pick-up location, as the performance metric for cross-validation. We consider three

10. We use the number of drop-offs in a period at a node as weights in the performance metric since μ_{it}^k and ν_{it}^k describe the relocation behavior of a taxi driver who just dropped off his customer. Recall that we only observe the sequence of pick-ups and drop-offs of the taxi drivers.

different similarity measures. The first similarity measure is the total variation distance

$$D_T(\mu, \nu) = \max_{E \subseteq \mathcal{E}} |\mu(E) - \nu(E)| = \frac{1}{2} \sum_{i=1}^n |\mu(i) - \nu(i)| \in [0, 1].$$

Total variation distance is the most well-known tool for quantifying the similarity between two probability distributions. It captures the largest difference in the probability assigned to sets $E \subseteq \mathcal{E}$. We also consider Bhattacharyya distance (Bhattacharyya [51])

$$D_B(\mu, \nu) = -\log \left(\sum_{i=1}^n \sqrt{\mu(i)\nu(i)} \right) \in [0, \infty]$$

that is commonly used in the information theory and pattern recognition literature; e.g., see Kailath [135], Basseville [43], and Aherne et al. [12]. Since the total variation and Bhattacharyya distances do not take advantage of the spatial aspect of the problem, we also consider earth mover's distance (also referred to as Wasserstein distance). Earth mover's distance, that is commonly used in the transportation literature, is defined as the minimum cost of turning one distribution (or pile of dirt) into the other, where the cost is assumed to be the mass moved from one node/location to the other times the distance between them (Carlsson et al. [64], Pele and Werman [173]). In other words,

$$D_E(\mu, \nu) = \inf_{\gamma \in \mathcal{P}(\mu, \nu)} \sum_{i=1}^n \sum_{j=1}^n \Delta_{ij} \gamma_{ij} \in [0, \max_{ij} \Delta_{ij}],$$

where $\mathcal{P}(\mu, \nu)$ denotes the collection of all probability measures on $\mathcal{E} \times \mathcal{E}$ with marginals μ and ν , and $[\Delta_{ij}]_{i,j \in \mathcal{E}}$ is a metric distance matrix. The distance matrix Δ must have zero diagonal entries, positive off-diagonal entries, be symmetric, and satisfy the triangle inequality. We use

$$\Delta_{ij} = \begin{cases} (d_{ij} + d_{ji})/2 & \text{for } i \neq j, \\ 0 & \text{for } i = j, \end{cases}$$

that satisfies the necessary conditions of a metric distance and it is derived from the travel distances (d_{ij}, d_{ji}) estimated from the data.

Table A.3 provides a summary of the cross-validation results. A performance metric of 0.22 under total variation distance can be interpreted as a maximum deviation of 22% between the ex ante and empirical distributions (in estimating the number of pick-ups in any set of nodes/locations). A performance metric of 0.52 under earth mover’s distance can be interpreted as an average deviation of 0.52 miles between the next pick-up location under the ex ante and empirical distributions. This performance metric is particularly meaningful when we compare it to the 2.26 miles average distance between (distinct) adjacent nodes or the 1.41 miles average distance between (distinct) non-airport adjacent nodes.

Table A.3: Cross-validation results.

Performance Metric	
Total variation distance	0.22
Bhattacharyya distance	0.08
Earth mover’s distance	0.52

A.7 Supplemental Material for Section 1.6

This section provides the mathematical formulation of the mechanisms discussed in Chapter 1 and a discussion on origin-destination spatial pricing.

A.7.1 Mathematical Formulations of the Mechanisms

This section discusses the mathematical formulation of the mechanisms introduced in Chapter 1. The mathematical formulations are introduced in the order they appear in Chapter 1. For numerical results and discussions on each mechanism, see Section 1.6.

Origin-Destination Pricing. Consider the following pricing formulation.

$$\underset{\eta_{ij}}{\text{maximize}} \quad \sum_{i=1}^n \sum_{j=1}^n \sum_{t=1}^T \text{CS}_{ij}^t(\eta_{ij}) \quad (\text{P3})$$

subject to (1.3) - (1.9) and

$$\begin{aligned} |\eta_{ij} - 1| &\leq \bar{\eta} \\ (F_{ij}, P_{ij}) &= \eta_{ij} (\bar{P}_{ij}, \bar{P}_{ij}) \\ \sum_{i=1}^n m_i^1 V(i, 0) &\geq \sum_{i=1}^n m_i^1 \bar{V}(i, 0). \end{aligned}$$

We call the solution $\{\eta_{ij}\}_{i,j \in \mathcal{V}}$ to Problem (P3) the optimal origin-destination price multipliers (or equivalently the optimal origin-destination prices) with a maximum price variation of $\bar{\eta}$. Note that in Problem (P2), price-multipliers only depend on the origin of the ride while in Problem (P3), price multipliers depend on the origin-destination pair.

Hybrid Mechanism. Let \mathcal{R} denote the set of nodes in which less than 80% of customers are served under the base prices and refer to them as under-served nodes. Consider the following pricing formulation.

$$\underset{\eta_i}{\text{maximize}} \quad \sum_{i=1}^n \sum_{j=1}^n \sum_{t=1}^T \text{CS}_{ij}^t(\eta_i) \quad (\text{P4})$$

subject to (1.4) - (1.9) and

$$(1.3b) \text{ and } |\eta_i - 1| \leq \bar{\eta} \quad \forall i \notin \mathcal{R} \quad (\text{P4a})$$

$$(1.3a) \text{ and } \eta_i = 1 \quad \forall i \in \mathcal{R} \quad (\text{P4b})$$

$$\begin{aligned} (F_{ij}, P_{ij}) &= \eta_i (\bar{P}_{ij}, \bar{P}_{ij}) \\ \sum_{i=1}^n m_i^1 V(i, 0) &\geq \sum_{i=1}^n m_i^1 \bar{V}(i, 0). \end{aligned}$$

We call the solution $\{\eta_i\}_{i \in \mathcal{V}}$ to Problem (P4) the optimal price multipliers for the hybrid mechanism (with origin-only pricing). In Problem (P4), local search friction is removed in under-served nodes and spatial pricing is used in well-served nodes. Constraint (P4a) ensures that in well-served nodes, friction is present and spatial pricing is used. In contrast, Constraint (P4b) ensures that in under-served nodes, friction is removed and spatial pricing is not used. The remaining constraints are similar to the ones in Problem (P2) in Section 1.6.1.

Proposed Mechanism. Similar to Section A.7.1, let \mathcal{R} denote the under-served nodes. Consider the following pricing formulation.

$$\underset{\eta_i}{\text{maximize}} \quad \sum_{i=1}^n \sum_{j=1}^n \sum_{t=1}^T \text{CS}_{ij}^t(\eta_i) \quad (\text{P5})$$

subject to (1.4) - (1.9) and

$$(1.3a) \text{ and } |\eta_i - 1| \leq \bar{\eta} \quad \forall i \notin \mathcal{R} \quad (\text{P5a})$$

$$(1.3a) \text{ and } \eta_i = 1 \quad \forall i \in \mathcal{R} \quad (\text{P5b})$$

$$\begin{aligned} (F_{ij}, P_{ij}) &= \eta_i (\bar{P}_{ij}, \bar{P}_{ij}) \\ \sum_{i=1}^n m_i^1 V(i, 0) &\geq \sum_{i=1}^n m_i^1 \bar{V}(i, 0). \end{aligned}$$

We call the solution $\{\eta_i\}_{i \in \mathcal{V}}$ to Problem (P5) the optimal price multipliers for the proposed mechanism (with origin-only pricing). In Problem (P5), local search friction is removed in all nodes while spatial pricing is limited to well-served nodes. Constraint (P5a) ensures that in well-served nodes, both friction removal and spatial pricing are used. In contrast, Constraint (P5b) ensures that in under-served nodes, friction is removed while spatial pricing is not used. The remaining constraints are similar to the ones in Problem (P2) in Section 1.6.1.

A.7.2 Further Discussion on Origin-Destination Pricing

Table A.4 presents the increase in consumer surplus for different maximum price variations $\bar{\eta}$ when using the optimal origin-destination pricing scheme. Similar to origin-only pricing, when using origin-destination pricing, larger values of $\bar{\eta}$ results in larger increases in consumer surplus. However, there are diminishing returns to change in $\bar{\eta}$. Origin-destination pricing with a maximum price variation of 50% results in a \$0.79 increase in consumer surplus per ride, which is 24.4% higher than the increase in consumer surplus from origin-only pricing. The gap between origin-destination pricing and origin-only pricing is increasing in the maximum price variation $\bar{\eta}$. Origin-destination pricing (with a maximum price variation of 50%) also results in a 3.2% increase in the number of served customers and an 7.4% increase in the miles traveled by customers.

Table A.4: Impact of maximum price variation when using origin-destination pricing on various performance metrics (maximum price multiplier $\hat{\eta} = 5$).

Maximum price variation ($\bar{\eta}$)	10%	20%	50%
Consumer surplus			
Total increase	\$69K	\$112K	\$168
Per ride	\$0.33	\$0.53	\$0.79
In terms of average fare	3.6%	5.8%	8.5%
Over origin-only pricing	11.3%	19.1%	24.4%
Number of served customers	1.3%	1.8%	3.2%
Miles traveled by customers	2.0%	3.9%	7.4%

A.7.3 Further Discussion on the Impact of Maximum price multiplier on optimal spatial prices

The optimal origin-only price multipliers for the maximum price multipliers of $\hat{\eta} = 5$ and $\hat{\eta} = 50$ are depicted in Figure A.11. The pattern of prices is not sensitive to $\hat{\eta}$. Therefore, in Chapter 1, we use maximum price multiplier $\hat{\eta} = 5$, which is equivalent to assuming that for a ride of \$9, no customer is willing to pay more than \$45. This is a conservative choice for $\hat{\eta}$ considering its impact on consumer surplus (see Table 1.2).

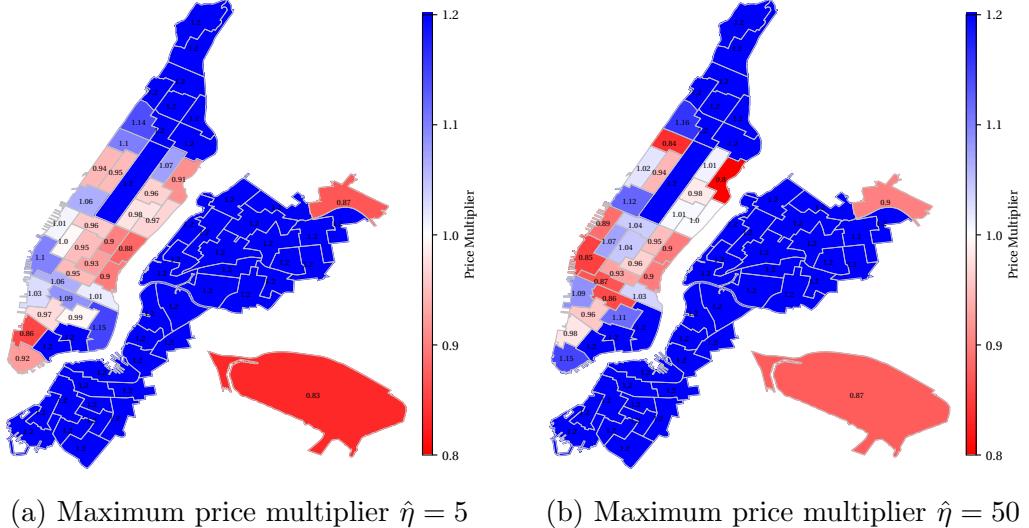


Figure A.11: Optimal origin-only price multipliers for different maximum price multipliers $\hat{\eta}$ (maximum price variation $\bar{\eta} = 0.2$).

A.8 Maximizing Drivers' Profit

This section focuses on the spatial prices that maximize drivers' profit. Consider the following pricing formulation.

$$\underset{\eta_i}{\text{maximize}} \quad \sum_{i=1}^n m_i^1 V(i, 0) \quad (\text{P6})$$

subject to (1.3) - (1.9) and

$$|\eta_i - 1| \leq \bar{\eta}$$

$$(F_{ij}, P_{ij}) = \eta_i (\bar{P}_{ij}, \bar{P}_{ij})$$

$$\sum_{i=1}^n \sum_{j=1}^n \sum_{t=1}^T \text{CS}_{ij}^t(\eta_i) \geq \sum_{i=1}^n \sum_{j=1}^n \sum_{t=1}^T \text{CS}_{ij}^t(1). \quad (\text{P6a})$$

We call the solution $\{\eta_i\}_{i \in \mathcal{V}}$ to Problem (P6) the optimal origin-only price multipliers (or equivalently the optimal origin-only prices) subject to minimum consumer surplus with a maximum price variation of $\bar{\eta}$. Note that in Problem (P6), price multipliers only depend on the origin of the ride. Constraint (P6a) ensures that the consumers surplus is not hurt by

the new pricing scheme. The remaining constraints are similar to the ones in Problem (P2) in Section 1.6.1.

The increase in drivers' profit for four plausible values of maximum price multiplier $\hat{\eta}$ is presented in Table A.5. As shown in Table A.5, using a origin-only pricing scheme with a maximum variation of 20% results in an increase in driver's profit of \$64,000 - \$144,000 on every day shift on weekdays. This is equivalent to an increase in drivers' profit of \$0.30 - \$0.71 per ride (3.4% - 7.9% of the average fare paid by the customers). Since similar to Section 1.6.1, the pattern of the optimal origin-only prices is insensitive to the maximum price multiplier $\hat{\eta}$, in the remainder of this section we use $\hat{\eta} = 5$. This is a conservative choice for $\hat{\eta}$ considering its impact on drivers' profit (see Table A.5).¹¹

Table A.5: Impact of origin-only prices on drivers' profit (maximum price variation $\bar{\eta} = 0.2$).

Maximum price multiplier ($\hat{\eta}$)	5	10	20	50
Total increase	\$64K	\$81K	\$105K	\$144K
Increase per ride	\$0.30	\$0.39	\$0.51	\$0.71
Increase in terms of average fare	3.4%	4.3%	5.6%	7.9%

Table A.6 presents the increase in drivers' profit for different maximum price variations $\bar{\eta}$. We observe that larger values of $\bar{\eta}$ results in larger increases in consumer surplus. Table A.6 indicates that the optimal origin-only pricing scheme (for maximizing drivers' profit) has minimal impact on the the number of served customers and the miles traveled by customers. This is not surprising since the objective function of Problem (P6) is drivers' profit as opposed to a measure of consumer welfare.

The optimal origin-only price multipliers (with a maximum price variation of 50%) for maximizing consumer surplus and drivers' profit are depicted in Figure A.12. The pattern of the optimal prices are similar. In upper Manhattan, Brooklyn, and Queens, prices increase since both customers and drivers benefit from the increase in prices. Higher prices in the high demand/fare areas in downtown and midtown Manhattan as well as the airports generate

11. In a similar setting, Cohen et al. [76, Section 4] makes the conservative assumption that no customer is willing to pay more than 4.9 times the base prices for an Uber ride.

Table A.6: Impact of maximum price variation on drivers' profit (maximum price multiplier $\hat{\eta} = 5$).

Maximum price variation ($\bar{\eta}$)	10%	20%	30%	40%	50%
Drivers' profit					
Total increase	\$43K	\$64K	\$79K	\$91K	\$100K
Per ride	\$0.20	\$0.30	\$0.38	\$0.43	\$0.47
In terms of average fare	2.3%	3.4%	4.2%	4.8%	5.3%
Number of served customers	-0.5%	-0.6%	-0.7%	-0.8%	-0.8%
Miles traveled by customers	0.3%	0.4%	0.4%	0.5%	0.5%

consumer surplus. This consumer surplus is then transformed into drivers' profits through higher prices in low demand/fare neighborhoods. In Figure A.12b, the increase in consumer surplus from better serving the under-served neighborhoods is transformed into drivers' profit through (slightly) higher prices compared to Figure A.12a in the well-served neighborhoods.

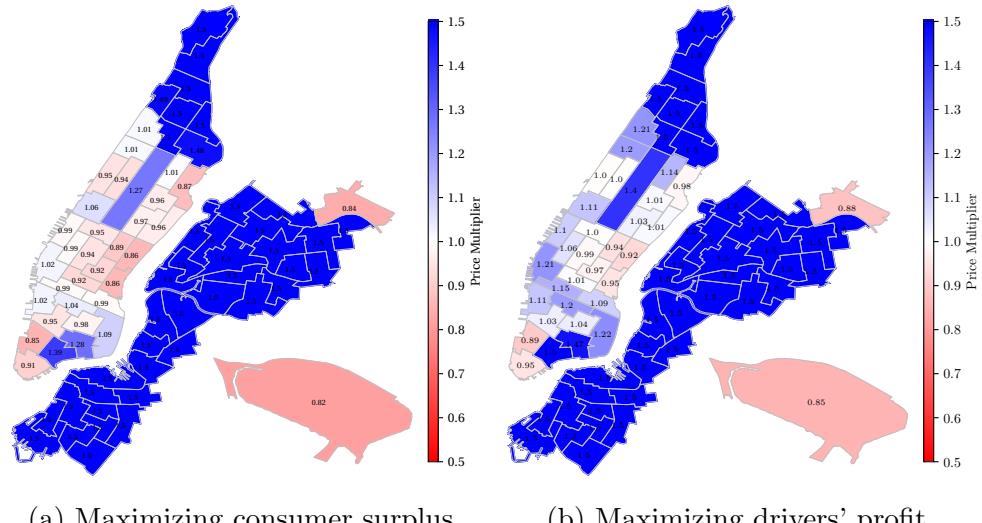


Figure A.12: Comparison of the optimal origin-only price multipliers for maximizing consumer surplus and drivers' profit (maximum price variation $\bar{\eta} = 0.5$ and maximum price multiplier $\hat{\eta} = 5$).

A.9 Miscellaneous Proofs and Derivations

This section provides a collection of proofs and derivations for the propositions, theorems, and equations discussed in Chapter 1. The proofs/derivations are discussed in the order they appear in the chapter.

Derivation of (1.3b) We start by deriving $G(\Lambda, m)$. Then, we discuss the non-monotonicity of $G(\cdot, m)$ at small values of Λ and use a linear approximation to resolve this issue. Consider node i in period t . Let us denote the de-normalized mass of empty taxis by m , the de-normalized potential demand by Λ , and the de-normalized satisfied demand by λ . Then,

$$m = \bar{M} m_i^t, \quad \Lambda = \bar{M} \Lambda_i^t(F, P), \quad \lambda = \bar{M} \lambda_i^t. \quad (\text{A.9})$$

Recall that the expected number of matches in node i is $\lambda = N_i \times \mathbb{E}[\min(X, Y)]$, where $X \sim \text{Binomial}(m, \frac{1}{N_i})$ and $Y \sim \text{Binomial}(\Lambda, \frac{1}{N_i})$. By approximating the Binomial distributions with Normal distributions, we obtain

$$\lambda \simeq N_i \times \mathbb{E}[\min(X, Y)], \quad (\text{A.10})$$

where $X \sim \mathcal{N}(\mu_m, \sigma_m)$ and $Y \sim \mathcal{N}(\mu_\Lambda, \sigma_\Lambda)$ are independent random variables with

$$\begin{aligned} \mu_m &= \frac{m}{N_i}, & \text{and} & \sigma_m^2 = \frac{m}{N_i} \left(1 - \frac{1}{N_i}\right), \\ \mu_\Lambda &= \frac{\Lambda}{N_i}, & \text{and} & \sigma_\Lambda^2 = \frac{\Lambda}{N_i} \left(1 - \frac{1}{N_i}\right). \end{aligned}$$

By substituting (A.9) into (A.10), we obtain

$$\lambda_i^t = \frac{1}{\bar{M}} \times N_i \times \mathbb{E}[\min(X, Y)], \quad (\text{A.11})$$

where $X \sim \mathcal{N}(\mu_m, \sigma_m)$ and $Y \sim \mathcal{N}(\mu_\Lambda, \sigma_\Lambda)$ are independent random variables with

$$\mu_m = \frac{\bar{M}m_i^t}{N_i} \quad \text{and} \quad \sigma_m^2 = \mu_m(1 - \frac{1}{N_i}), \quad (\text{A.12})$$

$$\mu_\Lambda = \frac{\bar{M}\Lambda_i^t}{N_i} \quad \text{and} \quad \sigma_\Lambda^2 = \mu_\Lambda(1 - \frac{1}{N_i}). \quad (\text{A.13})$$

Using *Equation (2)* of Clark [74], (A.11) simplifies to

$$\begin{aligned} \lambda_i^t &= -\frac{N_i}{\bar{M}} \left[-\mu_m \Phi\left(\frac{\mu_\Lambda - \mu_m}{\sqrt{\sigma_m^2 + \sigma_\Lambda^2}}\right) - \mu_\Lambda \Phi\left(\frac{\mu_m - \mu_\Lambda}{\sqrt{\sigma_m^2 + \sigma_\Lambda^2}}\right) + \sqrt{\sigma_m^2 + \sigma_\Lambda^2} \phi\left(\frac{\mu_\Lambda - \mu_m}{\sqrt{\sigma_m^2 + \sigma_\Lambda^2}}\right) \right] \\ &= \frac{N_i}{\bar{M}} \left[\mu_m \Phi\left(\frac{\mu_\Lambda - \mu_m}{\sqrt{\sigma_m^2 + \sigma_\Lambda^2}}\right) + \mu_\Lambda \Phi\left(\frac{\mu_m - \mu_\Lambda}{\sqrt{\sigma_m^2 + \sigma_\Lambda^2}}\right) - \sqrt{\sigma_m^2 + \sigma_\Lambda^2} \phi\left(\frac{\mu_m - \mu_\Lambda}{\sqrt{\sigma_m^2 + \sigma_\Lambda^2}}\right) \right], \end{aligned} \quad (\text{A.14})$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ are the cumulative distribution function and probability density function of the standard normal distribution. Then, by plugging $\nu \triangleq (\mu_m - \mu_\Lambda)/\sqrt{\sigma_m^2 + \sigma_\Lambda^2}$ into (A.14), we obtain

$$\begin{aligned} \lambda_i^t &= \frac{N_i}{\bar{M}} \left(\mu_m \Phi(-\nu) + \mu_\Lambda \Phi(\nu) - \sqrt{\sigma_m^2 + \sigma_\Lambda^2} \phi(\nu) \right), \\ &= \frac{N_i}{\bar{M}} \left(\mu_m \Phi(-\nu) + \mu_\Lambda \Phi(\nu) - \frac{\mu_m - \mu_\Lambda}{\nu} \phi(\nu) \right), \end{aligned} \quad (\text{A.15})$$

$$= \frac{N_i}{\bar{M}} \left(\frac{\bar{M}m_i^t}{N_i} \Phi(-\nu) + \frac{\bar{M}\Lambda_i^t}{N_i} \Phi(\nu) - \frac{\frac{\bar{M}m_i^t}{N_i} - \frac{\bar{M}\Lambda_i^t}{N_i}}{\nu} \phi(\nu) \right), \quad (\text{A.16})$$

$$= m_i^t \Phi(-\nu) + \Lambda_i^t \Phi(\nu) - \frac{m_i^t - \Lambda_i^t}{\nu} \phi(\nu), \quad (\text{A.17})$$

where (A.15) follows from the definition of ν and (A.16) follows from (A.12)-(A.13). Next, we rewrite ν in terms of m_i^t and Λ_i^t as follows:

$$\nu = \frac{(\mu_m - \mu_\Lambda)}{\sqrt{\sigma_m^2 + \sigma_\Lambda^2}} = \frac{(\mu_m - \mu_\Lambda)}{\sqrt{(1 - \frac{1}{N_i})(\mu_m + \mu_\Lambda)}}, \quad (\text{A.18})$$

$$\begin{aligned} &= \frac{\left(\frac{\bar{M}m_i^t}{N_i} - \frac{\bar{M}\Lambda_i^t}{N_i}\right)}{\sqrt{1 - \frac{1}{N_i}} \sqrt{\frac{\bar{M}m_i^t}{N_i} + \frac{\bar{M}\Lambda_i^t}{N_i}}}, \\ &= \frac{\frac{\bar{M}}{N_i} (m_i^t - \Lambda_i^t)}{\sqrt{1 - \frac{1}{N_i}} \sqrt{\frac{\bar{M}}{N_i}} \sqrt{m_i^t + \Lambda_i^t}}, \\ &= \frac{\sqrt{\frac{\bar{M}}{N_i}} \frac{m_i^t - \Lambda_i^t}{\sqrt{m_i^t + \Lambda_i^t}}}{\sqrt{1 - \frac{1}{N_i}}}, \\ &= \sqrt{\frac{\bar{M}}{N_i - 1}} \frac{m_i^t - \Lambda_i^t}{\sqrt{m_i^t + \Lambda_i^t}}, \end{aligned} \quad (\text{A.19})$$

where (A.18)-(A.19) follow from (A.12)-(A.13).

For small value of m_i^t , (A.17) is not monotone in Λ_i^t for small values of Λ_i^t . This issue follows from the inaccuracy of the normal approximation for the binomial distributions of X and Y at small values of m_i^t and Λ_i^t . We use a linear approximation to resolve the issue. In other words, we use

$$\lambda_i^t = \begin{cases} G(\Lambda_i^t, m_i^t) & \text{if } \Lambda_i^t \geq \hat{\Lambda}_{m_i^t}, \\ G(\hat{\Lambda}_{m_i^t}, m_i^t) \frac{\Lambda_i^t}{\hat{\Lambda}_{m_i^t}} & \text{otherwise,} \end{cases}$$

where $G(\Lambda, m) = \Lambda \Phi(\nu) + m \Phi(-\nu) - \frac{m - \Lambda}{\nu} \phi(\nu)$ follows from (A.17) and $\hat{\Lambda}_m = \min \{ \Lambda \geq 0 : G(\Lambda, m) \text{ is strictly increasing on } (\Lambda, \infty) \}$.

Proof of Proposition 1 First, we derive the expected payoff of the driver in each state from each action. Then, we derive the state transition matrix of the driver. We conclude

the proof by writing out the Bellman equation and simplifying it. Consider the empty infinitesimal driver at state $s = (i, t, \epsilon)$. For readability and ease of notation, in the remainder of this proof, we write $\epsilon^{(i)}$ as ϵ and ϵ_{ij} as ϵ_j . Assume that the driver chooses action $j \in \mathcal{A}(s) = \{j : S_{ij} > 0\}$. The driver picks up a customer with probability $p(s) = \frac{\lambda_i^t}{m_i^t}$ and relocates to node j with probability $1 - p(s)$. If the driver picks up a customer, his expected payoff will be $\sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - c_{il})$, where $c_{il} = c d_{il} + \epsilon_l$ is the cost of traveling from node i to node l at state $s = (i, t, \epsilon)$. If the driver does not pick up a customer, he will receive the payoff $-c_{ij} = -c d_{ij} - \epsilon_j$. Therefore, the expected (immediate) payoff of the driver at state $s = (i, t, \epsilon)$ from action $j \in \mathcal{A}(s)$ is

$$\begin{aligned} u_j(s) &\triangleq p(s) \left[\sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - c_{il}) \right] - [1 - p(s)] c_{ij} \\ &= p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - c d_{il} - \epsilon_l) - [1 - p(s)] (c d_{ij} + \epsilon_j) \end{aligned}$$

and the state transition probabilities are

$$\mathbb{P}(s' = (i', t', \epsilon') \mid s = (i, t, \epsilon), a = j) = \mathbb{P}(x' = (i', t') \mid x = (i, t), a = j) \times g(\epsilon'),$$

where

$$\mathbb{P}(x' = (i', t') \mid x = (i, t), a = j) = \begin{cases} p(s) \pi_{ii'}^t(F, P) & \text{if } i' \neq j \text{ and } t = t + d_{ii'}, \\ p(s) \pi_{ii'}^t(F, P) + (1 - p(s)) & \text{if } i' = j \text{ and } t = t + d_{ii'}, \\ 0 & \text{otherwise.} \end{cases}$$

Let $v(i, t, \epsilon)$ denote the value function of the driver at state $s = (i, t, \epsilon)$. Then, by Bellman's equation, the value function must satisfy

$$\begin{aligned} v(s) &= \max_{j \in \mathcal{A}(s)} \left\{ u_j(s) + \mathbb{E}_{s'} \left[v(s' | s, a = j) \right] \right\} \\ &= \max_{j \in \mathcal{A}(s)} \left\{ p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - c d_{il} - \epsilon_l) - [1 - p(s)] (c d_{ij} + \epsilon_j) \right. \\ &\quad \left. + \int \left(p(s) \sum_{l=1}^n \pi_{il}^t(F, P) v(l, t + \tau_{il}, \epsilon') - [1 - p(s)] v(t + \tau_{ij}, j, \xi') \right) g(\epsilon') d\epsilon' \right\}. \end{aligned}$$

Let $V(i, t) = \int v(i, t, \epsilon) d\epsilon$ denote the integrated value function of the driver at (the observable state) $x = (i, t)$. Then,

$$\begin{aligned} V(i, t) &= \int v(i, t, \epsilon) g(\epsilon) d\epsilon \\ &= \int \max_{j \in \mathcal{A}(s)} \left\{ u_j(s) + \mathbb{E}_{s'} \left[v(s' | s, a = j) \right] \right\} g(\epsilon) d\epsilon \\ &= \int \max_{j \in \mathcal{A}(s)} \left\{ p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - c d_{il} - \epsilon_l) - [1 - p(s)] (c d_{ij} + \epsilon_j) \right. \\ &\quad \left. + \int \left(p(s) \sum_{l=1}^n \pi_{il}^t(F, P) v(l, t + \tau_{il}, \epsilon') + [1 - p(s)] v(j, t + \tau_{ij}, \epsilon') \right) \right. \\ &\quad \left. g(\epsilon') d\epsilon' \right\} g(\epsilon) d\epsilon \\ &= \int \max_{j \in \mathcal{A}(s)} \left\{ p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - c d_{il} - \epsilon_l) - [1 - p(s)] (c d_{ij} + \epsilon_j) \right. \\ &\quad \left. + \int p(s) \sum_{l=1}^n \pi_{il}^t(F, P) v(l, t + \tau_{il}, \epsilon') g(\epsilon') d\epsilon' \right. \\ &\quad \left. + \int [1 - p(s)] v(j, t + \tau_{ij}, \epsilon') g(\epsilon') d\epsilon' \right\} g(\epsilon) d\epsilon, \end{aligned}$$

where in the last equality we have separated the integral over ϵ' into two parts. Then, by linearity of integration, we have

$$\begin{aligned} V(i, t) = & \int \max_{j \in \mathcal{A}(s)} \left\{ p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - cd_{il} - \epsilon_l) - [1 - p(s)] (cd_{ij} + \epsilon_j) \right. \\ & + p(s) \sum_{l=1}^n \pi_{il}^t(F, P) \int v(l, t + \tau_{il}, \epsilon') g(\epsilon') d\epsilon' \\ & \left. + [1 - p(s)] \int v(j, t + \tau_{ij}, \epsilon') g(\epsilon') d\epsilon' \right\} g(\epsilon) d\epsilon. \end{aligned}$$

Next, using the definition of the integrated value function, we obtain

$$\begin{aligned} V(i, t) = & \int \max_{j \in \mathcal{A}(s)} \left\{ p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - cd_{il} - \epsilon_l) - [1 - p(s)] (cd_{ij} + \epsilon_j) \right. \\ & + p(s) \sum_{l=1}^n \pi_{il}^t(F, P) V(l, t + \tau_{il}) + [1 - p(s)] V(j, t + \tau_{ij}) \left. \right\} g(\epsilon) d\epsilon. \end{aligned}$$

The first term in the maximum does not depend on j . Therefore, we can take it out. Thus,

$$\begin{aligned} V(i, t) = & \int \left(p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - cd_{il} - \epsilon_l) + \max_{j \in \mathcal{A}(s)} \left\{ - [1 - p(s)] (cd_{ij} + \epsilon_j) \right. \right. \\ & \left. \left. + p(s) \sum_{l=1}^n \pi_{il}^t(F, P) V(l, t + \tau_{il}) + [1 - p(s)] V(j, t + \tau_{ij}) \right\} \right) g(\epsilon) d\epsilon \\ = & \int p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - cd_{il} - \epsilon_l) g(\epsilon) d\epsilon \\ & + \int \max_{j \in \mathcal{A}(s)} \left\{ - [1 - p(s)] (cd_{ij} + \epsilon_j) \right. \\ & \left. + p(s) \sum_{l=1}^n \pi_{il}^t(F, P) V(l, t + \tau_{il}) + [1 - p(s)] V(j, t + \tau_{ij}) \right\} g(\epsilon) d\epsilon. \end{aligned}$$

The expected value of the cost shocks is zero. Thus, we can simplify the first term of the equation. Furthermore, the second term in the maximum does not depend on j and ϵ . Therefore, we can take it out of the maximization and integration. Combining these two

steps gives

$$\begin{aligned}
V(i, t) &= p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - c d_{il}) + p(s) \sum_{l=1}^n \pi_{il}^t(F, P) V(l, t + \tau_{il}) \\
&\quad + \int \max_{j \in \mathcal{A}(s)} \left\{ - [1 - p(s)] (c d_{ij} + \epsilon_j) - [1 - p(s)] V(j, t + \tau_{ij}) \right\} g(\epsilon) d\epsilon \\
&= p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - c d_{il}) + p(s) \sum_{l=1}^n \pi_{il}^t(F, P) V(l, t + \tau_{il}) \\
&\quad + [1 - p(s)] \int \max_{j \in \mathcal{A}(s)} \left\{ V(j, t + \tau_{ij}) - c d_{ij} - \epsilon_j^t \right\} g(\epsilon) d\epsilon \\
&= p(s) \sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + P_{il} d_{il} - c d_{il}) + p(s) \sum_{l=1}^n \pi_{il}^t(F, P) V(l, t + \tau_{il}) \\
&\quad + \sigma [1 - p(s)] \log \left[\sum_{j \in \mathcal{A}(s)} \exp \left(\frac{V(j, t + \tau_{ij}) - c d_{ij}}{\sigma} \right) \right] \\
&= p(s) \left(\sum_{l=1}^n \pi_{il}^t(F, P) (F_{il} + [P_{il} - c] d_{il}) + \sum_{l=1}^n \pi_{il}^t(F, P) V(l, t + \tau_{il}) \right) \\
&\quad + \sigma [1 - p(s)] \log \left[\sum_{j \in \mathcal{A}(s)} \exp \left(\frac{V(j, t + \tau_{ij}) - c d_{ij}}{\sigma} \right) \right],
\end{aligned}$$

where $p(s) = \lambda_i^t / m_i^t$, and

$$q_{ij}^t = \begin{cases} \frac{\exp \left([V(j, t + \tau_{ij}) - c d_{ij}] / \sigma \right)}{\sum_{l \in \mathcal{A}(s)} \exp \left([V(l, t + \tau_{il}) - c d_{il}] / \sigma \right)} & \text{for } j \in \mathcal{A}(s), \\ 0 & \text{otherwise.} \end{cases}$$

Proof of Theorem 1 First, we show that a mean field equilibrium is a solution $(m_i^t, V(i, t))$ to Equations (1.6) and (1.8), where λ_i^t , m_i^t , f_{ij}^t , and q_{ij}^t are characterized by Equations (1.3), (1.4), (1.5), and (1.9). Next, we use Brouwer's Fixed Point Theorem to show that there exists a solution $(m_i^t, V(i, t))$ to Equations (1.6) and (1.8).

Assume that the problem primitives \bar{M} , M , N_i , c , F_{ij} , P_{ij} , τ_{ij} , d_{ij} , S_{ij} , A_{ij}^t , α_i , β_i , k , σ , and the initial distribution of (empty) cars m_i^1 are given. By Definition 1, a mean

field equilibrium is a solution $(\lambda_i^t, m_i^t, m_{ij}^t, f_{ij}^t, q_{ij}^t, V(i, t))$ to Equations (1.3)-(1.9). Given $(m_i^t, V(i, t))$, the equilibrium values λ_i^t , m_{ij}^t , f_{ij}^t , and q_{ij}^t are uniquely determined by Equations (1.3), (1.4), (1.5), and (1.9). In fact, Equations (1.3), (1.4), (1.5), and (1.9), provide no further information besides determining the values of λ_i^t , m_{ij}^t , f_{ij}^t , and q_{ij}^t given $(m_i^t, V(i, t))$. Therefore, we can think of a mean field equilibrium as a solution $(m_i^t, V(i, t))$ to Equations (1.6) and (1.8), where λ_i^t , m_{ij}^t , f_{ij}^t , and q_{ij}^t are characterized by Equations (1.3), (1.4), (1.5), and (1.9). Note that Equation (1.7) is satisfied by any solution to Equations (1.4)-(1.6) and we need not worry about it.

Let

$$\begin{aligned} V &= \left(V(i, t) : i \in \{1, \dots, n\}, t \in \{1, \dots, T\} \right) \\ m &= \left(m_i^t : i \in \{1, \dots, n\}, t \in \{1, \dots, T\} \right) \end{aligned}$$

denote the $n \times T$ dimensional vectors of the value function and the distribution of the empty taxis. Furthermore, let

$$\begin{aligned} S_V &= \left\{ V : |V(i, t)| \leq \left(\max_{ij} \{F_{ij} + P_{ij}d_{ij}\} + \log(n) \right) \times (T - i) \right\} \\ S_m &= \{m : \|m\|_1 \leq T, m \geq 0\} \end{aligned}$$

Let $\mathcal{G} : S_m \times S_V \rightarrow S_V$ be the continuous function on the right-hand side of (1.8), i.e.,

$$\begin{aligned} \mathcal{G}(m, V)_{it} &= \frac{\lambda_i^t}{m_i^t} \left(\sum_{j=1}^n \pi_{ij}^t(F, P) \left(F_{ij} + [P_{ij} - c] d_{ij} \right) + \sum_{j=1}^n \pi_{ij}^t(F, P) V(j, t + \tau_{ij}) \right) \\ &\quad + \sigma \left(1 - \frac{\lambda_i^t}{m_i^t} \right) \log \left[\sum_{j \in \mathcal{A}(i)} \exp \left(\frac{V(j, t + \tau_{ij}) - c d_{ij}}{\sigma} \right) \right]. \end{aligned}$$

Similarly, let $\mathcal{F} : S_m \times S_V \rightarrow S_m$ be (a continuous function) such that the (i, t) -th element

of \mathcal{F} for $t = 1$ is the initial condition, m_i^1 , and (i, t) -th element of \mathcal{F} for $t > 1$ is given by

$$\mathcal{F}_{it}(m, V) = \sum_{j \in \mathcal{A}(i)} q_{ji}^{t-\tau_{ji}} (m_j^{t-\tau_{ji}} - \lambda_j^{t-\tau_{ji}}) + \sum_{j=1}^n \lambda_j^{t-\tau_{ji}} \pi_{ji}^{t-\tau_{ji}}(F, P),$$

where λ_i^t and q_{ij}^t are given by (1.3b) and (1.9). A mean field equilibrium is a solution $(m_i^t, V(i, t))$ to Equations (1.6) and (1.8), which is equivalent to a solution (m, V) to the fixed point equation

$$(m, V) = (\mathcal{F}(m, V), \mathcal{G}(m, V)) \quad (\text{A.20})$$

given the initial distribution m_i^1 and terminal value $V(i, t) = 0$ for $t > T$. Since \mathcal{F} and \mathcal{G} are continuous mappings, $(m, V) \rightarrow (\mathcal{F}(m, V), \mathcal{G}(m, V))$ is a continuous mapping from the compact convex set $S_m \times S_V$ onto itself. Therefore, by Brouwer's Fixed Point Theorem (see Royden and Fitzpatrick [180, Section 10.3]), there exists a solution to (A.20). Hence, there exists a mean field equilibrium.

Proof of Proposition 2 Let $H_{ij} = F_{ij} + P_{ij}d_{ij}$ denote the total fare paid for a ride from node i to node j . Following the notation introduced in Section 1.6, \bar{H}_{ij} denotes the fare paid under the base prices $(\bar{P}_{ij}, \bar{F}_{ij})$, and H_{ij} denotes the fare paid under (P_{ij}, F_{ij}) . By the definition of η_{ij} , it follows that $H_{ij} = \eta_{ij} \bar{H}_{ij}$ for all i, j . By Equation (1.1), the demand for rides from i to j in period t at fare level H_{ij} is equal to $\Lambda_{ij}^t(H_{ij}) = A_{ij}^t H_{ij}^\alpha$. Therefore, the consumer surplus¹² generated by all customers who want a ride from node i to node j

12. See Van Zandt [199, Page 59] for an introduction to consumer surplus.

in period t at fare level $H_{ij} < \hat{\eta} \bar{H}_{ij}$ is equal to

$$\tilde{CS}_{ij}^t(\eta_{ij}) = \int_{H_{ij}}^{\hat{\eta} \bar{H}_{ij}} \Lambda_{ij}^t(h) dh \quad (\text{A.21})$$

$$\begin{aligned} &= \int_{H_{ij}}^{\hat{\eta} \bar{H}_{ij}} A_{ij}^t h^\alpha dh \\ &= \int_{H_{ij}}^{\hat{\eta} \bar{H}_{ij}} A_{ij}^t [\eta \bar{H}_{ij}]^\alpha d(\eta \bar{H}_{ij}) \end{aligned} \quad (\text{A.22})$$

$$\begin{aligned} &= \int_1^{\hat{\eta}} A_{ij}^t \bar{H}_{ij}^{(\alpha+1)} \eta^\alpha d\eta \\ &= A_{ij}^t \bar{H}_{ij}^{(\alpha+1)} \int_1^{\hat{\eta}} \eta^\alpha d\eta \\ &= A_{ij}^t \bar{H}_{ij}^{(\alpha+1)} \frac{[\hat{\eta}^{(\alpha+1)} - \eta^{(\alpha+1)}]}{1+\alpha} \\ &= \Lambda_{ij}^t(\bar{H}_{ij}) \bar{H}_{ij} \frac{[\hat{\eta}^{(\alpha+1)} - \eta^{(\alpha+1)}]}{1+\alpha} \\ &= \Lambda_{ij}^t(\bar{F}, \bar{P}) [\bar{F}_{ij} + \bar{P}_{ij} d_{ij}] \frac{[\hat{\eta}^{(\alpha+1)} - \eta^{(\alpha+1)}]}{1+\alpha}. \end{aligned} \quad (\text{A.24})$$

Equation (A.21) follows from the definition of consumer surplus,¹³ Equations (A.22)-(A.23) use the change of variables $h = \eta \bar{H}_{ij}$, and Equation (A.24) uses the definition of demand curve and H_{ij} .

Since $\lambda_i^t / \Lambda_i^t(F, P)$ fraction of the customers who want a ride originating at node i in period t are served, the consumer surplus generated by the served customers for rides from node i to node j in period t is equal to

$$\begin{aligned} CS_{ij}^t(\eta_{ij}) &= \frac{\lambda_i^t}{\Lambda_i^t(F, P)} \tilde{CS}_{ij}^t(\eta_{ij}) \\ &= \frac{\lambda_i^t}{\Lambda_i^t(F, P)} \Lambda_{ij}^t(\bar{F}, \bar{P}) [\bar{F}_{ij} + \bar{P}_{ij} d_{ij}] \frac{[\hat{\eta}^{(\alpha+1)} - \eta^{(\alpha+1)}]}{1+\alpha}. \end{aligned}$$

13. Consumer surplus is defined as the difference between the total amount that consumers are willing to pay for a ride (indicated by the demand curve) and the total amount that they pay (fare).

APPENDIX B

APPENDIX OF CHAPTER 2

B.1 Supplemental Material

B.1.1 Derivation of the Routing Process

This section provides a derivation for the routing process introduced in Section 2.3 to describe the transition of jobs between buffers upon service completion. Let $\zeta_i^j(l)$ be a random variable such that

$$\zeta_i^j(l) = \begin{cases} 1 & \text{if the } l\text{-th job processed by dispatch activity } j \text{ becomes a class } i \text{ job next,} \\ 0 & \text{otherwise.} \end{cases}$$

Let $\zeta^j(l)$ denote the q -vector with these components. Following Harrison [110, Section 4], we call $\{\zeta^j(1), \zeta^j(2), \dots\}$ the sequence of routing indicators for dispatch activity j . We assume the routing indicators for a given class form an i.i.d sequence, and the service processes and routing indicators of various activities are mutually independent. Letting $\phi^j(l) = e^{i(j)} - \zeta^j(l)$, where e^i is the q -vector with a one in the i -th component and zeros elsewhere, define the partial sums $\Phi^j(l) = \phi^j(1) + \dots + \phi^j(l)$, with $\Phi^j(0) = 0$ (the q -vector of zeros). Then, $\phi_i^j(l)$ is the decrease in the number of class i jobs as a result of the l -th job processed by dispatch activity j and $\Phi^j(l)$ is the cumulative decrease in the number of class i jobs as a result of the first l jobs processed by dispatch activity j .

B.1.2 Additional Model Primitives for the Motivating Example

This section provides additional model primitives for the NYC application introduced in Section 2.1. We use the system parameter of $r = 8200^2$, which corresponds to the base case discussed in Section 2.9, i.e., the case with $\Lambda = \Lambda_0 = 8200$ cars; see Equation (2.65). The service rates of the nine servers in the Brownian Control Problem (BCP) are as follows:

$\mu_1 = 0.46, \mu_2 = 1.30, \mu_3 = 0.84, \mu_4 = 0.05, \mu_5 = 0.07, \mu_6 = 0.01, \mu_7 = 0.05, \mu_8 = 0.05$, and $\mu_9 = 0.03$ jobs per year; see Equation (2.13). In the BCP, we set $\tilde{\mu} = 1$. The matrix R of the BCP is as follows:

$$R = \begin{bmatrix} 0.29 & -0.17 & -0.20 & -0.04 & -0.04 & -0.04 & -0.00 & -0.02 & -0.00 & -0.00 & -0.01 & -0.00 & 1.10 & -0.20 & -0.00 & -0.00 & -0.00 & -0.01 \\ -0.24 & 0.22 & 0.49 & 0.51 & -0.33 & -0.33 & -0.02 & -0.03 & -0.01 & -0.02 & -0.02 & -0.03 & -0.81 & -0.82 & -0.02 & -0.02 & -0.01 & -0.03 \\ -0.03 & -0.03 & -0.21 & -0.42 & 0.42 & -0.42 & -0.02 & -0.00 & -0.00 & -0.01 & -0.01 & -0.01 & -0.21 & 1.10 & 0.03 & -0.01 & -0.00 & -0.01 \\ -0.00 & -0.00 & -0.01 & -0.03 & -0.03 & 0.81 & 0.04 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & -0.01 & -0.01 & -0.01 & -0.00 & -0.00 & -0.00 \\ -0.01 & -0.01 & -0.01 & -0.00 & -0.00 & -0.00 & 0.05 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & -0.01 & -0.01 & -0.00 & -0.00 & -0.00 & -0.00 \\ -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & 0.01 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & -0.05 & -0.00 \\ -0.00 & -0.00 & -0.01 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & 0.03 & -0.02 & -0.00 & -0.00 & -0.01 & -0.01 & -0.02 & 0.01 & 0.05 & -0.00 \\ -0.01 & -0.01 & -0.03 & -0.01 & -0.01 & -0.00 & -0.00 & -0.00 & -0.00 & 0.05 & 0.05 & -0.00 & -0.03 & -0.03 & -0.00 & -0.00 & -0.00 & -0.00 \\ -0.00 & -0.00 & -0.02 & -0.01 & -0.01 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & -0.00 & 0.03 & -0.02 & -0.02 & -0.00 & -0.00 & -0.00 & -0.00 \end{bmatrix} ;$$

see Equation (2.13). The capacity consumption matrix for the dispatch activities A is as follows:

$$A = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} .$$

Table B.1 provides a list of the seventy two relocation activities. The resulting 9×72 relocation activity input-output matrix \tilde{R} is defined as follows:

$$\tilde{R}^j = e^{o(j)} - e^{d(j)}, \quad j = 1, \dots, 72,$$

Table B.1: The list of the relocation activities of the NYC application.

Number	Origin	Destination	Number	Origin	Destination	Number	Origin	Destination
1	1	5	25	3	9	49	6	9
2	9	8	26	4	1	50	7	1
3	1	2	27	4	2	51	7	2
4	1	3	28	4	3	52	7	3
5	1	4	29	4	5	53	7	4
6	1	6	30	4	6	54	7	5
7	1	7	31	4	7	55	7	6
8	1	8	32	4	8	56	7	8
9	1	9	33	4	9	57	7	9
10	2	1	34	5	1	58	8	1
11	2	3	35	5	2	59	8	2
12	2	4	36	5	3	60	8	3
13	2	5	37	5	4	61	8	4
14	2	6	38	5	6	62	8	5
15	2	7	39	5	7	63	8	6
16	2	8	40	5	8	64	8	7
17	2	9	41	5	9	65	8	9
18	3	1	42	6	1	66	9	1
19	3	2	43	6	2	67	9	2
20	3	4	44	6	3	68	9	3
21	3	5	45	6	4	69	9	4
22	3	6	46	6	5	70	9	5
23	3	7	47	6	7	71	9	6
24	3	8	48	6	8	72	9	7

where \tilde{R}^j denotes the j -th column of \tilde{R} . The value vector for the dispatch activities is

$$v = (2.1 \ 1.7 \ 1.7 \ 1.5 \ 1.7 \ 1.4 \ 2.5 \ 3.3 \ 2.7 \ 2.5 \ 1.9 \ 6.3 \ 11.2 \ 1.4 \ 1.5 \ 2.3 \ 1.9 \ 2.4 \ 5.9)'$$

and the cost vector for the relocation activities is

$$c = (4.1 \ 14.7 \ 3.3 \ 6.0 \ 10.1 \ 4.5 \ 10.3 \ 15.0 \ 22.9 \ 3.5 \ 2.6 \ 6.9 \ 7.7 \ 6.2 \ 7.4 \ 12.6 \ 21.0 \ 6.9 \ 2.7 \ 3.4 \ 10.7 \ 8.0 \ 6.4 \ 11.4 \ 23.3 \ 10.3 \ 6.6 \ 3.2 \ 14.9 \ 12.2 \ 8.4 \ 10.2 \ 22.6 \ 4.1 \ 7.7 \ 10.5 \ 14.7 \ 5.0 \ 10.5 \ 13.6 \ 21.5 \ 4.2 \ 6.2 \ 7.4 \ 12.2 \ 4.8 \ 5.9 \ 9.4 \ 20.0 \ 9.4 \ 5.9 \ 5.4 \ 8.0 \ 9.6 \ 4.0 \ 5.6 \ 17.7 \ 15.7 \ 12.7 \ 11.5 \ 10.2 \ 13.7 \ 9.5 \ 5.7 \ 15.2 \ 23.3 \ 21.3 \ 23.2 \ 23.0 \ 20.8 \ 20.1 \ 18.1)'.$$

The covariance matrix of the BCP is as follows:

$$\Sigma = \begin{bmatrix} 0.57 & -0.43 & -0.07 & 0.00 & -0.03 & -0.01 & 0.00 & -0.02 & -0.01 \\ -0.43 & 1.16 & -0.54 & -0.03 & -0.03 & -0.01 & -0.02 & -0.07 & -0.03 \\ -0.07 & -0.54 & 0.71 & -0.05 & -0.01 & 0.00 & -0.01 & -0.02 & -0.01 \\ 0.00 & -0.03 & -0.05 & 0.10 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ -0.03 & -0.03 & -0.01 & 0.00 & 0.07 & 0.00 & 0.00 & 0.00 & 0.00 \\ -0.01 & -0.01 & 0.00 & 0.00 & 0.00 & 0.02 & 0.00 & 0.00 & 0.00 \\ 0.00 & -0.02 & -0.01 & 0.00 & 0.00 & 0.00 & 0.04 & -0.01 & 0.00 \\ -0.02 & -0.07 & -0.02 & 0.00 & 0.00 & 0.00 & -0.01 & 0.12 & 0.00 \\ -0.01 & -0.03 & -0.01 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.06 \end{bmatrix}.$$

Matrix K the BCP is as follows:

Next, we numerically verify that Assumption 2 holds for the NYC application, i.e.,

$$\{Ry + \tilde{R}\tilde{y} : Ky \geq 0, \tilde{y}_N \leq 0, y \in \mathbb{R}^n, \tilde{y} \in \mathbb{R}^{\tilde{n}}\} = \mathbb{R}^q. \quad (\text{B.1})$$

Since the left-hand side of (B.1) is a convex cone, it holds if

$$\pm e^i \in \{Ry + \tilde{R}\tilde{y} : Ky \geq 0, \tilde{y}_N \leq 0, y \in \mathbb{R}^n, \tilde{y} \in \mathbb{R}^{\tilde{n}}\}, \quad i = 1, \dots, q, \quad (\text{B.2})$$

where e^i is the q -vector with a one in the i -th component and zeros elsewhere. We verify (B.2) by numerically showing that there exists (y^1, \tilde{y}^1) and (y^2, \tilde{y}^2) for $i = 1 \dots, q$ such that¹

$$Ry^1 + \tilde{R}\tilde{y}^1 = -e^i, \quad Ky^1 \geq 0, \quad \tilde{y}_N^1 \leq 0, \quad y^1 \in \mathbb{R}^n, \quad \tilde{y}^1 \in \mathbb{R}^{\tilde{n}}, \quad (\text{B.3})$$

$$Ry^2 + \tilde{R}\tilde{y}^2 = +e^i, \quad Ky^2 \geq 0, \quad \tilde{y}_N^2 \leq 0, \quad y^2 \in \mathbb{R}^n, \quad \tilde{y}^2 \in \mathbb{R}^{\tilde{n}}. \quad (\text{B.4})$$

The matrices G and \tilde{G} of the Reduced Brownian Control Problem (RBCP) are as follows:²

$$G = \begin{bmatrix} 0.02 & 0.06 & 0.02 & 0.00 & 0.01 & -0.01 & -0.02 & -0.05 & -0.02 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.01 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & -0.05 & 0.01 & 0.00 \\ -0.01 & -0.06 & -0.02 & 0.00 & 0.00 & 0.00 & 0.02 & 0.05 & 0.02 & 0.00 & 0.00 & 0.00 & 0.05 & -0.01 & 0.00 \end{bmatrix}$$

and $\tilde{G} = [\tilde{G}^L \ \tilde{G}^R]$, where

$$\tilde{G}^L = \begin{bmatrix} 0 & 0 & 0 & -1 & -1 & -1 & 1 & 0 & 0 & 0 & -1 & -1 & -1 & 1 & 0 & 0 & 0 & -1 & -1 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

$$\tilde{G}^R = \begin{bmatrix} -1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & -1 & -1 & -1 & -1 & -1 & -1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & -1 & -1 & -1 & -1 & -1 & 1 & 0 & -1 & -1 & -1 & -1 & -1 & -1 & 1 & 0 & 0 & 0 \end{bmatrix}.$$

The vectors π , κ , and $\tilde{\kappa}$ of the RBCP are as follows:

$$\pi = (12.93 \ 12.62 \ 12.78 \ 12.52 \ 17.04 \ 16.05 \ 16.26 \ 15.68 \ 1.00)',$$

$$\kappa = (0.87 \ 2.3 \ 1.37 \ 0.13 \ 0.0 \ 0.01 \ 0.03 \ 0.2 \ 0.65 \ 0.78 \ 0.46 \ 0.03 \ 0.02 \ 0.01 \ 0.05)',$$

1. The existence of (y^1, \tilde{y}^1) and (y^2, \tilde{y}^2) satisfying (B.3)-(B.4) can be numerically verified using a linear program, where the objective is zero and the constraints are (B.3)-(B.4).

2. This is one particular choice of the matrix G . As discussed in Section 2.5, any solution to Equation (2.35) is acceptable.

$$\tilde{\kappa} = (3.61 \ 6.18 \ 10.54 \ 1.37 \ 6.95 \ 12.28 \ 34.82 \ 3.24 \ 2.42 \ 7.04 \ 3.26 \ 2.8 \ 3.73 \ 9.56 \ 32.61 \ 6.74 \ 2.85 \ 3.68 \ 6.4 \ 4.73 \ 2.89 \ 8.52 \ 35.05 \ 9.9 \ 6.47 \ 2.96 \ 10.41 \ 8.69 \ 4.71 \ 7.07 \ 34.12 \ 8.20 \ 12.16 \ 14.75 \ 19.24 \ 5.95 \ 11.26 \ 14.98 \ 37.53 \ 7.33 \ 9.63 \ 10.7 \ 15.74 \ 3.78 \ 5.72 \ 9.79 \ 35.06 \ 12.69 \ 9.53 \ 8.9 \ 11.71 \ 8.82 \ 4.19 \ 6.13 \ 32.98 \ 18.4 \ 15.75 \ 14.4 \ 13.38 \ 12.34 \ 9.13 \ 5.12 \ 29.87 \ 11.36 \ 9.65 \ 11.41 \ 11.47 \ 4.76 \ 5.08 \ 2.83)'$$

B.1.3 No Opportunity for Arbitrage

Proposition 8 shows that there is no opportunity for arbitrage in the BCP (2.32). To be specific, it states that the platform cannot reduce costs by deviating from the nominal plan without creating a queue length displacement, idling servers, or taking part in non-basic dispatch or relocation activities.

Proposition 8 (Harrison [114, Proposition 1]). *There is no opportunity for arbitrage, i.e.,*

$$\underset{y, \tilde{y}}{\text{minimize}} \ v'y - c'\tilde{y} \text{ subject to } Ry + \tilde{R}\tilde{y} = 0, Ky \geq 0, \text{ and } \tilde{y}_N \leq 0 \quad (\text{B.5})$$

has the unique optimal solution of $y^* = 0$ and $\tilde{y}^* = 0$.

To see the relationship between Problem (B.5) and the opportunity for arbitrage, assume that there exists (y, \tilde{y}) such that $Ry + \tilde{R}\tilde{y} = 0$, $Ky \geq 0$, $\tilde{y}_N \leq 0$, and $v'y - c'\tilde{y} < 0$. Consider a control (Y, \tilde{Y}) that has jumps $Y(t) - Y(t^-) = y$ and $\tilde{Y}(t) - \tilde{Y}(t^-) = \tilde{y}$ at some fixed time $t > 0$. From Equation (2.22) and equality $Ry + \tilde{R}\tilde{y} = 0$, we have that $Z(t) - Z(t^-) = R(Y(t) - Y(t^-)) + \tilde{R}(\tilde{Y}(t) - \tilde{Y}(t^-)) = 0$, which means that the jumps in Y and \tilde{Y} do not violate the state space constraint (2.26). Also, Equations (2.23)-(2.24) and inequalities $Ky \geq 0$, and $\tilde{y}_N \leq 0$ ensure that Equation (2.25) holds. Finally, inequality $v'y - c'\tilde{y} < 0$ states that the control jumps of Y and \tilde{Y} generate a strictly negative jump in the cumulative cost so by considering arbitrarily large multiples of y and \tilde{y} as control jumps, we can drive the cost to $-\infty$.

B.1.4 A Simplification of the Reduced Brownian Formulation

The reduced Brownian formulation introduced in Harrison and Williams [123] includes the extra condition that $U(t) \in \{Ky : y \in \mathbb{R}^n\}$ for $t \geq 0$. This condition is necessary for ensuring that there exists an admissible control (Y, \tilde{Y}) to (2.22)-(2.27) whose extended state process includes U . Lemma 12 shows that $\{Ky : y \in \mathbb{R}^n\} = \mathbb{R}^{m+n-b}$ in our case. Therefore, we can omit this condition; see Pesic and Williams [174, Lemma 3.1] for a similar result in the context of parallel server systems. The crucial feature of our model that gives rise to this property is the fact that each dispatch activity is associated with a unique server.

Lemma 12. *We have that $\{Ky : y \in \mathbb{R}^n\} = \mathbb{R}^{m+n-b}$.*

Proof. The proof resembles the proof of Pesic and Williams [174, Lemma 3.1]. For simplicity, we use the notation $\text{Range}(K)$ to denote $\{Ky : y \in \mathbb{R}^n\}$. It is sufficient to show that $e^l \in \text{Range}(K)$ for $l = 1, \dots, m+n-b$. We show this for $l = 1, \dots, m$ and $l = m+1, \dots, m+n-b$, separately.

Case 1. Consider $l \in \{1, \dots, m\}$. By Assumption 1, $Ax^* = e$. Therefore, there exists a basic dispatch activity j such that $k(j) = l$. Then, by Equations (2.3), (2.19), and (2.28), we have $K^j = e^l$, which gives $e^l \in \text{Range}(K)$.

Case 2. Consider $l \in \{m+1, \dots, m+n-b\}$ and the non-basic dispatch activity $j = l - m + b$. It follows from Equations (2.3), (2.19), and (2.28) that $K^j = e^{k(j)} - e^l$. Since $e^{k(j)} \in \text{Range}(K)$, by Case 1, and $K^j \in \text{Range}(K)$, we have $e^{k(j)} - K^j = e^l \in \text{Range}(K)$. \square

B.1.5 Equivalence of the BCP and the RBCP

Lemma 13 establishes an equivalence between the BCP (2.32) and the RBCP (2.42). In particular, it shows that given the state $Z(t)$ in the BCP, one sets $W(t) = MZ(t)$ to arrive at the equivalent state descriptor in the RBCP. To go in the other direction, given the state $W(t)$ in the RBCP, one sets $Z(t) = \Delta(W(t))$ to arrive at the equivalent state descriptor in the BCP.

Lemma 13 (Harrison and Williams [123, Theorem 6.1]). *The RBCP (2.42) is equivalent to the BCP (2.32) in the following sense: Fix $\epsilon \geq 0$.*

- (i) *Suppose that (Y, \tilde{Y}) is an ϵ -optimal control for the BCP (2.32) with Brownian motion X and extended state process (Z, U, \tilde{U}) . Then (U, \tilde{U}) is an ϵ -optimal control for the RBCP (2.42) with Brownian motion X and state process W given by $X(t) = MX(t)$ and $W(t) = MZ(t)$ for $t \geq 0$.*
- (ii) *Suppose that (U, \tilde{U}) is an ϵ -optimal control for the RBCP (2.42) with Brownian motion X and state process W . Then, there exists a Brownian motion X such that $MX(t) = X(t)$ for $t \geq 0$. Furthermore, there exists a pair (Y, \tilde{Y}) that is an ϵ -optimal control for the BCP (2.32) with Brownian motion X and extended state process (Z, U, \tilde{U}) , where $Z(t) = \Delta(W(t))$ for $t \geq 0$.*

B.1.6 The Value Function of a Region of Inaction Type Policy.

Consider a closed set \mathcal{B} and let $\partial\mathcal{B}^1, \dots, \partial\mathcal{B}^d$ be a partition of its boundary such that pushing control l is used for $w \in \partial\mathcal{B}^l$. Given the initial workload $w \in \hat{\mathcal{W}}$, consider the following control policy, denoted by \hat{U}_w : Exert no control in the interior of \mathcal{B} and keep the state process \hat{W} in \mathcal{B} by using pushing control l on $\partial\mathcal{B}^l$ minimally for $l = 1, \dots, d$. Moreover, if $w \notin \mathcal{B}$, move the state process \hat{W} instantaneously from w to the point $w' \in \partial\mathcal{B}$ for which the cost associated with this instantaneous move is the smallest; see Equations (2.57)-(2.59). On the other hand, consider a function $f \in C^2(\hat{\mathcal{W}})$ and assume that it solves the following PDE on \mathcal{B} :

$$\mathcal{L}f(w) - \alpha f(w) + h(w) = 0, \quad w \in \text{int}(\mathcal{B}) \quad (\text{B.6})$$

subject to the boundary condition

$$\nabla f(w)' \hat{G}^l + \hat{\kappa}_l = 0, \quad w \in \partial\mathcal{B}^l \text{ and } l = 1, \dots, d, \quad (\text{B.7})$$

and that it is defined on $\hat{\mathcal{W}} \setminus \mathcal{B}$ as follows:

$$f(w) = \min_{w' \in \partial\mathcal{B}} \{C(w, w') + f(w')\}, \quad w \in \hat{\mathcal{W}} \setminus \mathcal{B}, \quad (\text{B.8})$$

where C is given by (2.57)-(2.59). Next, we show that $f(w)$ is the expected discounted cost of policy \hat{U}_w that starts at state w . To this end, note that by Lemma 18, we have

$$f(w) = -\mathbb{E} \left[\int_0^\infty e^{-\alpha t} (\mathcal{L}f(w) - \alpha f(w)) dt \right] - \mathbb{E} \left[\sum_{l=1}^d \int_0^\infty e^{-\alpha t} \nabla f(\hat{W}(t))' \hat{G}^l d\hat{U}_{w,l}(t) \right] \quad (\text{B.9})$$

for $w \in \hat{\mathcal{W}}$. Substituting Equations (B.6)-(B.8) into Equation (B.9) gives the desired result:

$$f(w) = \mathbb{E} \left[\int_0^\infty e^{-\alpha t} h(w) dt \right] + \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(\hat{\kappa}' \hat{U}_w)(t) \right], \quad w \in \hat{\mathcal{W}}.$$

B.1.7 Numerical Solution of the Partial Differential Equation

This section develops a numerical method for solving the PDE discussed in Section 2.7. To be specific, suppose we are given a closed convex set $\mathcal{B} \subseteq \hat{\mathcal{W}}$, referred to as the region of inaction, and a partition $\partial\mathcal{B}^1, \dots, \partial\mathcal{B}^d$ of its boundary, $\partial\mathcal{B}$, such that pushing control l is used for $w \in \partial\mathcal{B}^l$. This section develops a numerical solution to the PDE

$$\mathcal{L}f(w) - \alpha f(w) + h(w) = 0, \quad w \in \text{int}(\mathcal{B}) \quad (\text{B.10})$$

subject to the boundary condition

$$\nabla f(w)' \hat{G}^l + \hat{\kappa}_l = 0, \quad w \in \partial\mathcal{B}^l \text{ and } l = 1, \dots, d. \quad (\text{B.11})$$

This PDE falls under the category of second-order linear elliptic PDE with oblique derivative boundary conditions; see Lieberman [156] and Gilbarg and Trudinger [101] for an overview. One can use any numerical method of choice to solve Equations (B.10)-(B.11). The specific

method used does not impact the algorithm discussed in Section 2.7. Following Kumar and Muthuraman [145], we choose to adapt the Finite Element Method (FEM) for two reasons: First, it provides an approximate solution with bounded error. Second, it can handle complex and irregular boundaries in high dimensions. Note that since irregular regions of inaction may arise in the algorithm discussed in Section 2.7, it is crucial to use a numerical method that can handle irregular boundaries; see Hughes [132] and Larson and Bengzon [151] for further details on FEM. For simplicity, in this section, we reuse some of the notation that was already introduced.

FEM transforms Equations (B.10)-(B.11) into a variational form and uses the notion of orthogonality in function spaces to obtain an approximate solution. To be specific, FEM constructs an approximate solution of the form

$$\hat{f}(w) = \sum_{m=1}^M \beta_m \psi^m(w), \quad w \in \mathcal{B}, \quad (\text{B.12})$$

where ψ^m are real-valued functions commonly referred to as the basis function, M denotes the number of (distinct) basis functions, and β_m are coefficients. The choice of the basis functions will be discussed later and for the moment, it is assumed the basis functions are given. FEM seeks to find the set of coefficients β_m for $m = 1, \dots, M$ that minimize an error criterion. We use the so-called Galerkin method, which uses a specific error criterion. Letting $\beta = (\beta_m)$ denote the vector of coefficients, we define the residual function as

$$\mathcal{R}(w, \beta) = \mathcal{L}\hat{f}(w) - \alpha\hat{f}(w) + h(w) \quad (\text{B.13})$$

at all points $w \in \text{int}(\mathcal{B})$ in which the derivatives exist. Ideally, we would like to ensure that Equation (B.10) holds, which is equivalent to ensuring that $\mathcal{R}(w, \beta) = 0$ for all $w \in \text{int}(\mathcal{B})$.

However, this need not be possible. The Galerkin method instead aims for the following:

$$\int_{\mathcal{B}} \mathcal{R}(w, \beta) \psi^k(w) dw = 0, \quad k = 1, \dots, M. \quad (\text{B.14})$$

Loosely speaking, this is equivalent to requiring the error to be orthogonal to the basis functions. Similarly, instead of requiring Equation (B.11) to hold for $w \in \partial\mathcal{B}$, the Galerkin method aims for

$$\sum_{l=1}^d \int_{\partial\mathcal{B}^l} (\nabla \hat{f}(w)' \hat{G}^l + \hat{\kappa}_l) \psi^k(w) dw = 0, \quad k = 1, \dots, M. \quad (\text{B.15})$$

Lemma 14 provides an integration by parts result that is crucially used in the Galerkin method; see Le Dret and Lucquin [154, Proposition 3.15 and Section 4.4] for further details. As a preliminary to Lemma 14, let $n(w)$ denote the normal unit exterior vector (the unit vector normal to the boundary in the exterior direction) for $w \in \partial\mathcal{B}$.

Lemma 14 (Green's Formula). *Suppose $f, \phi \in L^2(\mathcal{B})$ such that $\partial_i f, \partial_{ij} f \in L^2(\mathcal{B})$ and $\partial_i \phi \in L^2(\mathcal{B})$ for $i, j = 1, \dots, p-1$. Then,*

$$\begin{aligned} \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_{ij} f(w) \psi(w) dw &= - \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j f(w) \partial_i \psi(w) dw \\ &\quad + \int_{\partial\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j f(w) n_i(w) \psi(w) dw \end{aligned} \quad (\text{B.16})$$

By substituting Equation (B.12) into Equation (B.14) and using Lemma 14 as well as Equation (B.15), Lemma 15 obtains a set of linear equations. As a preliminary to stating

Lemma 15, we introduce the following notation: Define D_{km} and q_k as follows:

$$\begin{aligned} D_{km} = & \sum_{i,j=1}^{p-1} \frac{1}{2} \hat{\sigma}_{ij} \int_{\mathcal{B}} \partial_j \psi^m(w) \partial_i \psi^k(w) dw + \alpha \int_{\mathcal{B}} \psi^m(w) \psi^k(w) dw \\ & - \sum_{i,j=1}^{p-1} \frac{1}{2} \hat{\sigma}_{ij} \int_{\partial \mathcal{B}} \partial_j \psi^m(w) n_i(w) \psi^k(w) dw - \sum_{l=1}^d \int_{\partial \mathcal{B}^l} \nabla \psi^m(w)' \hat{G}^l \psi^k(w) dw, \quad (\text{B.17}) \end{aligned}$$

where $n_i(w)$ denotes the i -th component of $n(w)$ and

$$q_k = \int_{\mathcal{B}} h(w) \psi^k(w) dw + \sum_{l=1}^d \int_{\partial \mathcal{B}^l} \hat{\kappa}_l \psi^k(w) dw \quad (\text{B.18})$$

for $k, m = 1, \dots, M$. Then, we define the $M \times M$ dimensional matrix $D = [D_{km}]$ and the M -dimensional vector $q = (q_k)$.

Lemma 15. *If the coefficient vector β satisfies (B.14)-(B.15), then $D\beta = q$.*

Proof. First, we use Lemma 14 to write the left-hand side of (B.14) in terms of first order derivatives. Then, we combine (B.14) and (B.15) to arrive at the desired result. It follows from Equation (B.13) and the definition of operator \mathcal{L} that for $k = 1, \dots, M$,

$$\begin{aligned} \int_{\mathcal{B}} \mathcal{R}(w, \beta) \psi^k(w) dw &= \int_{\mathcal{B}} \mathcal{L} \hat{f}(w) \psi^k(w) dw - \int_{\mathcal{B}} \alpha \hat{f}(w) \psi^k(w) dw + \int_{\mathcal{B}} h(w) \psi^k(w) dw \\ &= \frac{1}{2} \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_{ij} \hat{f}(w) \psi^k(w) dw - \int_{\mathcal{B}} \alpha \hat{f}(w) \psi^k(w) dw \\ &\quad + \int_{\mathcal{B}} h(w) \psi^k(w) dw. \quad (\text{B.19}) \end{aligned}$$

Substituting (B.16) into (B.19) gives

$$\begin{aligned}
\int_{\mathcal{B}} \mathcal{R}(w, \beta) \psi^k(w) dw &= -\frac{1}{2} \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j \hat{f}(w) \partial_i \psi^k(w) dw \\
&\quad + \frac{1}{2} \int_{\partial \mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j \hat{f}(w) n_i(w) \psi^k(w) dw \\
&\quad - \int_{\mathcal{B}} \alpha \hat{f}(w) \psi^k(w) dw + \int_{\mathcal{B}} h(w) \psi^k(w) dw, \quad k = 1, \dots, M. \quad (\text{B.20})
\end{aligned}$$

Substituting (B.20) into (B.14) gives

$$\begin{aligned}
&-\frac{1}{2} \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j \hat{f}(w) \partial_i \psi^k(w) dw + \frac{1}{2} \int_{\partial \mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j \hat{f}(w) n_i(w) \psi^k(w) dw \\
&\quad - \int_{\mathcal{B}} \alpha \hat{f}(w) \psi^k(w) dw + \int_{\mathcal{B}} h(w) \psi^k(w) dw = 0, \quad k = 1, \dots, M. \quad (\text{B.21})
\end{aligned}$$

Adding the left-hand side of (B.15) to the left-hand side of (B.21), the right-hand side of (B.15) to the right-hand side of (B.21) gives

$$\begin{aligned}
&-\frac{1}{2} \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j \hat{f}(w) \partial_i \psi^k(w) dw + \frac{1}{2} \int_{\partial \mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j \hat{f}(w) n_i(w) \psi^k(w) dw \\
&\quad - \int_{\mathcal{B}} \alpha \hat{f}(w) \psi^k(w) dw + \int_{\mathcal{B}} h(w) \psi^k(w) dw \\
&\quad + \sum_{l=1}^d \int_{\partial \mathcal{B}^l} (\nabla \hat{f}(w)' \hat{G}^l + \hat{\kappa}_l) \psi^k(w) dw = 0, \quad k = 1, \dots, M.
\end{aligned}$$

Then, rearranging the terms gives

$$\begin{aligned}
& \frac{1}{2} \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j \hat{f}(w) \partial_i \psi^k(w) dw + \int_{\mathcal{B}} \alpha \hat{f}(w) \psi^k(w) dw \\
& - \frac{1}{2} \int_{\partial\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j \hat{f}(w) n_i(w) \psi^k(w) dw - \sum_{l=1}^d \int_{\partial\mathcal{B}^l} \nabla \hat{f}(w)' \hat{G}^l \psi^k(w) dw \\
& = \int_{\mathcal{B}} h(w) \psi^k(w) dw + \sum_{l=1}^d \int_{\partial\mathcal{B}^l} \hat{\kappa}_l \psi^k(w) dw, \quad k = 1, \dots, M. \tag{B.22}
\end{aligned}$$

Substituting Equation (B.12) into Equation (B.22) gives

$$\begin{aligned}
& \frac{1}{2} \sum_{m=1}^M \sum_{i,j=1}^{p-1} \int_{\mathcal{B}} \hat{\sigma}_{ij} \beta_m \partial_j \psi^m(w) \partial_i \psi^k(w) dw + \sum_{m=1}^M \int_{\mathcal{B}} \alpha \beta_m \psi^m(w) \psi^k(w) dw \\
& - \frac{1}{2} \sum_{m=1}^M \sum_{i,j=1}^{p-1} \int_{\partial\mathcal{B}} \hat{\sigma}_{ij} \beta_m \partial_j \psi^m(w) n_i(w) \psi^k(w) dw - \sum_{m=1}^M \sum_{l=1}^d \int_{\partial\mathcal{B}^l} \beta_m \nabla \psi^m(w)' \hat{G}^l \psi^k(w) dw \\
& = \int_{\mathcal{B}} h(w) \psi^k(w) dw + \sum_{l=1}^d \int_{\partial\mathcal{B}^l} \hat{\kappa}_l \psi^k(w) dw, \quad k = 1, \dots, M. \tag{B.23}
\end{aligned}$$

Finally, we write (B.23) as $D\beta = q$ using the definitions of D and q ; see Equations (B.17)-(B.18). \square

Motivated by Lemma 15, the Galerkin method seeks to find the coefficient vector β that solves $D\beta = q$. Before discussing the choice of the basis functions, let us show that the approximation error of FEM is bounded. To do so, let \mathcal{F} denote the space spanned by the basis functions ψ^1, \dots, ψ^M and define the symmetric bilinear form $b(\cdot, \cdot)$ as follows:

$$\begin{aligned}
b(\psi, \eta) &= \sum_{i,j=1}^{p-1} \frac{1}{2} \hat{\sigma}_{ij} \int_{\mathcal{B}} \partial_j \psi(w) \partial_i \eta(w) dw + \alpha \int_{\mathcal{B}} \psi(w) \eta(w) dw \\
& - \sum_{i,j=1}^{p-1} \frac{1}{2} \hat{\sigma}_{ij} \int_{\partial\mathcal{B}} \partial_j \psi(w) n_i(w) \eta(w) dw - \sum_{l=1}^d \int_{\partial\mathcal{B}^l} \nabla \psi(w)' \hat{G}^l \eta(w) dw
\end{aligned}$$

for ψ and η that are twice continuously differentiable almost everywhere on $\hat{\mathcal{W}}$. This bilinear form defines a seminorm on the space of functions that are twice continuously differentiable almost everywhere on $\hat{\mathcal{W}}$.³ Proposition 9, which extends Kumar and Muthuraman [145, Proposition 5] to our case, shows that the function \hat{f} characterized by the the solution to $D\beta = q$ is the "best approximation" to the solution to Equations (B.10)-(B.11) in the space \mathcal{F} as measured by the seminorm defined by b .

Proposition 9. *Assume f satisfies (B.10)-(B.11) and \hat{f} is characterized by the solution to $D\beta = q$. Then,*

$$b(f - \hat{f}, f - \hat{f}) \leq b(f - \psi, f - \psi), \quad \psi \in \mathcal{F}. \quad (\text{B.24})$$

Proof. The proof closely resembles the proof of Kumar and Muthuraman [145, Proposition 5]. First, we assume $b(f - \hat{f}, \psi) = 0$ for $\psi \in \mathcal{F}$ and use the bilinearity of b to arrive at (B.24). Then, we prove $b(f - \hat{f}, \psi) = 0$ for $\psi \in \mathcal{F}$. Assume we have $b(f - \hat{f}, \psi) = 0$ for $\psi \in \mathcal{F}$. Then, by the bilinearity of b , for $\eta \in \mathcal{F}$, we have

$$\begin{aligned} b(f - \hat{f} + \eta, f - \hat{f} + \eta) &= b(f - \hat{f}, f - \hat{f}) + 2b(f - \hat{f}, \eta) + b(\eta, \eta) \\ &= b(f - \hat{f}, f - \hat{f}) + b(\eta, \eta) \\ &\geq b(f - \hat{f}, f - \hat{f}), \end{aligned}$$

where the equality follows from $b(f - \hat{f}, \eta) = 0$ and the inequality follows from the fact that the bilinear form defines a seminorm. The result then follows from this and the fact that any $\psi \in \mathcal{F}$ can be written as $\hat{f} - \eta$ for some $\eta \in \mathcal{F}$.

It just remains to show that $b(f - \hat{f}, \psi) = 0$ for $\psi \in \mathcal{F}$. We prove this in two steps. In the

3. The bilinear form b is said to be a seminorm if (i) $b(\psi, \eta) \leq b(\psi, \psi) + b(\eta, \eta)$ and (ii) $b(a\psi, a\psi) = |a|b(\psi, \psi)$ for $a \in \mathbb{R}$ and ψ, η that are twice continuously differentiable almost everywhere on $\hat{\mathcal{W}}$.

first step, we show that

$$b(f, \psi) = \int_{\mathcal{B}} h(w)\psi(w)dw + \sum_{l=1}^d \int_{\partial\mathcal{B}^l} \hat{\kappa}_l \psi(w)dw, \quad \psi \in \mathcal{F}. \quad (\text{B.25})$$

In the second step, we show that

$$b(\hat{f}, \psi) = \int_{\mathcal{B}} h(w)\psi(w)dw + \sum_{l=1}^d \int_{\partial\mathcal{B}^l} \hat{\kappa}_l \psi(w)dw, \quad \psi \in \mathcal{F}. \quad (\text{B.26})$$

The desired equality, i.e., $b(f - \hat{f}, \psi) = 0$ for $\psi \in \mathcal{F}$, follows from (B.25)-(B.26).

Step 1. Since f satisfies (B.10)-(B.11), we have

$$\int_{\mathcal{B}} (\mathcal{L}f(w) - \alpha f(w) + h(w))\psi(w)dw = 0, \quad \psi \in \mathcal{F}. \quad (\text{B.27})$$

$$\sum_{l=1}^d \int_{\partial\mathcal{B}^l} (\nabla f(w)' \hat{G}^l + \hat{\kappa}_l) \psi(w)dw = 0, \quad \psi \in \mathcal{F}. \quad (\text{B.28})$$

First, we use Lemma 14 to write (B.27) in terms of first order derivatives. Then, we combine (B.27) and (B.28) to arrive at the desired result. By writing out the differential operator \mathcal{L} in (B.27), we obtain

$$\frac{1}{2} \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_{ij} f(w) \psi(w) dw - \int_{\mathcal{B}} \alpha f(w) \psi(w) dw + \int_{\mathcal{B}} h(w) \psi(w) dw = 0, \quad \psi \in \mathcal{F}. \quad (\text{B.29})$$

Substituting (B.16) into (B.29) gives

$$\begin{aligned} & -\frac{1}{2} \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j f(w) \partial_i \psi(w) dw + \frac{1}{2} \int_{\partial\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j f(w) n_i(w) \psi(w) dw \\ & - \int_{\mathcal{B}} \alpha f(w) \psi(w) dw + \int_{\mathcal{B}} h(w) \psi(w) dw = 0, \quad \psi \in \mathcal{F} \end{aligned} \quad (\text{B.30})$$

Adding the left-hand side of (B.28) to the left-hand side of (B.30), the right-hand side of

(B.28) to the right-hand side of (B.30) gives

$$\begin{aligned}
& -\frac{1}{2} \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j f(w) \partial_i \psi(w) dw + \frac{1}{2} \int_{\partial \mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j f(w) n_i(w) \psi(w) dw \\
& - \int_{\mathcal{B}} \alpha f(w) \psi(w) dw + \int_{\mathcal{B}} h(w) \psi(w) dw + \sum_{l=1}^d \int_{\partial \mathcal{B}^l} (\nabla f(w)' \hat{G}^l + \hat{\kappa}_l) \psi(w) dw = 0, \quad \psi \in \mathcal{F}.
\end{aligned}$$

Then, by rearranging the terms, we obtain

$$\begin{aligned}
& \frac{1}{2} \int_{\mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j f(w) \partial_i \psi(w) dw + \int_{\mathcal{B}} \alpha f(w) \psi(w) dw - \frac{1}{2} \int_{\partial \mathcal{B}} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_j f(w) n_i(w) \psi(w) dw \\
& - \sum_{l=1}^d \int_{\partial \mathcal{B}^l} \nabla f(w)' \hat{G}^l \psi(w) dw = \int_{\mathcal{B}} h(w) \psi(w) dw + \sum_{l=1}^d \int_{\partial \mathcal{B}^l} \hat{\kappa}_l \psi(w) dw, \quad \psi \in \mathcal{F},
\end{aligned}$$

which concludes step 1.

Step 2. This step is almost identical to the first step with the exception that we use (B.14)-(B.15) instead of (B.27)-(B.28). \square

Next, we discuss the choice of the basis functions ψ^m for $m = 1, \dots, M$. First, we choose a finite number of closed sets $\mathcal{T}^1, \dots, \mathcal{T}^N$ such that their intersection has a Lebesgue measure zero and their union is the region of inaction \mathcal{B} . These sets will be used to decompose the integrals over \mathcal{B} into sums of integrals over $\mathcal{T}^1, \dots, \mathcal{T}^N$. Such a partition of \mathcal{B} is called a mesh or triangulation on \mathcal{B} . The sets $\mathcal{T}^1, \dots, \mathcal{T}^N$ are called the elements and their vertices are called nodes. For the NYC application, we use triangular elements. Triangular elements can accommodate any polygonal \mathcal{B} , which provides great flexibility for the algorithm discussed in Section 2.7. Following Larson and Bengzon [151, Section 1.1.3], we first discretize the region of inaction \mathcal{B} ; see Figure B.2 for an illustration of the discretized region of inaction in two dimensions. Then, we construct a mesh of the region of inaction \mathcal{B} such that the nodes coincide with the discretized points. We use the Delaunay triangulation algorithm of Barber et al. [42] to construct the mesh. We denote the nodes by w^m for $m = 1, \dots, M$ and define

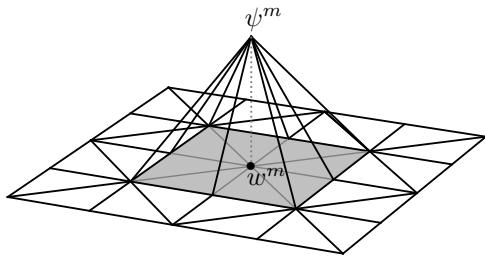


Figure B.1: An example of a hat function in two dimensions. The shaded area depicts the support of the hat function.

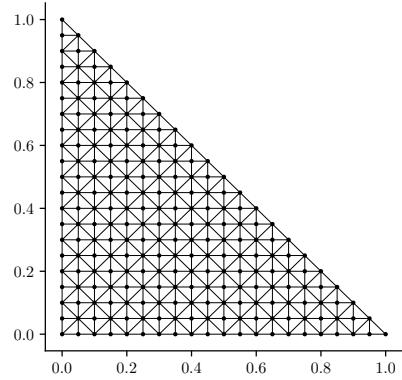


Figure B.2: An example of a mesh of $\hat{\mathcal{W}}$ in two dimensions, i.e., $p = 3$. Each triangle is an element.

the basis functions on the nodes as follows:

$$\psi^m(w) = \begin{cases} 1 & \text{if } w = w^m, \\ 0 & \text{if } w = w^l, l \neq m. \end{cases} \quad (\text{B.31})$$

The basis functions ψ^m can be defined in a variety of manners at points in \mathcal{B} other than the mesh points. For second-order PDEs, linear interpolation in each element \mathcal{T}^n is the most common approach. These (linearly interpolated) basis functions are referred to as the "hat" functions; see Figure B.1 for an illustration of the hat functions in two dimensions. Each hat function corresponds to a node. The hat function is non-zero only on the elements for which the corresponding node is a vertex. The elements on which the hat function is non-zero are shaded in Figure B.1.

We conclude this section by providing closed-form expressions for matrix D and vector q defined in Equations (B.17)-(B.18), in the two-dimensional case, which is the relevant case for the NYC application. Our discussion closely follows Larson and Bengzon [151, Sections 3.5-3.6 and 4.6]. We first express the integrals over \mathcal{B} and $\partial\mathcal{B}$ as sums of integrals over the

elements \mathcal{T}^n for $n = 1, \dots, N$, i.e.,

$$D_{km} = \sum_{n=1}^N \sum_{i,j=1}^{p-1} \frac{1}{2} \hat{\sigma}_{ij} \int_{\mathcal{T}^n} \partial_j \psi^m(w) \partial_i \psi^k(w) dw + \alpha \sum_{n=1}^N \int_{\mathcal{T}^n} \psi^m(w) \psi^k(w) dw \\ - \sum_{n=1}^N \sum_{i,j=1}^{p-1} \frac{1}{2} \hat{\sigma}_{ij} \int_{\partial \mathcal{B} \cap \mathcal{T}^n} \partial_j \psi^m(w) n_i(w) \psi^k(w) dw - \sum_{n=1}^N \sum_{l=1}^d \int_{\partial \mathcal{B}^l \cap \mathcal{T}^n} \nabla \psi^m(w)' \hat{G}^l \psi^k(w) dw \quad (\text{B.32})$$

and

$$q_k = \sum_{n=1}^N \int_{\mathcal{T}^n} h(w) \psi^k(w) dw + \sum_{n=1}^N \sum_{l=1}^d \hat{\kappa}_l \int_{\partial \mathcal{B}^l \cap \mathcal{T}^n} \psi^k(w) dw \quad (\text{B.33})$$

for $k, m = 1, \dots, M$. Then, we use the following identities to compute the matrix D and the vector q : For $k, m = 1, \dots, M$ and $n = 1, \dots, N$, we have

$$\int_{\mathcal{T}^n} \psi^m(w) \psi^k(w) dw = \begin{cases} \text{Area}(\mathcal{T}^n)/12 & \text{if } w^m \text{ and } w^k \text{ are vertices of } \mathcal{T}^n \text{ and } m \neq k, \\ \text{Area}(\mathcal{T}^n)/6 & \text{if } w^m \text{ and } w^k \text{ are vertices of } \mathcal{T}^n \text{ and } m = k, \\ 0 & \text{otherwise,} \end{cases}$$

$$\int_{\mathcal{T}^n} h(w) \psi^k(w) dw \simeq \begin{cases} h(w^k) \text{Area}(\mathcal{T}^n)/3 & \text{if } w^k \text{ is a vertex of } \mathcal{T}^n, \\ 0 & \text{otherwise,} \end{cases}$$

$$\int_{\mathcal{T}^n} \psi^k(w) dw = \begin{cases} \text{Area}(\mathcal{T}^n)/3 & \text{if } w^k \text{ is a vertex of } \mathcal{T}^n, \\ 0 & \text{otherwise,} \end{cases}$$

where $a \simeq b$ denotes $a = b + o(\text{Area}(\mathcal{T}^n))$.⁴ Now, we fix $n \in \{1, \dots, N\}$ and let w^1, w^2 , and w^3 denote be the vertices of \mathcal{T}^n . Then, $\psi^i(w) = (a^i + b^i w_1 + c^i w_2)/2\text{Area}(\mathcal{T}^n)$ for $i = 1, 2, 3$

4. We can ensure that the sum (over n) of the error terms is negligible, i.e., converges to zero as $n \rightarrow \infty$, by making sure that the mesh satisfies $\max_n \text{Area}(\mathcal{T}^n) = O(1/n)$.

and $w \in \mathcal{T}^n$, where

$$\begin{aligned} a^1 &= w_1^2 w_2^3 - w_1^3 w_2^2, & b^1 &= w_2^2 - w_2^3, & c^1 &= w_1^3 - w_1^2 \\ a^2 &= w_1^3 w_2^1 - w_1^1 w_2^3, & b^2 &= w_2^3 - w_2^1, & c^2 &= w_1^1 - w_1^3 \\ a^3 &= w_1^1 w_2^2 - w_1^2 w_2^1, & b^3 &= w_2^1 - w_2^2, & c^3 &= w_1^2 - w_1^1, \end{aligned}$$

and $\nabla \psi^i(w) = (b^i/2\text{Area}(\mathcal{T}^n), c^i/2\text{Area}(\mathcal{T}^n))'$ for $w \in \mathcal{T}^n$.

B.1.8 Solving the Equivalent Workload Formulation using the Markov Chain Approximation Method

This section describes the Markov chain approximation method that provides an approximate numerical solution to the EWF (2.47)-(2.50). This method constructs a Markov chain and an associated control problem whose value function approximates the value function of the EWF (2.47)-(2.50) at a discrete subset of the states space, i.e., the state space of the Markov chain. Then, it numerically solves the control problem and interprets its solution in the context of the EWF (2.47)-(2.50) in order to obtain an approximate solution for the EWF (2.47)-(2.50). This method was initially introduced in Kushner [146]; also see Kushner and Martins [148] and Kushner and Dupuis [147]. We first discuss the intuition behind the approximation method. Then, we describe the state space, the action space, and the state transitions of the approximating Markov chain. In doing so, we focus on the two-dimensional case as this is the relevant case for the NYC application. In the remainder of this section, similar to Section 2.7, dummy variables w and u are used instead of \hat{w} and \hat{u} .

The Markov chain approximation method takes advantage of the hypothesis that the (optimal) value function of the EWF (2.47)-(2.50) is the unique solution of the following

dynamic programming equation:

$$(\mathcal{L}\phi(w) - \alpha\phi(w) + h(w)) \wedge \inf_{\|\hat{G}u\|=1} \{\nabla\phi(w)' \hat{G}u + \hat{\kappa}'u\} = 0, \quad w \in \hat{\mathcal{W}}. \quad (\text{B.34})$$

The Markov chain approximation method seeks a discrete-time, discrete state-space Markov chain and an associated control problem, whose value function approximates the solution of the dynamic programming equation (B.34).

State Space. Let δ denote the discretization step size and define the δ -grid $\bar{\mathcal{W}}$ and the extended δ -grid $\bar{\mathcal{W}}^+$ as follows:⁵

$$\bar{\mathcal{W}} = \{(i\delta, j\delta) \in \hat{\mathcal{W}} : i, j \in \mathbb{N}\}$$

and

$$\bar{\mathcal{W}}^+ = \{(i\delta + i'\delta, j\delta + j'\delta) : i', j' \in \{-1, 0, +1\}, (i\delta, j\delta) \in \bar{\mathcal{W}}\}.$$

The extended grid is the super-set of $\bar{\mathcal{W}}$ and includes the states that the approximating Markov chain could transition to from $w \in \bar{\mathcal{W}}$. The state space of the approximating Markov chain is the extended grid $\bar{\mathcal{W}}^+$. We define $\partial\bar{\mathcal{W}}^+ = \bar{\mathcal{W}}^+ \setminus \bar{\mathcal{W}}$ and call it the reflecting boundary. The grid $\bar{\mathcal{W}}$ and the reflecting boundary $\partial\bar{\mathcal{W}}^+$ for the NYC application are depicted in Figure B.3; black dots represent the grid $\bar{\mathcal{W}}$ and gray circles represent the reflecting boundary $\partial\bar{\mathcal{W}}^+$. The inclusion of this reflecting boundary in the state space is essential for capturing the behavior of the Brownian motion $\hat{\chi}$ on the boundary of the state space $\hat{\mathcal{W}}$ of the EWF (2.47)-(2.50). Let $\bar{\phi}(w)$ denote the (optimal) value function of the approximating Markov chain at state $w \in \bar{\mathcal{W}}^+$, i.e., the minimum achievable expected discounted cost starting at w .

5. We choose $\delta = 1/n \in (0, 1)$ so that $1/\delta = n$ is an integer.

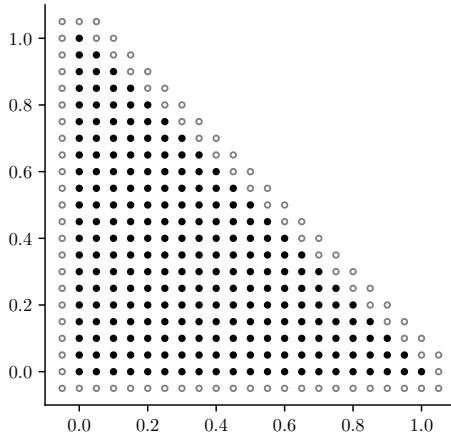


Figure B.3: State space of the approximating Markov chain with $\delta = 0.05$ for the NYC application. Black dots represent the grid $\bar{\mathcal{W}}$ and gray circles represent the reflecting boundary $\partial\bar{\mathcal{W}}^+$.

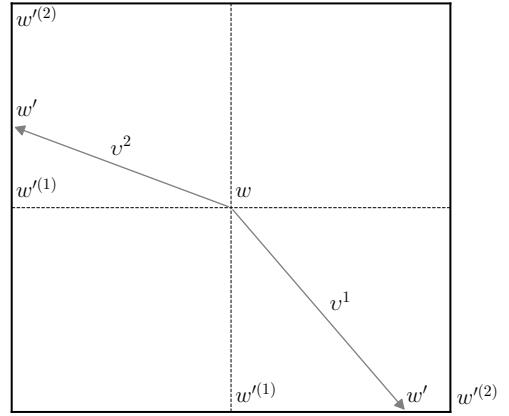


Figure B.4: Two examples of controls in the approximating Markov chain method. Each control in the approximating Markov chain method results in a random transition to one of two neighboring states.

Actions and Transitions on $\bar{\mathcal{W}}$. At state $w \in \bar{\mathcal{W}}$, we can take one of the following actions: (i) exert no control, or (ii) exert control in some direction v . When we exert no control, we want the value function of the Markov chain $\bar{\phi}$ to locally approximate the solution ϕ to the PDE

$$\mathcal{L}\phi(w) - \alpha\phi(w) + h(w) = 0. \quad (\text{B.35})$$

When we exert control, we want the value function of the Markov chain $\bar{\phi}$ to locally approximate the solution ϕ to

$$\inf_{\|\hat{G}u\|=1} \{\nabla\phi(w)' \hat{G}u + \hat{\kappa}'u\} = 0. \quad (\text{B.36})$$

Equation (B.36) states that when control is exerted, the change in value function (as a result of the control displacement) must cancel out the control cost.

Next, we discuss the transitions (for the approximating Markov chain) that give rise to this

desired behavior.

The Case of No Control. In this case, we set the period length for the approximating Markov chain as $\Delta t = \delta^2/Q$, where

$$Q = \sum_{i=1}^{p-1} \hat{\sigma}_{ii} - \sum_{\substack{i,j=1 \\ j \neq i}}^{p-1} |\hat{\sigma}_{ij}|/2 > 0.$$

Also, we let the (immediate) cost to be

$$\bar{C}(w, \text{No Control}) = h(w) \Delta t.$$

Lastly, the transition probabilities are as follows:

$$\mathbb{P}(w' | w, \text{No Control}) = \begin{cases} 2\hat{\sigma}_{ii} - \sum_{i,j=1, j \neq i}^{p-1} |\hat{\sigma}_{ij}|/2Q & \text{if } w' = w \pm e^i \delta, \\ \hat{\sigma}_{i,j}^+/2Q & \text{if } w' = w + e^i \delta + e^j \delta, i \neq j, \\ \hat{\sigma}_{i,j}^-/2Q & \text{if } w' = w - e^i \delta - e^j \delta, i \neq j, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{B.37})$$

By replacing the differential operator \mathcal{L} in Equation (B.35) with finite difference approximations, Lemma 16 obtains a discrete approximation to Equation (B.35). Letting $x^+ = \max(x, 0)$ and $x^- = \min(x, 0)$, we next state the lemma.

Lemma 16. *The solution to (B.35) satisfies*

$$\begin{aligned}
\phi(w) = & \exp(-\alpha \frac{\delta^2}{Q}) \sum_{i=1}^{p-1} \frac{1}{2Q} \left(2\hat{\sigma}_{ii} - \sum_{\substack{j=1 \\ j \neq i}}^{p-1} |\hat{\sigma}_{ij}| \right) \phi(w + e^i \delta) \\
& + \exp(-\alpha \frac{\delta^2}{Q}) \sum_{i=1}^{p-1} \frac{1}{2Q} \left(2\hat{\sigma}_{ii} - \sum_{\substack{j=1 \\ j \neq i}}^{p-1} |\hat{\sigma}_{ij}| \right) \phi(w - e^i \delta) \\
& + \exp(-\alpha \frac{\delta^2}{Q}) \sum_{\substack{i,j=1 \\ j \neq i}}^{p-1} \frac{1}{2Q} \hat{\sigma}_{ij}^+ \phi(w + e^i \delta + e^j \delta) \\
& + \exp(-\alpha \frac{\delta^2}{Q}) \sum_{\substack{i,j=1 \\ j \neq i}}^{p-1} \frac{1}{2Q} \hat{\sigma}_{ij}^+ \phi(w - e^i \delta - e^j \delta) \\
& + \exp(-\alpha \frac{\delta^2}{Q}) \sum_{\substack{i,j=1 \\ j \neq i}}^{p-1} \frac{1}{2Q} \hat{\sigma}_{ij}^- \phi(w - e^i \delta + e^j \delta) \\
& + \exp(-\alpha \frac{\delta^2}{Q}) \sum_{\substack{i,j=1 \\ j \neq i}}^{p-1} \frac{1}{2Q} \hat{\sigma}_{ij}^- \phi(w + e^i \delta - e^j \delta) \\
& + \exp(-\alpha \frac{\delta^2}{Q}) \frac{\delta^2}{Q} h(w) + O(\delta^4). \tag{B.38}
\end{aligned}$$

Proof. Note that Equation (B.35) can be written more explicitly as follows:

$$\frac{1}{2} \sum_{i,j=1}^{p-1} \hat{\sigma}_{ij} \partial_{ij} \phi(w) - \alpha \phi(w) + h(w) = 0.$$

which we then rewrite as below.

$$\frac{1}{2} \sum_{i=1}^{p-1} \hat{\sigma}_{ii} \partial_{ii} \phi(w) + \frac{1}{2} \sum_{\substack{i=1 \\ \hat{\sigma}_{ij} \geq 0}}^{p-1} \sum_{j \neq i} \hat{\sigma}_{ij} \partial_{ij} \phi(w) + \frac{1}{2} \sum_{i=1}^{p-1} \sum_{\substack{j \neq i \\ \hat{\sigma}_{ij} < 0}} \hat{\sigma}_{ij} \partial_{ij} \phi(w) - \alpha \phi(w) + h(w) = 0. \tag{B.39}$$

Consider the following finite difference approximations (see e.g., Kushner and Martins [148,

Equations (5.2)-(5.4)]): For $i, j = 1, \dots, p-1$, $i \neq j$, and $\hat{\sigma}_{ij} \geq 0$,

$$\begin{aligned}\partial_{ii} \phi(w) &= \frac{\phi(w + e^i \delta) + \phi(w - e^i \delta) - 2\phi(w)}{\delta^2} + O(\delta^2), \\ \partial_{ij} \phi(w) &= \frac{2\phi(w) + \phi(w + e^i \delta + e^j \delta) + \phi(w - e^i \delta - e^j \delta)}{2\delta^2} \\ &\quad - \frac{\phi(w + e^i \delta) + \phi(w - e^i \delta) + \phi(w + e^j \delta) + \phi(w - e^j \delta)}{2\delta^2} + O(\delta^2), \quad i \neq j, \hat{\sigma}_{ij} \geq 0, \\ \partial_{ij} \phi(w) &= - \frac{2\phi(w) + \phi(w + e^i \delta - e^j \delta) + \phi(w - e^i \delta + e^j \delta)}{2\delta^2} \\ &\quad + \frac{\phi(w + e^i \delta) + \phi(w - e^i \delta) + \phi(w + e^j \delta) + \phi(\hat{w} - e^j \delta)}{2\delta^2} + O(\delta^2), \quad i \neq j, \hat{\sigma}_{ij} < 0.\end{aligned}$$

Substituting these finite difference approximations into Equation (B.39) gives

$$\begin{aligned}& \frac{1}{2} \sum_{i=1}^{p-1} \frac{\hat{\sigma}_{ii}}{\delta^2} (\phi(w + e^i \delta) + \phi(w - e^i \delta) - 2\phi(w)) \\ &+ \frac{1}{2} \sum_{\substack{i,j=1, j \neq i \\ \hat{\sigma}_{ij} \geq 0}}^{p-1} \frac{\hat{\sigma}_{ij}}{2\delta^2} (2\phi(w) + \phi(\hat{w} + e^i \delta + e^j \delta) + \phi(w - e^i \delta - e^j \delta)) \\ &- \frac{1}{2} \sum_{\substack{i,j=1, j \neq i \\ \hat{\sigma}_{ij} \geq 0}}^{p-1} \frac{\hat{\sigma}_{ij}}{2\delta^2} (\phi(w + e^i \delta) + \phi(w - e^i \delta) + \phi(w + e^j \delta) + \phi(w - e^j \delta)) \\ &- \frac{1}{2} \sum_{\substack{i,j=1, j \neq i \\ \hat{\sigma}_{ij} < 0}}^{p-1} \frac{\hat{\sigma}_{ij}}{2\delta^2} (2\phi(w) + \phi(w + e^i \delta - e^j \delta) + \phi(w - e^i \delta + e^j \delta)) \\ &+ \frac{1}{2} \sum_{\substack{i,j=1, j \neq i \\ \hat{\sigma}_{ij} < 0}}^{p-1} \frac{\hat{\sigma}_{ij}}{2\delta^2} (\phi(w + e^i \delta) + \phi(w - e^i \delta) + \phi(w + e^j \delta) + \phi(w - e^j \delta)) \\ &- \alpha \phi(w) + h(w) + O(\delta^2) = 0.\end{aligned}$$

Multiplying both sides by δ^2 and rearranging the terms gives

$$\begin{aligned}
\left(\sum_{i=1}^{p-1} \hat{\sigma}_{ii} - \frac{1}{2} \sum_{i=1}^{p-1} \sum_{j \neq i} |\hat{\sigma}_{ij}| + \alpha \delta^2 \right) \phi(w) &= \sum_{i=1}^{p-1} \left(\hat{\sigma}_{ii} - \frac{1}{2} \sum_{j \neq i} |\hat{\sigma}_{ij}| \right) \phi(w + e^i \delta) \\
&\quad + \sum_{i=1}^{p-1} \left(\hat{\sigma}_{ii} - \frac{1}{2} \sum_{j \neq i} |\hat{\sigma}_{ij}| \right) \phi(w - e^i \delta) \\
&\quad + \sum_{i=1}^{p-1} \sum_{j \neq i} \frac{1}{2} \hat{\sigma}_{ij}^+ \phi(w + e^i \delta + e^j \delta) \\
&\quad + \sum_{i=1}^{p-1} \sum_{j \neq i} \frac{1}{2} \hat{\sigma}_{ij}^+ \phi(w - e^i \delta - e^j \delta) \\
&\quad + \sum_{i=1}^{p-1} \sum_{j \neq i} \frac{1}{2} \hat{\sigma}_{ij}^- \phi(w - e^i \delta + e^j \delta) \\
&\quad + \sum_{i=1}^{p-1} \sum_{j \neq i} \frac{1}{2} \hat{\sigma}_{ij}^- \phi(w + e^i \delta - e^j \delta) \\
&\quad + \delta^2 h(w) + O(\delta^4). \tag{B.40}
\end{aligned}$$

Dividing both sides of Equation (B.40) by Q gives

$$\begin{aligned}
\left(1 + \frac{\alpha \delta^2}{Q} \right) \phi(w) &= \sum_{i=1}^{p-1} \frac{1}{2Q} \left(2\hat{\sigma}_{ii} - \sum_{j \neq i} |\hat{\sigma}_{ij}| \right) \phi(w + e^i \delta) + \sum_{i=1}^{p-1} \frac{1}{2Q} \left(2\hat{\sigma}_{ii} - \sum_{j \neq i} |\hat{\sigma}_{ij}| \right) \phi(w - e^i \delta) \\
&\quad + \sum_{i=1}^{p-1} \sum_{j \neq i} \frac{1}{2Q} \hat{\sigma}_{ij}^+ \phi(w + e^i \delta + e^j \delta) + \sum_{i=1}^{p-1} \sum_{j \neq i} \frac{1}{2Q} \hat{\sigma}_{ij}^+ \phi(w - e^i \delta - e^j \delta) \\
&\quad + \sum_{i=1}^{p-1} \sum_{j \neq i} \frac{1}{2Q} \hat{\sigma}_{ij}^- \phi(w - e^i \delta + e^j \delta) + \sum_{i=1}^{p-1} \sum_{j \neq i} \frac{1}{2Q} \hat{\sigma}_{ij}^- \phi(w + e^i \delta - e^j \delta) \\
&\quad + \frac{\delta^2}{Q} h(w) + O(\delta^4). \tag{B.41}
\end{aligned}$$

Then, substituting the approximation $\exp(\alpha \delta^2 / Q) = 1 + \alpha \delta^2 / Q + O(\delta^4)$ into Equation (B.41) and multiplying both sides by $\exp(-\alpha \delta^2 / Q)$ gives the result. \square

We interpret Equation (B.38) as the Bellman equation of a discounted infinite horizon

Markov chain with period length $\Delta t = \delta^2/Q$, immediate cost $\bar{C}(w, \text{No Control}) = h(w) \Delta t$, and transition probabilities given by (B.37); see Kushner and Dupuis [147, Section 5.3.1] for further details.

As a preliminary to discussing the case where we exert control at direction v , let w' denote the intersection of the ray emanating from w in direction v with the grid lines surrounding w as depicted in Figure B.4. Note that w' is uniquely determined by the state w and the direction v . If $w' \in \bar{\mathcal{W}}$, the Markov chain transitions with probability one to w' . However, w' need not belong to $\bar{\mathcal{W}}$, i.e., it need not be a grid point. To address this issue, we define the states $w'^{(1)}, w'^{(2)} \in \mathcal{W}$ as the two grid points that are neighboring w and closest to it on either side on the grid line depicted in Figure B.4.

The Case of Exerting Control in Direction v . In this case, the Markov chain transitions from state w to $w'' \in \{w'^{(1)}, w'^{(2)}\}$ instantaneously with the following transition probabilities:

$$\mathbb{P}(w'' | w, \text{Control } v) = \begin{cases} |w' - w'^{(2)}|/\delta & \text{if } w'' = w'^{(1)}, \\ |w' - w'^{(1)}|/\delta & \text{if } w'' = w'^{(2)}, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{B.42})$$

Since controls create an instantaneous displacement in the EWF (2.47)-(2.50), the (immediate) cost of control is $\bar{C}(w, \text{Control } v) = C(w, w')$, where C is defined in Equations (2.57)-(2.59). Figure B.4 illustrates the transitions in this case. When we exert control at direction v , the transition function of the approximating Markov chain is chosen by a randomization between states $w'^{(1)}$ and $w'^{(2)}$ such that the mean increment is $w' - w$. In other words, $\mathbb{P}(w'^{(1)} | w, \text{Control } v)$ and $\mathbb{P}(w'^{(2)} | w, \text{Control } v)$ are chosen such that

$$(w'^{(1)} - w) \times \mathbb{P}(w'^{(1)} | w, \text{Control } v) + (w'^{(2)} - w) \times \mathbb{P}(w'^{(2)} | w, \text{Control } v) = w' - w,$$

i.e., mean increment is $w' - w$, and $\mathbb{P}(w'^{(1)} | w, \text{Control } v) + \mathbb{P}(w'^{(2)} | w, \text{Control } v) = 1$. This

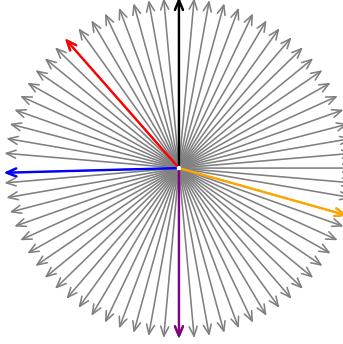


Figure B.5: The discretized control directions for the NYC application.

gives rise to the transition probabilities given in Equation (B.42). Note that $w'^{(1)}$ and $w'^{(2)}$ are introduced merely to address the problem that w' need not belong to $\bar{\mathcal{W}}$. The states $w'^{(1)}$ and $w'^{(2)}$ play no role in determining the (immediate) cost of control. The (immediate) cost of control is derived directly from w' . To be more specific, the (immediate) cost of control at direction v is $\bar{C}(w, \text{Control } v) = C(w, w')$; see Equations (2.57)-(2.59).

Kushner and Martins [148] and Kushner and Dupuis [147] assume that control has to be exerted in one of finitely many directions. In our case, however, there are an infinite number of directions available to the decision maker as the decision maker can use any linear combination of the pushing controls to move the state process. We remedy this problem by discretizing the set of control directions; see Figure B.5 for an illustration. For the NYC application, we use a discretization step size of five degrees. Let \mathcal{C} denote the set of discretized control directions. The action set at $w \in \bar{\mathcal{W}}$ is

$$\bar{\mathcal{A}}(w) = \{\text{No Control}\} \cup \mathcal{C}, \quad w \in \bar{\mathcal{W}}.$$

Actions and Transitions on $\partial\bar{\mathcal{W}}^+$. At $w \in \partial\bar{\mathcal{W}}^+$, the decision maker has to exert control. However, only a limited set of control directions are feasible. The action set at

$w \in \partial\bar{\mathcal{W}}^+$ is

$$\bar{\mathcal{A}}(w) = \{v \in \mathcal{C} : \exists \epsilon > 0 \text{ such that } w + \epsilon v \in \hat{\mathcal{W}}\}, \quad w \in \partial\bar{\mathcal{W}}^+.$$

Given a feasible control direction $v \in \bar{\mathcal{A}}(w)$, the state transitions are the same as those given by Equation (B.42).

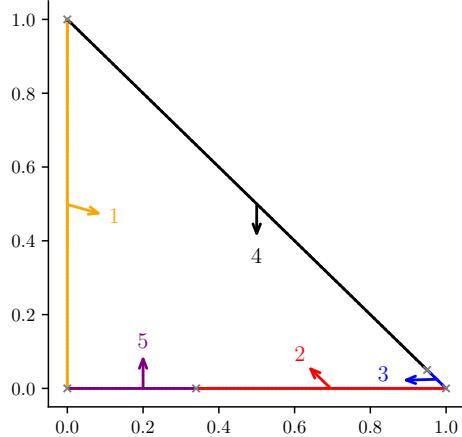
Bellman Equation. The (optimal) value function of the approximating Markov chain must satisfy

$$\bar{\phi}(w) = \min_{a \in \bar{\mathcal{A}}(w)} \left\{ \bar{C}(w, a) + \sum_{w' \in \bar{\mathcal{W}}^+} \exp(-\alpha \Delta t_a) \mathbb{P}(w' | w, a) \bar{\phi}(w') \right\}, \quad w \in \bar{\mathcal{W}}^+. \quad (\text{B.43})$$

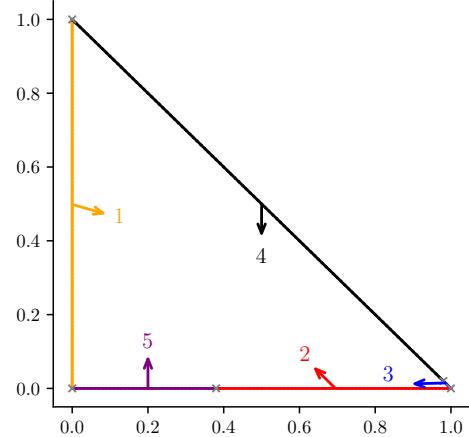
Note that when the decision maker does not exert control, $\Delta t_a = \delta^2/Q > 0$. However, when the decision exerts control (at any direction), we set $\Delta t = 0$ because the transition is instantaneous.

Kushner and Martins [148, Theorem 7.1] shows that as $\delta \rightarrow 0$, the (optimal) value function of the control problem discussed above converges pointwise to the (optimal) value function of the EWF (2.47)-(2.50), i.e., $\bar{\phi}(w) \rightarrow \phi(w)$ for $w \in \bar{\mathcal{W}}$.

Interpretation. We interpret the optimal solution to the control problem discussed above as a control policy for the EWF (2.47)-(2.50) as follows: In the interior points of the state space, i.e., $w \in \text{int}(\hat{\mathcal{W}})$, we use the action prescribed by the Bellman equation (Equation (B.43)) at the grid point w' of the state space $\bar{\mathcal{W}}$ of the Markov chain that is closest to w , i.e., $w' = \text{argmin}_{\bar{w} \in \bar{\mathcal{W}}} \|\bar{w} - w\|$. Recall that $\bar{\mathcal{W}}$ is the δ -grid of $\hat{\mathcal{W}}$. At the boundary points of the state space, i.e., $w \in \partial\hat{\mathcal{W}}$, we use the action prescribed by the Bellman equation at the grid point w' of the reflecting boundary $\partial\bar{\mathcal{W}}^+$ of the Markov chain that is closest to w , i.e., $w' = \text{argmin}_{\bar{w} \in \partial\bar{\mathcal{W}}^+} \|\bar{w} - w\|$. In both cases, ties are broken in favor of the smaller (pushing control) indices.



(a) Our computational method.



(b) Markov chain approximation method.

Figure B.6: The control policies derived by our computational method and the Markov chain approximation method for the NYC application ($\delta = 0.01$).

Table B.2: Pushing controls used under the control policy derived by the Markov chain approximation method for the EWF (2.47)-(2.50) of the NYC application ($\delta = 0.01$).

State	Pushing Control	
	Number	Action
$\{w : w_1 = 0, 0 < w_2 < 1\}$	1	Idle Server 5
$\{w : w_2 = 0, w_1 < 0.385\}$	5	Serve Buffer 7 with Server 6
$\{w : w_2 = 0, w_1 \geq 0.385\}$	2	Idle Server 6
$\{w : w_1 + w_2 = 1, w_1 < 0.975\}$	4	Serve Buffer 6 with Server 7
$\{w : w_1 + w_2 = 1, w_1 \geq 0.975, w_2 > 0\}$	3	Idle Server 7

Comparison with Our Method. We conclude this section by comparing the performances of the two policies: the policy derived by our computational method (presented in Section 2.7) versus the one derived by the Markov chain approximation method. The control policy derived by our computational method is depicted in Figure B.6a and further explained in Table 2.1. The control policy derived by the Markov chain approximation method is depicted in Figure B.6b and further explained in Table B.2. Similar to the policy derived by our computational method, the policy derived by the Markov chain approximation method is a region-of-inaction type policy. In fact, the two control policies are quite similar. They only differ on two parts of the boundary $\partial\hat{\mathcal{W}}$. The first part is $\{w : w_2 = 0, 0.35 \leq w_1 < 0.385\}$. The policy derived by our computational method pre-

Table B.3: The expected discounted cost of the control policies derived by our computational method and the Markov chain approximation method for different discretization step sizes.

δ	Our Computational Method		Markov Chain Approximation	
	Cost (in millions)	Percentage Difference	Cost (in millions)	Percentage Difference
0.04	4.389	1.64%	4.422	2.40%
0.02	4.337	0.43%	4.419	2.33%
0.01	4.318	0.00%	4.326	0.17%

scribes pushing control 2 for $w \in \{w : w_2 = 0, 0.35 \leq w_1 < 0.385\}$ while the policy derived by the Markov chain approximation method prescribes pushing control 5. The second part is $\{w : w_1 + w_2 = 1, 0.95 \leq w_1 < 0.975\}$. The policy derived by our computational method prescribes pushing control 3 for $w \in \{w : w_1 + w_2 = 1, 0.95 \leq w_1 < 0.975\}$ while the policy derived by the Markov chain approximation method prescribes pushing control 4.

Table B.3 reports the expected discounted cost of the control policies derived by the two methods for different discretization step sizes.⁶ We take the control policy derived by our computational method with $\delta = 0.01$ as reference. We observe that for all $w \in \hat{\mathcal{W}}$, the value function of the control policy derived by our computational method is lower than the value function of the control policy derived by the Markov chain approximation method. Table B.3 reports the minimal percentage difference between the value function of the control policy derived by the Markov chain approximation method and the value function of the control policy derived by our computational method. Table B.3 demonstrates that our computational method, outperforms the Markov chain approximation method in terms of the expected discounted cost.

We use the policy iteration algorithm to solve the Bellman equation (B.43). Our computational method terminates in three steps for all discretization step sizes considered while the Markov chain approximation method requires seven to nine (policy iteration) steps. This demonstrates that our computational method is (considerably) more computationally effi-

6. Note that the expected discounted costs in Table B.3 are reported in thousands. Then, since $\max_{w \in \hat{\mathcal{W}}} f(w) - \min_{w \in \hat{\mathcal{W}}} f(w) << 1000$ for the region-of-inaction type policies considered here, we have not specified the initial workload.

cient than the Markov chain approximation method. Recall that our computational method takes advantage of the special structure of the EWF (2.47)-(2.50) while the Markov chain approximation method is a general purpose method (used for solving many stochastic control problems). Note that the steps of our computational method are no more computationally expensive than the steps of the Markov chain approximation method. To see this, note that the policy iteration algorithm solves a system of linear equations (in the policy evaluation step) that has one variable for each state in the state space $\bar{\mathcal{W}}^+$ of the Markov chain. Our computational method solves a similar system of linear equations with a slightly smaller number of variables (one variable for each state in $\bar{\mathcal{W}}$); see Appendix B.1.7 for details. Moreover, the Markov chain approximation method iterates over all states (in the policy improvement step) while our computational method only iterates over the boundary points.

B.1.9 Policy Improvement

Lemma 17 shows that under certain assumptions, the control policy obtained in each step of the computational method discussed in Section 2.7 is an improvement over the previous control policy.

Lemma 17. *Suppose in step k no boundary point $w \in \partial\mathcal{B}^k$ is moved outwards and the (updated) region of inaction \mathcal{B}^{k+1} does not require convexification. Moreover, suppose the PDEs solved in steps k and $k+1$ have twice continuously differentiable solutions denoted by f_k and f_{k+1} , respectively. Furthermore, suppose there exists a region-of-inaction type control policy associated with the region of inaction \mathcal{B}^{k+1} and the partition $\partial\mathcal{B}^{k+1,1}, \dots, \partial\mathcal{B}^{k+1,d}$ of its boundary $\partial\mathcal{B}^{k+1}$. Then, $\mathcal{B}^{k+1} \subseteq \mathcal{B}^k$, i.e., the region of inaction has not expanded. Moreover, $f_{k+1}(w) \leq f_k(w)$ for $w \in \mathcal{B}^{k+1}$.*

Proof. Consider a step $k+1$ in which no boundary point $w \in \partial\mathcal{B}^k$ is moved outwards and the (updated) region of inaction \mathcal{B}^{k+1} does not require convexification. Note that in this step, the region of inaction has not expanded, i.e., $\mathcal{B}^{k+1} \subseteq \mathcal{B}^k$. First, we show that $g = f_{k+1} - f_k$

satisfies

$$\mathcal{L}g(w) - \alpha g(w) = 0, \quad w \in \text{int}(\mathcal{B}^{k+1}), \quad (\text{B.44})$$

$$\nabla g(w)' \hat{G}^l \geq 0, \quad w \in \partial\mathcal{B}^{k+1,l} \text{ and } l = 1, \dots, d. \quad (\text{B.45})$$

Then, we use Lemma 18 to show that $g(w) \leq 0$ for $w \in \mathcal{B}^{k+1}$, which gives $f_{k+1}(w) - f_k(w) \leq 0$ for $w \in \mathcal{B}^{k+1}$.

Step 1. Define $g \in C^2(\mathcal{B}^{k+1})$ as follows: $g(w) = f_{k+1}(w) - f_k(w)$ for $w \in \mathcal{B}^{k+1}$. It follows from the fact that f_k and f_{k+1} satisfy (2.60) for $w \in \text{int}(\mathcal{B}^{k+1})$ that g satisfies (B.44). It remains to show that g satisfies (B.45). Consider a pushing control $l = 1, \dots, d$ and a boundary point $w \in \partial\mathcal{B}^{k+1,l}$. Since the region of inaction has not expanded, one of the following updates has occurred in step k : (i) some boundary point $w' \in \partial\mathcal{B}^{k,l}$ has been moved inwards along pushing direction \hat{G}^l to $w \in \partial\mathcal{B}^{k+1,l}$, (ii) the boundary point w has not moved in the update, but the pushing control has been updated to l , or (iii) the boundary point w has not moved in the update and the pushing control has not been updated.

Next, we show that $\nabla g(w)' \hat{G}^l \geq 0$ in all three cases. Since $\nabla f_{k+1}(w)' \hat{G}^l = -\hat{\kappa}_l$ by Equation (2.61) and $g = f_{k+1} - f_k$, it suffices to show that $\nabla f_k(w)' \hat{G}^l \leq -\hat{\kappa}_l$. In the first case, $\nabla f_k(w)' \hat{G}^l < -\hat{\kappa}_l$ as illustrated in Figure 2.10a. In the second case, since the pushing control has been updated to l , we have $\nabla f_k(w)' \hat{G}^l \leq -\hat{\kappa}_l$ as illustrated in Figures 2.10a and 2.10b. In the third case, since the region of inaction has not expanded and the pushing control has not been updated, we have $\nabla f_k(w)' \hat{G}^l = -\hat{\kappa}_l$.

Step 2. Fix $w \in \mathcal{B}^{k+1}$. Let \hat{U} denote the region-of-inaction type control policy for initial workload w , the region of inaction \mathcal{B}^{k+1} , and the partition $\partial\mathcal{B}^{k+1,1}, \dots, \partial\mathcal{B}^{k+1,d}$ of its boundary $\partial\mathcal{B}^{k+1}$. We will use the generalized Ito's formula to show the desired result. Let \hat{U}^c denote the continuous part of \hat{U} , $\Delta\hat{U}$ denote its jumps, and $\mathcal{D}(\hat{U})$ denote the set of its

discontinuities. To be specific, for $t \geq 0$,

$$\Delta\hat{U}(t) = \hat{U}(t) - \hat{U}(t^-), \quad \mathcal{D}(\hat{U}) = \{t : \Delta\hat{U}(t) \neq 0\}, \quad \text{and} \quad \hat{U}^c(t) = \hat{U}(t) - \sum_{s \in \mathcal{D}(\hat{U}), s \leq t} \Delta\hat{U}(s).$$

Moreover, let \hat{W} denote state process associated with \hat{U} . By Lemma 18, we have

$$\begin{aligned} -g(w) = & \mathbb{E} \left[\int_0^\infty e^{-\alpha t} (\mathcal{L}g - \alpha g)(\hat{W}(t)) dt \right] + \mathbb{E} \left[\sum_{l=1}^d \int_0^\infty e^{-\alpha t} \nabla g(\hat{W}(t))' \hat{G}^l d\hat{U}_l^c(t) \right] \\ & + \mathbb{E} \left[\sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} (g(\hat{W}(t^-) + \hat{G} \Delta\hat{U}(t)) - g(\hat{W}(t^-))) \right]. \end{aligned} \quad (\text{B.46})$$

Since \hat{U} is a region-of-inaction type control policy with region of inaction \mathcal{B}^{k+1} and the initial workload $w \in \mathcal{B}^{k+1}$, we have $\hat{W}(t) \in \mathcal{B}^{k+1}$ for $t \geq 0$. Then, it follows from (B.44) that the first term on the right-hand side of (B.46) is zero. Moreover, since the initial workload $w \in \mathcal{B}^{k+1}$, the third term on the right-hand side of (B.46) is zero. Therefore,

$$-g(w) = \mathbb{E} \left[\sum_{l=1}^d \int_0^\infty e^{-\alpha t} \nabla g(\hat{W}(t))' \hat{G}^l d\hat{U}_l^c(t) \right]. \quad (\text{B.47})$$

Then, it follows from (B.45) and the fact that pushing control l is only used when $\hat{W}(t) \in \partial\mathcal{B}^{k+1,l}$ that $g(w) \leq 0$ for $w \in \mathcal{B}^{k+1}$. \square

B.1.10 Detailed Description of the Alternative Policies

In this section, we describe the alternative policies considered in Section 2.9 in more detail.

Static Policy. Under the static policy, when server k for $k = 1, \dots, 9$ becomes available, it engages in one of the basic dispatch activities available to it that serve a non-empty buffer, i.e., $\{j : k(j) = k, Q_{i(j)}(t) > 0, j = 1, \dots, 13\}$, with probabilities proportional to the basic activity rates x_j^* . If no such activity exists, the server idles. Moreover, the platform engages

in basic relocation activities $j = 1, 2$ at the rate \tilde{x}_j^* as long as there are jobs available at the origin to relocate, i.e., $Q_{o(j)}(t) > 0$.

Closest-Driver Policy. Let d_{ik} denote the travel distance (in miles) from area i to area k for $i, k = 1, \dots, 9$. Under the closest-driver policy, when server k for $k = 1, \dots, 9$ becomes available, it undertakes the dispatch activity that serves the closest non-empty buffer, i.e., the dispatch activity $j^* = \operatorname{argmin}\{d_{i(j)k} : k(j) = k, Q_{i(j)}(t) > 0, j = 1, \dots, 19\}$. If no such activity exists, the server idles.

MaxWeight Policy. Under the MaxWeight policy, when server k for $k = 1, \dots, 9$ becomes available, it engages in the dispatch activity that serves the buffer with the largest queue, i.e., the dispatch activity $j^* = \operatorname{argmax}\{Q_{i(j)}(t) : k(j) = k, Q_{i(j)}(t) > 0, j = 1, \dots, 19\}$. In case of a tie, the activity with the smallest distance $d_{i(j)k}$ is undertaken i.e., ties are broken in favor of proximity. If no such activity exists, the server idles.

Penalized MaxWeight Policy. Under the penalized MaxWeight policy, when server k for $k = 1, \dots, 9$ becomes available, it engages in the dispatch activity that serves the buffer with the largest queue minus a linear travel cost, i.e., the dispatch activity $j^* = \operatorname{argmax}\{Q_{i(j)}(t) - ad_{i(j)k} : k(j) = k, Q_{i(j)}(t) > 0, j = 1, \dots, 19\}$, where a is a tuning parameter. In case of a tie, the activity with the smallest distance $d_{i(j)k}$ is undertaken i.e., ties are broken in favor of proximity. If no such activity exists, the server idles.

Mirror Backpressure Policy. The mirror backpressure (MBP) policies are parameterized by a congestion function $\eta : (0, \infty) \rightarrow \mathbb{R}$. Following Kanoria and Qian [136], we consider the following congestion functions: (i) linear, i.e., $\eta(x) = ax$, (ii) logarithmic, i.e., $\eta(x) = a \log(x)$, and (iii) inverse square-root, i.e., $\eta(x) = -a/\sqrt{x}$, where $a > 0$ is tuning parameter that is calibrated via simulation. Letting $\tilde{Z}_i(t) = (Q_i(t) + \sqrt{\Lambda})/(\Lambda + m\sqrt{\Lambda})$ denote the centered and normalized queue length at time t in buffer i , the congestion cost of buffer i at time t is defined as $\eta(\tilde{Z}_i(t))$ for $i = 1, \dots, 9$ and $t \geq 0$. Unlike the other policies discussed above, the mirror backpressure policies require knowledge of the buffer the next

job served by a server will transition into. That is, under mirror backpressure policies, when server k for $k = 1, \dots, 9$ becomes available, it knows the buffer its next job will transition into (independent of the dispatch activity the server undertakes). In other words, the mirror backpressure policies use the destination of the customers, whereas the other policies including our proposed policy do not require the destination information. We use the notation l to denote a customer's destination, which also corresponds to the next buffer for the car transporting the customer. In order to decide the activity j^* , that server k is to engage at time t , the mirror backpressure policy ranks activities with respect to the index $v_j + \eta(\tilde{Z}_{i(j)}(t))$, i.e., the value of service per job plus the congestion cost at the origin. In doing so, it only considers the non-empty buffers for which the index $v_j + \eta(\tilde{Z}_{i(j)}(t))$ is greater than or equal to the congestion cost at the destination, i.e., $\eta(\tilde{Z}_l(t))$. That is, the dispatch activity $j^* = \operatorname{argmax}\{v_j + \eta(\tilde{Z}_{i(j)}(t)) : k(j) = k, Q_{i(j)}(t) > 0, v_j + \eta(\tilde{Z}_{i(j)}(t)) \geq \eta(\tilde{Z}_l(t)), j = 1, \dots, 19\}$ is undertaken. If no such activity exists, the server idles.

Modified Policies. Under the modified closest-driver policy, the modified MaxWeight policy, the modified penalized MaxWeight policy, and the modified mirror backpressure policy, in addition to the dispatch activities described above, the platform engages in basic relocation activities $j = 1, 2$ at the rate \tilde{x}_j^* as long as there are jobs available at the origin to relocate, i.e., $Q_{o(j)}(t) > 0$.

B.1.11 Auxiliary Simulation Results

This section reports the results of the simulation study discussed in Section 2.9. Tables B.4-B.5 report the expected discounted cost of the various policies as defined in Section 2.4; see Equations (2.11), (2.21), (2.30), and (2.31). To be more specific, Tables B.4-B.5 report the unscaled expected discounted costs ΛJ_z for each policy; see Equation (2.30) for the definition of J_z . Tables B.6-B.7 report the scaled expected discounted costs, i.e., J_z or the expected discounted cost per car. Tables B.8-B.9 report the percentage difference of the various policies considered from the proposed policy. Tables B.10-B.11 report the demand

dropping probabilities.

Table B.4: The expected discounted cost ΛJ_z in millions of dollars under the proposed policy, the static policy, and the original alternative policies (the values denoted in the parentheses indicate the 95% confidence interval).

Excess Supply			Original				
Total	Percentage	Proposed	Static	Closest-Driver	MaxWeight	MaxWeight+	MBP
-4100	-50%	3642.5 (± 1.6)	3691.4 (± 1.6)	4026.3 (± 0.1)	4107.0 (± 0.1)	3937.7 (± 0.1)	3363.0 (± 0.1)
-3280	-40%	2837.8 (± 1.4)	2908.4 (± 2.2)	3393.7 (± 0.1)	3518.3 (± 0.1)	3344.4 (± 0.1)	2613.3 (± 0.1)
-2460	-30%	2050.7 (± 0.8)	2157.8 (± 0.3)	2748.3 (± 0.1)	2928.5 (± 0.1)	2759.2 (± 0.1)	1869.0 (± 0.1)
-1640	-20%	1300.3 (± 0.4)	1432.5 (± 0.2)	2074.9 (± 0.2)	2338.9 (± 0.1)	2182.1 (± 0.1)	1155.4 (± 0.1)
-820	-10%	610.2 (± 0.3)	723.6 (± 0.2)	1156.6 (± 1.0)	1751.5 (± 0.1)	1607.3 (± 0.1)	717.8 (± 0.8)
-410	-5%	318.5 (± 0.3)	399.5 (± 0.2)	659.6 (± 0.7)	1460.2 (± 0.1)	1323.4 (± 0.1)	636.0 (± 1.4)
-246	-3%	218.3 (± 0.0)	289.3 (± 0.3)	513.2 (± 1.2)	1344.7 (± 0.1)	1211.1 (± 0.1)	586.9 (± 0.7)
-82	-1%	137.6 (± 0.3)	200.7 (± 0.5)	466.3 (± 1.8)	1230.2 (± 0.1)	1100.1 (± 0.1)	537.4 (± 0.8)
0	0%	106.6 (± 0.4)	166.0 (± 0.7)	462.5 (± 2.0)	1173.6 (± 0.1)	1045.0 (± 0.1)	526.1 (± 0.9)
82	1%	82.9 (± 0.5)	138.4 (± 0.6)	461.6 (± 1.7)	1117.3 (± 0.1)	990.5 (± 0.1)	519.7 (± 1.0)
246	3%	51.9 (± 0.5)	99.2 (± 0.7)	460.2 (± 2.0)	1007.3 (± 0.1)	883.0 (± 0.1)	508.3 (± 1.4)
410	5%	34.9 (± 0.5)	75.6 (± 0.8)	457.2 (± 1.6)	903.8 (± 0.2)	780.2 (± 0.2)	463.7 (± 1.2)
820	10%	17.4 (± 0.5)	45.6 (± 0.7)	451.9 (± 1.7)	761.7 (± 1.5)	601.6 (± 1.2)	352.3 (± 1.6)
1640	20%	4.7 (± 0.2)	24.8 (± 0.6)	440.8 (± 1.9)	750.9 (± 1.4)	582.6 (± 1.5)	333.7 (± 0.2)
2460	30%	3.2 (± 0.2)	16.2 (± 0.4)	431.2 (± 1.6)	742.1 (± 1.5)	575.4 (± 1.3)	232.4 (± 0.2)
3280	40%	2.7 (± 0.2)	12.5 (± 0.4)	419.7 (± 1.8)	733.9 (± 1.7)	568.6 (± 1.4)	198.9 (± 0.3)
4100	50%	2.5 (± 0.3)	10.3 (± 0.3)	410.1 (± 1.8)	726.9 (± 1.6)	561.5 (± 1.3)	179.1 (± 0.3)

Table B.5: The expected discounted cost ΛJ_z in millions of dollars under the proposed policy, the static policy, and the modified alternative policies (the values denoted in the parentheses indicate the 95% confidence interval).

Excess Supply			Modified				
Total	Percentage	Proposed	Static	Closest-Driver	MaxWeight	MaxWeight+	MBP
-4100	-50%	3642.5 (± 1.6)	3691.4 (± 1.6)	4081.1 (± 0.1)	4168.6 (± 0.1)	4139.7 (± 0.1)	3494.1 (± 0.1)
-3280	-40%	2837.8 (± 1.4)	2908.4 (± 2.2)	3462.7 (± 0.1)	3595.2 (± 0.1)	3544.5 (± 0.5)	2768.5 (± 0.1)
-2460	-30%	2050.7 (± 0.8)	2157.8 (± 0.3)	2833.6 (± 0.2)	3021.9 (± 0.1)	2957.6 (± 0.5)	2047.3 (± 0.1)
-1640	-20%	1300.3 (± 0.4)	1432.5 (± 0.2)	2181.6 (± 0.2)	2449.9 (± 0.1)	2378.1 (± 0.8)	1355.9 (± 0.1)
-820	-10%	610.2 (± 0.3)	723.6 (± 0.2)	1338.3 (± 1.2)	1881.1 (± 0.1)	1802.7 (± 0.5)	1142.5 (± 1.1)
-410	-5%	318.5 (± 0.3)	399.5 (± 0.2)	829.8 (± 0.8)	1598.6 (± 0.1)	1518.6 (± 0.5)	982.1 (± 0.9)
-246	-3%	218.3 (± 0.0)	289.3 (± 0.3)	633.5 (± 0.9)	1486.3 (± 0.1)	1406.1 (± 0.5)	969.0 (± 1.3)
-82	-1%	137.6 (± 0.3)	200.7 (± 0.5)	442.9 (± 0.9)	1374.7 (± 0.1)	1294.4 (± 0.4)	964.8 (± 1.4)
0	0%	106.6 (± 0.4)	166.0 (± 0.7)	355.0 (± 0.8)	1318.9 (± 0.1)	1239.4 (± 0.6)	911.1 (± 0.4)
82	1%	82.9 (± 0.5)	138.4 (± 0.6)	276.3 (± 0.6)	1263.7 (± 0.1)	1183.9 (± 0.5)	863.9 (± 0.3)
246	3%	51.9 (± 0.5)	99.2 (± 0.7)	162.6 (± 0.5)	1153.5 (± 0.1)	1074.3 (± 0.5)	758.0 (± 0.4)
410	5%	34.9 (± 0.5)	75.6 (± 0.8)	108.7 (± 0.5)	1044.6 (± 0.1)	965.8 (± 0.4)	648.0 (± 0.3)
820	10%	17.4 (± 0.5)	45.6 (± 0.7)	73.2 (± 0.3)	782.1 (± 0.1)	707.2 (± 0.6)	376.1 (± 0.3)
1640	20%	4.7 (± 0.2)	24.8 (± 0.6)	70.4 (± 0.2)	438.9 (± 0.5)	338.2 (± 3.1)	36.9 (± 0.5)
2460	30%	3.2 (± 0.2)	16.2 (± 0.4)	70.5 (± 0.2)	405.0 (± 0.6)	280.4 (± 2.1)	17.1 (± 0.5)
3280	40%	2.7 (± 0.2)	12.5 (± 0.4)	70.1 (± 0.2)	398.3 (± 0.4)	267.4 (± 1.8)	10.7 (± 0.5)
4100	50%	2.5 (± 0.3)	10.3 (± 0.3)	70.0 (± 0.3)	395.6 (± 0.3)	265.7 (± 0.5)	7.7 (± 0.2)

Table B.6: The expected discounted cost per driver J_z in thousands of dollars under the proposed policy, the static policy, and the original alternative policies (the values inside the parentheses indicate the 95% confidence interval).

Excess Supply			Original				
Total	Percentage	Proposed	Static	Closest-Driver	MaxWeight	MaxWeight+	MBP
-4100	-50%	888.4 (± 0.4)	900.3 (± 0.4)	982.0 (± 0.0)	1001.7 (± 0.0)	960.4 (± 0.0)	820.2 (± 0.0)
-3280	-40%	576.8 (± 0.3)	591.1 (± 0.4)	689.8 (± 0.0)	715.1 (± 0.0)	679.8 (± 0.0)	531.2 (± 0.0)
-2460	-30%	357.3 (± 0.1)	375.9 (± 0.1)	478.8 (± 0.0)	510.2 (± 0.0)	480.7 (± 0.0)	325.6 (± 0.0)
-1640	-20%	198.2 (± 0.1)	218.4 (± 0.0)	316.3 (± 0.0)	356.5 (± 0.0)	332.6 (± 0.0)	176.1 (± 0.0)
-820	-10%	82.7 (± 0.0)	98.0 (± 0.0)	156.7 (± 0.1)	237.3 (± 0.0)	217.8 (± 0.0)	97.3 (± 0.1)
-410	-5%	40.9 (± 0.0)	51.3 (± 0.0)	84.7 (± 0.1)	187.4 (± 0.0)	169.9 (± 0.0)	81.6 (± 0.2)
-246	-3%	27.4 (± 0.0)	36.4 (± 0.0)	64.5 (± 0.1)	169.1 (± 0.0)	152.3 (± 0.0)	73.8 (± 0.1)
-82	-1%	17.0 (± 0.0)	24.7 (± 0.1)	57.4 (± 0.2)	151.5 (± 0.0)	135.5 (± 0.0)	66.2 (± 0.1)
0	0%	13.0 (± 0.0)	20.2 (± 0.1)	56.4 (± 0.2)	143.1 (± 0.0)	127.4 (± 0.0)	64.2 (± 0.1)
82	1%	10.0 (± 0.1)	16.7 (± 0.1)	55.7 (± 0.2)	134.9 (± 0.0)	119.6 (± 0.0)	62.8 (± 0.1)
246	3%	6.1 (± 0.1)	11.7 (± 0.1)	54.5 (± 0.2)	119.3 (± 0.0)	104.5 (± 0.0)	60.2 (± 0.2)
410	5%	4.0 (± 0.1)	8.8 (± 0.1)	53.1 (± 0.2)	105.0 (± 0.0)	90.6 (± 0.0)	53.9 (± 0.1)
820	10%	1.9 (± 0.1)	5.1 (± 0.1)	50.1 (± 0.2)	84.5 (± 0.2)	66.7 (± 0.1)	39.1 (± 0.2)
1640	20%	0.5 (± 0.0)	2.5 (± 0.1)	44.8 (± 0.2)	76.3 (± 0.1)	59.2 (± 0.2)	33.9 (± 0.0)
2460	30%	0.3 (± 0.0)	1.5 (± 0.0)	40.4 (± 0.2)	69.6 (± 0.1)	54.0 (± 0.1)	21.8 (± 0.0)
3280	40%	0.2 (± 0.0)	1.1 (± 0.0)	36.6 (± 0.2)	63.9 (± 0.1)	49.5 (± 0.1)	17.3 (± 0.0)
4100	50%	0.2 (± 0.0)	0.8 (± 0.0)	33.3 (± 0.1)	59.1 (± 0.1)	45.6 (± 0.1)	14.6 (± 0.0)

Table B.7: The expected discounted cost per driver J_z in thousands of dollars under the proposed policy, the static policy, and the modified alternative policies (the values inside the parentheses indicate the 95% confidence interval).

Excess Supply			Modified				
Total	Percentage	Proposed	Static	Closest-Driver	MaxWeight	MaxWeight+	MBP
-4100	-50%	888.4 (± 0.4)	900.3 (± 0.4)	995.4 (± 0.0)	1016.7 (± 0.0)	1009.7 (± 0.0)	852.2 (± 0.0)
-3280	-40%	576.8 (± 0.3)	591.1 (± 0.4)	703.8 (± 0.0)	730.7 (± 0.0)	720.4 (± 0.1)	562.7 (± 0.0)
-2460	-30%	357.3 (± 0.1)	375.9 (± 0.1)	493.7 (± 0.0)	526.5 (± 0.0)	515.3 (± 0.1)	356.7 (± 0.0)
-1640	-20%	198.2 (± 0.1)	218.4 (± 0.0)	332.6 (± 0.0)	373.5 (± 0.0)	362.5 (± 0.1)	206.7 (± 0.0)
-820	-10%	82.7 (± 0.0)	98.0 (± 0.0)	181.3 (± 0.2)	254.9 (± 0.0)	244.3 (± 0.1)	154.8 (± 0.1)
-410	-5%	40.9 (± 0.0)	51.3 (± 0.0)	106.5 (± 0.1)	205.2 (± 0.0)	194.9 (± 0.1)	126.1 (± 0.1)
-246	-3%	27.4 (± 0.0)	36.4 (± 0.0)	79.7 (± 0.1)	186.9 (± 0.0)	176.8 (± 0.1)	121.8 (± 0.2)
-82	-1%	17.0 (± 0.0)	24.7 (± 0.1)	54.6 (± 0.1)	169.3 (± 0.0)	159.5 (± 0.1)	118.8 (± 0.2)
0	0%	13.0 (± 0.0)	20.2 (± 0.1)	43.3 (± 0.1)	160.8 (± 0.0)	151.2 (± 0.1)	111.1 (± 0.1)
82	1%	10.0 (± 0.1)	16.7 (± 0.1)	33.4 (± 0.1)	152.6 (± 0.0)	143.0 (± 0.1)	104.3 (± 0.0)
246	3%	6.1 (± 0.1)	11.7 (± 0.1)	19.2 (± 0.1)	136.6 (± 0.0)	127.2 (± 0.1)	89.8 (± 0.0)
410	5%	4.0 (± 0.1)	8.8 (± 0.1)	12.6 (± 0.1)	121.3 (± 0.0)	112.2 (± 0.0)	75.3 (± 0.0)
820	10%	1.9 (± 0.1)	5.1 (± 0.1)	8.1 (± 0.0)	86.7 (± 0.0)	78.4 (± 0.1)	41.7 (± 0.0)
1640	20%	0.5 (± 0.0)	2.5 (± 0.1)	7.2 (± 0.0)	44.6 (± 0.0)	34.4 (± 0.3)	3.7 (± 0.1)
2460	30%	0.3 (± 0.0)	1.5 (± 0.0)	6.6 (± 0.0)	38.0 (± 0.1)	26.3 (± 0.2)	1.6 (± 0.0)
3280	40%	0.2 (± 0.0)	1.1 (± 0.0)	6.1 (± 0.0)	34.7 (± 0.0)	23.3 (± 0.2)	0.9 (± 0.0)
4100	50%	0.2 (± 0.0)	0.8 (± 0.0)	5.7 (± 0.0)	32.2 (± 0.0)	21.6 (± 0.0)	0.6 (± 0.0)

Table B.8: The expected discounted cost in millions of dollars under the proposed policy and the percentage difference of the static policy as well as the original alternative policies from the proposed policy (the values inside the parentheses indicate the 95% confidence interval).

Excess Supply			Original				
Total	Percentage	Proposed	Static	Closest-Driver	MaxWeight	MaxWeight+	MBP
-4100	-50%	3642.5 (± 1.6)	1% ($\pm 0.1\%$)	10% ($\pm 0.0\%$)	12% ($\pm 0.0\%$)	8% ($\pm 0.0\%$)	-7% ($\pm 0.0\%$)
-3280	-40%	2837.8 (± 1.4)	2% ($\pm 0.2\%$)	19% ($\pm 0.1\%$)	23% ($\pm 0.1\%$)	17% ($\pm 0.1\%$)	-7% ($\pm 0.0\%$)
-2460	-30%	2050.7 (± 0.8)	5% ($\pm 0.1\%$)	34% ($\pm 0.1\%$)	42% ($\pm 0.1\%$)	34% ($\pm 0.1\%$)	-8% ($\pm 0.0\%$)
-1640	-20%	1300.3 (± 0.4)	10% ($\pm 0.0\%$)	59% ($\pm 0.1\%$)	79% ($\pm 0.1\%$)	67% ($\pm 0.1\%$)	-11% ($\pm 0.0\%$)
-820	-10%	610.2 (± 0.3)	18% ($\pm 0.1\%$)	89% ($\pm 0.3\%$)	187% ($\pm 0.2\%$)	163% ($\pm 0.1\%$)	17% ($\pm 0.3\%$)
-410	-5%	318.5 (± 0.3)	25% ($\pm 0.2\%$)	107% ($\pm 0.5\%$)	358% ($\pm 0.4\%$)	315% ($\pm 0.4\%$)	99% ($\pm 0\%$)
-246	-3%	218.3 (± 0.0)	32% ($\pm 0.3\%$)	135% ($\pm 1\%$)	515% ($\pm 0.1\%$)	454% ($\pm 0.1\%$)	168% ($\pm 0\%$)
-82	-1%	137.6 (± 0.3)	45% ($\pm 0\%$)	238% ($\pm 2\%$)	793% ($\pm 1\%$)	699% ($\pm 1\%$)	290% ($\pm 1\%$)
0	0%	106.6 (± 0.4)	55% ($\pm 1\%$)	334% ($\pm 3\%$)	1001% ($\pm 3\%$)	880% ($\pm 3\%$)	393% ($\pm 2\%$)
82	1%	82.9 (± 0.5)	66% ($\pm 1\%$)	456% ($\pm 5\%$)	1247% ($\pm 8\%$)	1094% ($\pm 7\%$)	526% ($\pm 4\%$)
246	3%	51.9 (± 0.5)	91% ($\pm 3\%$)	787% ($\pm 12\%$)	1842% ($\pm 20\%$)	1602% ($\pm 17\%$)	880% ($\pm 11\%$)
410	5%	34.9 (± 0.5)	116% ($\pm 5\%$)	1211% ($\pm 21\%$)	2492% ($\pm 39\%$)	2138% ($\pm 33\%$)	1230% ($\pm 21\%$)
820	10%	17.4 (± 0.5)	161% ($\pm 11\%$)	2490% ($\pm 76\%$)	4266% ($\pm 125\%$)	3348% ($\pm 99\%$)	1919% ($\pm 60\%$)
1640	20%	4.7 (± 0.2)	422% ($\pm 33\%$)	9208% ($\pm 415\%$)	15757% ($\pm 697\%$)	12203% ($\pm 542\%$)	6946% ($\pm 308\%$)
2460	30%	3.2 (± 0.2)	399% ($\pm 44\%$)	13181% ($\pm 979\%$)	22759% ($\pm 1680\%$)	17623% ($\pm 1303\%$)	7059% ($\pm 525\%$)
3280	40%	2.7 (± 0.2)	355% ($\pm 47\%$)	15247% ($\pm 1302\%$)	26735% ($\pm 2268\%$)	20691% ($\pm 1757\%$)	7173% ($\pm 614\%$)
4100	50%	2.5 (± 0.3)	303% ($\pm 51\%$)	16000% ($\pm 1838\%$)	28439% ($\pm 3252\%$)	21943% ($\pm 2512\%$)	6933% ($\pm 801\%$)

Table B.9: The expected discounted cost in millions of dollars under the proposed policy and the percentage difference of the static policy as well as the modified alternative policies from the proposed policy (the values inside the parentheses indicate the 95% confidence interval).

Excess Supply			Modified				
Total	Percentage	Proposed	Static	Closest-Driver	MaxWeight	MaxWeight+	MBP
-4100	-50%	3642.5 (±1.6)	1% (±0.1%)	12% (±0.0%)	14% (±0.0%)	13% (±0.0%)	-4% (±0.0%)
-3280	-40%	2837.8 (±1.4)	2% (±0.2%)	22% (±0.1%)	26% (±0.1%)	24% (±0.1%)	-2% (±0.0%)
-2460	-30%	2050.7 (±0.8)	5% (±0.1%)	38% (±0.1%)	47% (±0.1%)	44% (±0.1%)	0% (±0.0%)
-1640	-20%	1300.3 (±0.4)	10% (±0.0%)	67% (±0.1%)	88% (±0.1%)	82% (±0.1%)	4% (±0.0%)
-820	-10%	610.2 (±0.3)	18% (±0.1%)	119% (±0.4%)	208% (±0.2%)	195% (±0.2%)	87% (±0.4%)
-410	-5%	318.5 (±0.3)	25% (±0.2%)	160% (±0%)	401% (±0.5%)	376% (±0%)	208% (±0%)
-246	-3%	218.3 (±0.0)	32% (±0.3%)	190% (±0%)	580% (±0.1%)	544% (±0.4%)	343% (±1%)
-82	-1%	137.6 (±0.3)	45% (±0%)	221% (±1%)	898% (±2%)	840% (±2%)	601% (±2%)
0	0%	106.6 (±0.4)	55% (±1%)	233% (±1%)	1137% (±4%)	1063% (±4%)	754% (±3%)
82	1%	82.9 (±0.5)	66% (±1%)	233% (±2%)	1424% (±9%)	1328% (±8%)	942% (±6%)
246	3%	51.9 (±0.5)	91% (±3%)	213% (±3%)	2124% (±23%)	1971% (±21%)	1361% (±15%)
410	5%	34.9 (±0.5)	116% (±5%)	211% (±5%)	2896% (±45%)	2670% (±41%)	1758% (±28%)
820	10%	17.4 (±0.5)	161% (±11%)	319% (±12%)	4383% (±127%)	3953% (±115%)	2055% (±61%)
1640	20%	4.7 (±0.2)	422% (±33%)	1387% (±65%)	9169% (±406%)	7043% (±339%)	679% (±40%)
2460	30%	3.2 (±0.2)	399% (±44%)	2070% (±159%)	12374% (±916%)	8538% (±647%)	426% (±50%)
3280	40%	2.7 (±0.2)	355% (±47%)	2464% (±217%)	14463% (±1229%)	9678% (±835%)	292% (±47%)
4100	50%	2.5 (±0.3)	303% (±51%)	2649% (±313%)	15433% (±1768%)	10330% (±1188%)	202% (±39%)

Table B.10: Percentage of the demand dropped under the proposed policy, the static policy, and the original alternative policies (the standard deviation of all estimates is less than or equal to 0.01%).

Excess Supply			Modified				
Total	Percentage	Proposed	Static	Closest-Driver	MaxWeight	MaxWeight+	MBP
-4100	-50%	50.25%	50.51%	54.94%	55.67%	55.39%	49.18%
-3280	-40%	39.19%	39.49%	46.02%	47.09%	46.83%	39.06%
-2460	-30%	28.60%	29.03%	37.05%	38.52%	38.46%	29.02%
-1640	-20%	18.51%	19.09%	27.87%	29.98%	30.22%	19.38%
-820	-10%	9.12%	9.53%	16.17%	21.50%	21.90%	13.25%
-410	-5%	4.83%	5.22%	9.87%	17.31%	17.77%	11.74%
-246	-3%	3.40%	3.76%	8.03%	15.65%	16.14%	10.57%
-82	-1%	2.13%	2.60%	7.45%	14.01%	14.51%	9.77%
0	0%	1.64%	2.14%	7.42%	13.19%	13.71%	9.63%
82	1%	1.26%	1.78%	7.42%	12.39%	12.91%	9.58%
246	3%	0.79%	1.27%	7.42%	10.82%	11.34%	9.42%
410	5%	0.54%	0.96%	7.41%	9.33%	9.84%	8.81%
820	10%	0.28%	0.58%	7.41%	7.33%	7.22%	7.17%
1640	20%	0.25%	0.30%	7.40%	7.30%	7.09%	2.21%
2460	30%	0.10%	0.21%	7.40%	7.29%	7.08%	1.41%
3280	40%	0.06%	0.16%	7.39%	7.28%	7.07%	1.30%
4100	50%	0.05%	0.12%	7.37%	7.27%	7.06%	1.27%

Table B.11: Percentage of the demand dropped under the proposed policy, the static policy, and the modified alternative policies (the standard deviation of all estimates is less than or equal to 0.01%).

Excess Supply			Modified				
Total	Percentage	Proposed	Static	Closest-Driver	MaxWeight	MaxWeight+	MBP
-4100	-50%	50.25%	50.51%	55.06%	55.85%	56.22%	49.81%
-3280	-40%	39.19%	39.49%	46.18%	47.31%	47.67%	39.80%
-2460	-30%	28.60%	29.03%	37.26%	38.80%	39.27%	29.86%
-1640	-20%	18.51%	19.09%	28.18%	30.33%	30.98%	20.34%
-820	-10%	9.12%	9.53%	17.02%	21.92%	22.66%	17.44%
-410	-5%	4.83%	5.22%	10.42%	17.75%	18.52%	13.98%
-246	-3%	3.40%	3.76%	7.85%	16.10%	16.87%	13.84%
-82	-1%	2.13%	2.60%	5.38%	14.45%	15.24%	13.80%
0	0%	1.64%	2.14%	4.25%	13.63%	14.42%	12.61%
82	1%	1.26%	1.78%	3.26%	12.81%	13.61%	11.92%
246	3%	0.79%	1.27%	1.85%	11.19%	12.00%	10.39%
410	5%	0.54%	0.96%	1.20%	9.59%	10.41%	8.80%
820	10%	0.28%	0.58%	0.78%	5.73%	6.59%	4.89%
1640	20%	0.25%	0.30%	0.75%	0.70%	1.19%	0.49%
2460	30%	0.10%	0.21%	0.75%	0.20%	0.36%	0.22%
3280	40%	0.06%	0.16%	0.75%	0.11%	0.17%	0.13%
4100	50%	0.05%	0.12%	0.75%	0.08%	0.13%	0.09%

B.2 Miscellaneous Proofs

Proof of Lemma 1. The proof closely follows the steps outlined in Bramson and Williams [54, Theorem 6.1]. First, we show that $\dim \mathcal{R} = b + \tilde{b} - m$. Then, using $\mathcal{M} = \mathcal{R}^\perp$, we arrive at $\dim \mathcal{M} = q - (b + \tilde{b} - m)$, which concludes the proof. By definition,

$$\mathcal{R} = \{Hy_B + \tilde{H}\tilde{y}_B : By_B = 0, y_B \in \mathbb{R}^b, \tilde{y}_B \in \mathbb{R}^{\tilde{b}}\}. \quad (\text{B.48})$$

Next, we show that the dimension of \mathcal{R} is equal to

$$\dim\{(y_B, \tilde{y}_B) : By_B = 0, y_B \in \mathbb{R}^b, \tilde{y}_B \in \mathbb{R}^{\tilde{b}}\} = b + \tilde{b} - \text{Rank}(B). \quad (\text{B.49})$$

It suffices to prove that $Hy_B + \tilde{H}\tilde{y}_B = 0$ and $By_B = 0$ only holds for $(y_B, \tilde{y}_B) = 0$. We prove this by contradiction. Assume that $Hy_B + \tilde{H}\tilde{y}_B = 0$ and $By_B = 0$ for some $(y_B, \tilde{y}_B) \neq 0$. Extend (y_B, \tilde{y}_B) to (y, \tilde{y}) , where $y \in \mathbb{R}^n$ and $\tilde{y} \in \mathbb{R}^{\tilde{n}}$, by setting $(y_N, \tilde{y}_N) = 0$. Then, since $y_N = 0$ and $By_B = 0$, we have $Ay = 0$. Similarly, since $(y_N, \tilde{y}_N) = 0$ and $Hy_B + \tilde{H}\tilde{y}_B = 0$, we have $Ry + \tilde{R}\tilde{y} = 0$. Since (x^*, \tilde{x}^*) is the optimal solution to the Static Planning Problem (2.14)-(2.17), it satisfies $x_B^* > 0$, $\tilde{x}_B^* > 0$, $Rx^* + \tilde{R}\tilde{x}^* = 0$, and $Ax^* = e$. Therefore, there exists a small enough $\epsilon > 0$ such that $(x, \tilde{x}) = (x^* + \epsilon y, \tilde{x}^* + \epsilon \tilde{y})$ satisfies $x \geq 0$, $\tilde{x} \geq 0$, $Rx + \tilde{R}\tilde{x} = 0$, and $Ax = e$. This contradicts the uniqueness of the optimal solution; see Assumption 1. Thus, such a $(y_B, \tilde{y}_B) \neq 0$ does not exist and the dimension of \mathcal{R} is equal to $b + \tilde{b} - \text{Rank}(B)$.

Now, it remains to prove that $\text{Rank}(B) = m$. This follows from the fact that B has more rows than columns and each column of B has exactly one non-zero component; see (2.3). \square

Proof of Lemma 2. The proof consists of two steps. In the first step, we show that the space of reversible displacements \mathcal{R} is spanned by the vectors $e^i - e^{i'}$ for buffers i and i' that communicate directly. In the second step, we use the result from the first step and $\mathcal{M} = \mathcal{R}^\perp$ to show that the rows of the matrix M defined in Lemma 2 form a basis for \mathcal{M} .

Step 1. Let

$$\mathcal{C}_1 = \{(j, j') : j \text{ and } j' \text{ are basic dispatch activities with } k(j) = k(j')\},$$

$$\tilde{\mathcal{C}}_1 = \{j : j \text{ is a basic relocation activity}\}.$$

Let e^i denote the unit vector (with appropriate dimension, which will be clear from the context) with a one in its i -th component and zeros elsewhere. We define $\mathcal{C}_2 = \{e^j - e^{j'} : (j, j') \in \mathcal{C}_1, e^j \in \mathbb{R}^n, e^{j'} \in \mathbb{R}^n\}$ and $\tilde{\mathcal{C}}_2 = \{e^j : j \in \tilde{\mathcal{C}}_1, e^j \in \mathbb{R}^{\tilde{n}}\}$. Then,

$$\text{Span}(\mathcal{C}_2) = \{y : Ay = 0, y_N = 0, y \in \mathbb{R}^n\} \quad \text{and} \quad \text{Span}(\tilde{\mathcal{C}}_2) = \{\tilde{y} : \tilde{y}_N = 0, \tilde{y} \in \mathbb{R}^{\tilde{n}}\}. \quad (\text{B.50})$$

To see the first equality, consider $y \in \{y : Ay = 0, y_N = 0, y \in \mathbb{R}^n\}$. By the definition of A in Equation (2.3) and the fact that $By_B = 0$, y can be written as a linear combination of the vectors in \mathcal{C}_2 . Similarly, since any linear combination of the vectors in \mathcal{C}_2 satisfies $By_B = 0$, it belongs to $\{y : Ay = 0, y_N = 0, y \in \mathbb{R}^n\}$. This proves the first equality in (B.50). The second equality follows from the fact that $\tilde{\mathcal{C}}_2$ is a basis for $\{\tilde{y} : \tilde{y}_N = 0, \tilde{y} \in \mathbb{R}^{\tilde{n}}\}$. The space of reversible displacements \mathcal{R} can be written as

$$\begin{aligned} \mathcal{R} &= \{Hy_B + \tilde{H}\tilde{y}_B : By_B = 0, y \in \mathbb{R}^n, \tilde{y} \in \mathbb{R}^{\tilde{n}}\} \\ &= \{Ry + \tilde{R}\tilde{y} : Ay = 0, y_N = 0, \tilde{y}_N = 0, y \in \mathbb{R}^n, \tilde{y} \in \mathbb{R}^{\tilde{n}}\} \\ &= \{Ry : Ay = 0, y_N = 0, y \in \mathbb{R}^n\} \oplus \{\tilde{R}\tilde{y} : \tilde{y}_N = 0, \tilde{y} \in \mathbb{R}^{\tilde{n}}\}. \end{aligned} \quad (\text{B.51})$$

It follows from (B.50)-(B.51) that $\mathcal{R} = \text{Span}(\mathcal{C}_3 \cup \tilde{\mathcal{C}}_3)$, where

$$\begin{aligned} \mathcal{C}_3 &= \{Ry : y \in \mathcal{C}_2\} = \{R(e^j - e^{j'}) : (j, j') \in \mathcal{C}_1, e^j \in \mathbb{R}^n, e^{j'} \in \mathbb{R}^n\} \\ &= \{\mu_{k(j)}(e^{i(j)} - e^{i(j')}) : (j, j') \in \mathcal{C}_1, e^{i(j)} \in \mathbb{R}^q, e^{i(j')} \in \mathbb{R}^q\} \end{aligned} \quad (\text{B.52})$$

and

$$\begin{aligned}
\tilde{\mathcal{C}}_3 &= \{\tilde{R}\tilde{y} : \tilde{y} \in \tilde{\mathcal{C}}_2\} \\
&= \{\tilde{R}e^j : j \in \tilde{\mathcal{C}}_1, e^j \in \mathbb{R}^{\tilde{n}}\} \\
&= \{e^{o(j)} - e^{d(j)} : j \in \tilde{\mathcal{C}}_1, e^{o(j)} \in \mathbb{R}^q, e^{d(j)} \in \mathbb{R}^q\}.
\end{aligned} \tag{B.53}$$

Equation (B.52) follows from (2.1) and the fact that $P_{jl} = P_{j'l}$ for $j, j' = 1, \dots, n$ with $k(j) = k(j')$ and $l = 1, \dots, q$. The last equality in (B.53) follows from (2.5). Let

$$\bar{\mathcal{C}} = \{e^i - e^{i'} : \text{buffers } i \text{ and } i' \text{ communicate directly, } e^i \in \mathbb{R}^q, e^{i'} \in \mathbb{R}^q\}. \tag{B.54}$$

Since $\mu_{k(j)} > 0$ for $j = 1, \dots, n$ and $\mathcal{R} = \text{Span}(\mathcal{C}_3 \cup \tilde{\mathcal{C}}_3)$, it follows from Definition 2 that $\mathcal{R} = \text{Span}(\bar{\mathcal{C}})$.

Step 2. Let \bar{p} denote the number of buffer pools and define $a^l \in \mathbb{R}^q$ for $l = 1, \dots, \bar{p}$ to be vectors such that for $i = 1, \dots, q$, we have

$$a_i^l = \begin{cases} 1 & \text{if buffer } i \text{ belongs to buffer pool } l, \\ 0 & \text{otherwise.} \end{cases} \tag{B.55}$$

Note that $a^1, \dots, a^{\bar{p}}$ are the rows of the matrix M defined in Lemma 2. If we show that $\mathcal{M} = \text{Span}(a^1, \dots, a^{\bar{p}})$, it would follow from the linear independence of $a^1, \dots, a^{\bar{p}}$ that $a^1, \dots, a^{\bar{p}}$ form a basis for \mathcal{M} . Moreover, the workload dimension would be equal to the number of buffer pools, i.e., $p = \bar{p}$ and the matrix M whose rows are a^1, \dots, a^p would be a workload matrix. It just remains to show that $\mathcal{M} = \text{Span}(a^1, \dots, a^{\bar{p}})$.

As a preliminary to showing that $\mathcal{M} = \text{Span}(a^1, \dots, a^{\bar{p}})$, note that since $\mathcal{M} = \mathcal{R}^\perp$ and $\mathcal{R} = \text{Span}(\bar{\mathcal{C}})$,

$$a \in \mathcal{M} \Leftrightarrow a \cdot z = 0 \text{ for all } z \in \bar{\mathcal{C}}. \tag{B.56}$$

It follows from (B.56) that if $a \in \mathcal{M}$, we have $a_i = a_{i'}$ for buffers i and i' that communicate. To see this, consider $a \in \mathcal{M}$. By (B.54) and (B.56), we have $a_i = a_{i'}$ for buffers i and i' that communicate directly. Then, it follows from Definition 2 that $a_i = a_{i'}$ for buffers i and i' that communicate, i.e., belong to the same buffer pool. Consequently, $\mathcal{M} \subseteq \text{Span}(a^1, \dots, a^{\bar{p}})$.

Next, we show that $\text{Span}(a^1, \dots, a^{\bar{p}}) \subseteq \mathcal{M}$. To prove this, it suffices to show that $a^l \in \mathcal{M}$ for $l = 1, \dots, \bar{p}$. By (B.54) and (B.56), this holds if and only if $a^l \cdot (e^i - e^{i'}) = 0$ for $l = 1, \dots, \bar{p}$ and buffers i and i' that communicate directly. This trivially holds since for $l = 1, \dots, \bar{p}$, we have

$$a^l \cdot (e^i - e^{i'}) = a_i^l - a_{i'}^l = 0,$$

where the second equality follows from (B.55) and the fact that buffers i and i' belong to the same buffer pool; see Definition 2. Combining $\mathcal{M} \subseteq \text{Span}(a^1, \dots, a^{\bar{p}})$ and $\text{Span}(a^1, \dots, a^{\bar{p}}) \subseteq \mathcal{M}$, we arrive at the desired result, i.e., $\mathcal{M} = \text{Span}(a^1, \dots, a^{\bar{p}})$. \square

Proof of Corollary 1. By the definition of \mathcal{W} , we have

$$\begin{aligned} \mathcal{W} &= \{Mz : z \in \mathcal{S}\} \\ &= \{w : w_l = \sum_{i \in \mathcal{P}_l} z_i \text{ for } l = 1, \dots, p, z'e = 1, z \in \mathbb{R}_+^q\} \\ &= \{w : w'e = 1, w \in \mathbb{R}_+^p\}, \end{aligned}$$

where the second equality follows from Lemma 2. \square

Proof of Lemma 4. By (2.39), for $w \in \mathcal{W}$, we have

$$g(w) = \inf_z \{\alpha \pi' z : Mz = w, z \in \mathcal{S}\} \quad (B.57)$$

$$\begin{aligned} &= \inf_z \{\alpha \pi' z : \sum_{i \in \mathcal{P}_l} z_i = w_l \text{ for } l = 1, \dots, p, z \geq 0\} \\ &= \alpha \sum_{l=1}^p \pi_{i^*(l)} w_l, \end{aligned} \quad (B.58)$$

where (B.57) follows from Lemma 2 and (B.58) follows from (2.40). \square

Proof of Lemma 5. We prove that $L : \mathcal{W} \rightarrow \hat{\mathcal{W}}$ is one-to-one by contradiction. Assume that $L : \mathcal{W} \rightarrow \hat{\mathcal{W}}$ is not one-to-one. Then, there exist distinct $w, \bar{w} \in \mathcal{W}$ such that $Lw = L\bar{w}$. By the definition of L , since $w \neq \bar{w}$ and $Lw = L\bar{w}$, we must have $w_p \neq \bar{w}_p$. This gives, $w'e \neq \bar{w}'e$, which contradicts the fact that $w, \bar{w} \in \mathcal{W}$, i.e., $w'e = \bar{w}'e = 1$. Therefore, the assumption that $L : \mathcal{W} \rightarrow \hat{\mathcal{W}}$ is not one-to-one is incorrect. To see that $L : \mathcal{W} \rightarrow \hat{\mathcal{W}}$ is onto, consider an arbitrarily chosen $\hat{w} \in \hat{\mathcal{W}}$ and let $w = (\hat{w}_1, \dots, \hat{w}_{p-1}, 1 - \sum_{l=1}^{p-1} \hat{w}_l)$. We observe that $w \in \mathcal{W}$ and $Lw = \hat{w}$. \square

Proof of Lemma 6. Assume l is a dominated pushing control and $\hat{w} \in \hat{\mathcal{W}}$ is an arbitrarily chosen initial workload. We want to show that for any control process \hat{U}^1 with $\hat{U}_l^1 \not\equiv 0$ that satisfies (2.48)-(2.50), there exists a control process \hat{U}^2 with $\hat{U}_l^2 \equiv 0$ that satisfies (2.48)-(2.50) and $\hat{J}_{\hat{w}}(\hat{U}^2) \leq \hat{J}_{\hat{w}}(\hat{U}^1)$. First, we introduce a candidate control process \hat{U}^2 . Then, we show that it satisfies $\hat{U}_l^2 \equiv 0$, (2.48)-(2.50), and $\hat{J}_{\hat{w}}(\hat{U}^2) \leq \hat{J}_{\hat{w}}(\hat{U}^1)$.

Since l is a dominated pushing control, By Definition 3, there exist a subset $\mathcal{L} \subset \{1, \dots, d\} \setminus \{l\}$ of pushing controls and scalars $\hat{u}_k > 0$ for $k \in \mathcal{L}$ that jointly satisfy $\hat{G}^l = \sum_{k \in \mathcal{L}} \hat{G}^k \hat{u}_k$ and $\hat{\kappa}_l \geq \sum_{k \in \mathcal{L}} \hat{\kappa}_k \hat{u}_k$. For ease of notation, define $\hat{u} \in \mathbb{R}^d$ to be a vector such that its k -th component is equal to the scalar \hat{u}_k (discussed above) for $k \in \mathcal{L}$ and its remaining components

are zero. Then,

$$\hat{G}^l = \hat{G}\hat{u} \quad \text{and} \quad \hat{\kappa}_l \geq \hat{\kappa}'\hat{u}. \quad (\text{B.59})$$

Let \hat{W} denote the state process associated with \hat{U}^1 and define

$$\hat{U}^2(t) = \hat{U}^1(t) + (\hat{u} - e^l) \hat{U}_l^1(t), \quad t \geq 0. \quad (\text{B.60})$$

It follows from (B.60) that $\hat{U}_l^2 \equiv 0$. Moreover,

$$\begin{aligned} \hat{G}\hat{U}^2(t) &= \hat{G}\hat{U}^1(t) + \hat{G}(\hat{u} - e^l) \hat{U}_l^1(t) \\ &= \hat{G}\hat{U}^1(t) + (\hat{G}\hat{u} - \hat{G}^l) \hat{U}_l^1(t) \\ &= \hat{G}\hat{U}^1(t), \quad t \geq 0, \end{aligned} \quad (\text{B.61})$$

where the last equality follows from (B.59). By Equation (B.61), the control process \hat{U}^2 satisfies (2.48) and (2.50) with state process \hat{W} , and by Equation (B.60) and the fact that $\hat{u} \geq 0$, it satisfies (2.49). It just remains to show that $\hat{J}_{\hat{w}}(\hat{U}^2) \leq \hat{J}_{\hat{w}}(\hat{U}^1)$. This follows from

$$\hat{J}(\hat{U}^2) = \mathbb{E} \left[\int_0^\infty e^{-\alpha t} h(\hat{W}_t) dt \right] + \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(\kappa' \hat{U}^2)(t) \right] \quad (\text{B.62})$$

$$= \mathbb{E} \left[\int_0^\infty e^{-\alpha t} h(\hat{W}_t) dt \right] + \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(\kappa' \hat{U}^1 + \kappa'(\hat{u} - e^l) \hat{U}_l^1)(t) \right] \quad (\text{B.63})$$

$$= \mathbb{E} \left[\int_0^\infty e^{-\alpha t} h(\hat{W}_t) dt \right] + \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(\kappa' \hat{U}^1)(t) \right] + \mathbb{E} \left[\int_0^\infty e^{-\alpha t} (\kappa' \hat{u} - \kappa_l) d\hat{U}_l^1(t) \right] \quad (\text{B.64})$$

$$= \hat{J}(\hat{U}^1),$$

where (B.62) follows from (2.46), (B.63) follows from (B.60), and (B.64) follows from (B.59). \square

Proof of Proposition 3. The proof resembles the proof of Fleming and Soner [90, Lemma 3.2] and Kumar and Muthuraman [145, Proposition 2]. Since $\hat{\mathcal{W}}$, the domain of the value function, is a convex polyhedron, to prove the convexity of the optimal value function, it suffices to show that for any two arbitrarily chosen initial workloads $w^1, w^2 \in \hat{\mathcal{W}}$ and scalar $\delta \in [0, 1]$, we have

$$\phi(\delta w^1 + (1 - \delta)w^2) \leq \delta\phi(w^1) + (1 - \delta)\phi(w^2). \quad (\text{B.65})$$

We prove this in two steps. First, we show that for any $\hat{U}^1 \in \mathcal{A}(w^1)$ and $\hat{U}^2 \in \mathcal{A}(w^2)$, there exists $\hat{U} \in \mathcal{A}(w)$, where $w = \delta w^1 + (1 - \delta)w^2$, such that $\hat{J}_w(\hat{U}) = \delta \hat{J}_{w^1}(\hat{U}^1) + (1 - \delta) \hat{J}_{w^2}(\hat{U}^2)$. Then, we use the definition of the optimal value function in (2.53) to arrive at (B.65).

Consider two arbitrarily chosen controls $\hat{U}^1 \in \mathcal{A}(w^1)$ and $\hat{U}^2 \in \mathcal{A}(w^2)$ and let \hat{W}^1 and \hat{W}^2 denote their respective state processes. For $t \geq 0$, define

$$w = \delta w^1 + (1 - \delta)w^2, \quad \hat{W}(t) = \delta \hat{W}^1(t) + (1 - \delta) \hat{W}^2(t), \quad \hat{U}(t) = \delta \hat{U}^1(t) + (1 - \delta) \hat{U}^2(t). \quad (\text{B.66})$$

By the convexity of $\hat{\mathcal{W}}$ and (B.66), we have $w \in \hat{\mathcal{W}}$ and $\hat{W}_t \in \hat{\mathcal{W}}$ for $t \geq 0$. Moreover, it follows from (B.66) that \hat{U} satisfies (2.48)-(2.51) with initial workload w and state process \hat{W} . Therefore, $\hat{U} \in \mathcal{A}(w)$. By Equation (2.43) and Lemma 4, the holding cost function h is linear. Therefore,

$$h(\delta \hat{W}_t^1 + (1 - \delta) \hat{W}_t^2) = \delta h(\hat{W}_t^1) + (1 - \delta) h(\hat{W}_t^2), \quad t \geq 0. \quad (\text{B.67})$$

Then, by (2.46), we have

$$\begin{aligned}
\hat{J}_w(\hat{U}) &= \mathbb{E} \left[\int_0^\infty e^{-\alpha t} h(\hat{W}_t) dt \right] + \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(\kappa' \hat{U})(t) \right] \\
&= \mathbb{E} \left[\int_0^\infty e^{-\alpha t} h(\delta \hat{W}_t^1 + (1-\delta) \hat{W}_t^2) dt \right] \\
&\quad + \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(\kappa' \delta \hat{U}^1 + \kappa' (1-\delta) \hat{U}^2)(t) \right] \tag{B.68}
\end{aligned}$$

$$\begin{aligned}
&= \delta \mathbb{E} \left[\int_0^\infty e^{-\alpha t} h(\hat{W}_t^1) dt \right] + (1-\delta) \mathbb{E} \left[\int_0^\infty e^{-\alpha t} h(\hat{W}_t^2) dt \right] \\
&\quad + \delta \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(\kappa' \hat{U}^1)(t) \right] + (1-\delta) \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(\kappa' \hat{U}^2)(t) \right] \tag{B.69}
\end{aligned}$$

$$= \delta \hat{J}_{w^1}(\hat{U}^1) + (1-\delta) \hat{J}_{w^2}(\hat{U}^2), \tag{B.70}$$

where (B.68) follows from (B.66) and (B.69) follows from (B.67). Thus,

$$\begin{aligned}
\delta \phi(w^1) + (1-\delta) \phi(w^2) &= \delta \inf_{\hat{U}^1 \in \mathcal{A}(w^1)} \hat{J}_{w^1}(\hat{U}^1) + (1-\delta) \inf_{\hat{U}^2 \in \mathcal{A}(w^2)} \hat{J}_{w^2}(\hat{U}^2) \\
&= \inf_{\substack{\hat{U}^1 \in \mathcal{A}(w^1) \\ \hat{U}^2 \in \mathcal{A}(w^2)}} (\delta \hat{J}_{w^1}(\hat{U}^1) + (1-\delta) \hat{J}_{w^2}(\hat{U}^2)) \\
&= \inf_{\substack{\hat{U}^1 \in \mathcal{A}(w^1) \\ \hat{U}^2 \in \mathcal{A}(w^2)}} \hat{J}_w(\delta \hat{U}^1 + (1-\delta) \hat{U}^2) \tag{B.71}
\end{aligned}$$

$$\begin{aligned}
&\geq \inf_{U \in \mathcal{A}(w)} \hat{J}_w(U) \\
&= \phi(w), \tag{B.72}
\end{aligned}$$

where (B.71) follows from (B.70) and (B.72) follows from the fact that $\{\delta \hat{U}^1 + (1-\delta) \hat{U}^2 : \hat{U}^1 \in \mathcal{A}(w^1), \hat{U}^2 \in \mathcal{A}(w^2)\} \subseteq \mathcal{A}(w)$. \square

Proof of Proposition 4.

Step 1. The proof of Equation (2.56) resembles the proof of Kumar and Muthuraman [145, Proposition 1]. Suppose the bounded $f \in C^2(\hat{\mathcal{W}})$ and $\epsilon \geq 0$ jointly satisfy Equations (2.54)-(2.55). Consider an initial workload $w \in \hat{W}$ and let $\hat{U} \in \mathcal{A}(w)$ be an optimal control for the initial workload w . First, we use the generalized Ito's formula to write $\hat{J}_w(\hat{U}) - f(w)$ in

terms of the sum of a few integrals. Then, we use (2.54)-(2.55) and the optimality of \hat{U} to arrive at the desired result. By Lemma 18, we have

$$\begin{aligned} -f(w) = & \mathbb{E} \left[\int_0^\infty e^{-\alpha t} (\mathcal{L}f - \alpha f)(\hat{W}(t)) dt \right] + \mathbb{E} \left[\sum_{l=1}^d \int_0^\infty e^{-\alpha t} \nabla f(\hat{W}(t))' \hat{G}^l d\hat{U}_l^c(t) \right] \\ & + \mathbb{E} \left[\sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} (f(\hat{W}(t^-) + \hat{G} \Delta \hat{U}(t)) - f(\hat{W}(t^-))) \right], \end{aligned} \quad (\text{B.73})$$

where \hat{U}^c denotes the continuous part of \hat{U} , $\Delta \hat{U}$ denotes its jumps, and $\mathcal{D}(\hat{U})$ denotes the set of its discontinuities. To be specific, for $t \geq 0$,

$$\Delta \hat{U}(t) = \hat{U}(t) - \hat{U}(t^-), \quad \mathcal{D}(\hat{U}) = \{t : \Delta \hat{U}(t) \neq 0\}, \quad \text{and} \quad \hat{U}^c(t) = \hat{U}(t) - \sum_{s \in \mathcal{D}(\hat{U}), s \leq t} \Delta \hat{U}(s).$$

Adding the left-hand side of Equation (2.46) to the left-hand side of Equation (B.73) and the right-hand side of Equation (2.46) to the right-hand side of Equation (B.73) gives

$$\begin{aligned} \hat{J}_w(\hat{U}) - f(w) = & \mathbb{E} \left[\int_0^\infty e^{-\alpha t} (\mathcal{L}f - \alpha f + h)(\hat{W}(t)) dt \right] \\ & + \mathbb{E} \left[\sum_{l=1}^d \int_0^\infty e^{-\alpha t} (\nabla f(\hat{W}(t))' \hat{G}^l + \kappa_l) d\hat{U}_l^c(t) \right] \\ & + \mathbb{E} \left[\sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} (f(\hat{W}(t^-) + \hat{G} \Delta \hat{U}(t)) - f(\hat{W}(t^-)) + \kappa' \Delta \hat{U}(t)) \right]. \end{aligned} \quad (\text{B.74})$$

Next, we provide a lower bound for the right-hand side of Equation (B.74). The lower bound for the first two terms is straightforward and follows from (2.54)-(2.55). In particular,

$$\mathbb{E} \left[\int_0^\infty e^{-\alpha t} (\mathcal{L}f - \alpha f + h)(\hat{W}(t)) dt \right] \geq -\epsilon \int_0^\infty e^{-\alpha t} dt \quad (\text{B.75})$$

and

$$\mathbb{E} \left[\sum_{l=1}^d \int_0^\infty e^{-\alpha t} (\nabla f(\hat{W}(t))' \hat{G}^l + \kappa_l) d\hat{U}_l^c(t) \right] \geq -\epsilon \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(e' \hat{U}^c)(t) \right]. \quad (\text{B.76})$$

For the third term on the right-hand side of Equation (B.74), we have

$$\begin{aligned} & \mathbb{E} \left[\sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} (f(\hat{W}(t^-) + \hat{G} \Delta \hat{U}(t)) - f(\hat{W}(t^-)) + \kappa' \Delta \hat{U}(t)) \right] \\ &= \mathbb{E} \left[\sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} \int_0^{\|\Delta \hat{U}(t)\|} \left(\nabla f(\hat{W}(t^-) + \hat{G} \frac{\Delta \hat{U}(t)}{\|\Delta \hat{U}(t)\|} z)' \hat{G} \frac{\Delta \hat{U}(t)}{\|\Delta \hat{U}(t)\|} + \kappa' \frac{\Delta \hat{U}(t)}{\|\Delta \hat{U}(t)\|} \right) dz \right], \quad (\text{B.77}) \end{aligned}$$

where $\|\cdot\|$ denotes the Euclidean norm. Substituting (2.55) into (B.77) gives

$$\begin{aligned} & \mathbb{E} \left[\sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} (f(\hat{W}(t^-) + \hat{G} \Delta \hat{U}(t)) - f(\hat{W}(t^-)) + \kappa' \Delta \hat{U}(t)) \right] \\ & \geq \mathbb{E} \left[\sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} \int_0^{\|\Delta \hat{U}(t)\|} -\epsilon e' \frac{\Delta \hat{U}(t)}{\|\Delta \hat{U}(t)\|} dz \right] \\ &= \mathbb{E} \left[-\epsilon \sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} e' \Delta \hat{U}(t) \right]. \quad (\text{B.78}) \end{aligned}$$

Substituting (B.75)-(B.76) and (B.78) into (B.74) and rearranging the terms gives the following bound:

$$f(w) \leq \hat{J}_w(\hat{U}) + \epsilon \int_0^\infty e^{-\alpha t} dt + \epsilon \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(e' \hat{U}^c)(t) + \sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} e' \Delta \hat{U}(t) \right]. \quad (\text{B.79})$$

Since \hat{U} is an optimal control for the initial workload w , we have $\hat{J}_w(\hat{U}) = \phi(w)$. Substituting this into Equation (B.79) gives

$$\begin{aligned} f(w) &\leq \phi(w) + \epsilon \int_0^\infty e^{-\alpha t} dt + \epsilon \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(e' \hat{U}^c)(t) + \sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} e' \Delta \hat{U}(t) \right] \\ &= \phi(w) + \epsilon \int_0^\infty e^{-\alpha t} dt + \epsilon \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(e' \hat{U})(t) \right], \end{aligned} \quad (\text{B.80})$$

where the right-hand side is finite by (2.51). Defining

$$\eta(w) = \int_0^\infty e^{-\alpha t} dt + \mathbb{E} \left[\int_0^\infty e^{-\alpha t} d(e' \hat{U})(t) \right]$$

and substituting it into (B.80) gives (2.56). \square

B.3 Auxiliary Results

This section discusses an auxiliary result that is used multiple times in Appendices B.1-B.2. As a preliminary to the result, consider an admissible control \hat{U} to the EWF (2.47)-(2.50). Let \hat{U}^c denote the continuous part of \hat{U} , $\Delta \hat{U}$ denote its jumps, and $\mathcal{D}(\hat{U})$ denote the set of its discontinuities. To be specific, for $t \geq 0$,

$$\Delta \hat{U}(t) = \hat{U}(t) - \hat{U}(t^-), \quad \mathcal{D}(\hat{U}) = \{t : \Delta \hat{U}(t) \neq 0\}, \quad \hat{U}^c(t) = \hat{U}(t) - \sum_{s \in \mathcal{D}(\hat{U}), s \leq t} \Delta \hat{U}(s). \quad (\text{B.81})$$

Lemma 18 extends Taksar [196, Lemma 7.1] to our setting.

Lemma 18. *Let \hat{U} be an admissible control to the EWF (2.47)-(2.50) and \hat{W} denote the*

associated state process. Then, for $f \in C^2(\hat{\mathcal{W}})$ and $w \in \hat{\mathcal{W}}$, we have

$$\begin{aligned} -f(w) = & \mathbb{E} \left[\int_0^\infty e^{-\alpha t} (\mathcal{L}f - \alpha f)(\hat{W}(t)) dt \right] + \mathbb{E} \left[\sum_{l=1}^d \int_0^\infty e^{-\alpha t} \nabla f(\hat{W}(t))' \hat{G}^l d\hat{U}_l^c(t) \right] \\ & + \mathbb{E} \left[\sum_{t \in \mathcal{D}(\hat{U})} e^{-\alpha t} (f(\hat{W}(t^-) + \hat{G} \Delta \hat{U}(t)) - f(\hat{W}(t^-))) \right]. \end{aligned} \quad (\text{B.82})$$

Proof. First, we use generalized Ito's formula to write f in terms of sum of a few integrals. Then, we use the Martingale property of the stochastic integral to simplify the equation. Finally, we take the limit as $t \rightarrow \infty$ and show that the various integrals converge to the integrals in (B.82). By the generalized Ito's formula (see e.g., Øksendal and Sulem [170, Section 1.2]), we have

$$\begin{aligned} \mathbb{E} \left[e^{-\alpha t} f(\hat{W}(t)) \right] - f(w) = & \mathbb{E} \left[\int_0^t e^{-\alpha s} \nabla f(\hat{W}(s)) d\hat{\chi}(s) \right] + \mathbb{E} \left[\int_0^t e^{-\alpha s} (\mathcal{L}f - \alpha f)(\hat{W}(s)) ds \right] \\ & + \mathbb{E} \left[\sum_{l=1}^d \int_0^t e^{-\alpha s} \nabla f(\hat{W}(s))' \hat{G}^l d\hat{U}_l^c(s) \right] \\ & + \mathbb{E} \left[\sum_{s \in \mathcal{D}(\hat{U}), s \leq t} e^{-\alpha s} (f(\hat{W}(s)) - f(\hat{W}(s^-))) \right], \quad t \geq 0. \end{aligned} \quad (\text{B.83})$$

Since $f \in C^2(\hat{\mathcal{W}})$ and $\hat{\mathcal{W}}$ is a compact set, f , ∇f , and $\mathcal{L}f$ are bounded. Thus,

$$\int_0^t e^{-\alpha s} \nabla f(\hat{W}(s)) d\hat{\chi}(s)$$

is a square integrable Martingale. Therefore, its expectation, i.e., the first term on the

right-hand side of (B.83), is zero. Substituting this into (B.83) gives

$$\begin{aligned}
\mathbb{E}\left[e^{-\alpha t}f(\hat{W}(t))\right] - f(w) &= \mathbb{E}\left[\int_0^t e^{-\alpha s}(\mathcal{L}f - \alpha f)(\hat{W}(s))ds\right] \\
&\quad + \mathbb{E}\left[\sum_{l=1}^d \int_0^t e^{-\alpha s} \nabla f(\hat{W}(s))' \hat{G}^l d\hat{U}_l^c(s)\right] \\
&\quad + \mathbb{E}\left[\sum_{s \in \mathcal{D}(\hat{U}), s \leq t} e^{-\alpha s} (f(\hat{W}(s)) - f(\hat{W}(s^-)))\right], \quad t \geq 0. \quad (\text{B.84})
\end{aligned}$$

Then, substituting $W(s) = W(s^-) + \hat{G}\Delta\hat{U}$ for $s \in \mathcal{D}(\hat{U})$ into (B.84) gives

$$\begin{aligned}
\mathbb{E}\left[e^{-\alpha t}f(\hat{W}(t))\right] - f(w) &= \mathbb{E}\left[\int_0^t e^{-\alpha s}(\mathcal{L}f - \alpha f)(\hat{W}(s))ds\right] \\
&\quad + \mathbb{E}\left[\sum_{l=1}^d \int_0^t e^{-\alpha s} \nabla f(\hat{W}(s))' \hat{G}^l d\hat{U}_l^c(s)\right] \\
&\quad + \mathbb{E}\left[\sum_{s \in \mathcal{D}(\hat{U}), s \leq t} e^{-\alpha s} (f(\hat{W}(s^-) + \hat{G}\Delta\hat{U}(s)) - f(\hat{W}(s^-)))\right] \quad (\text{B.85})
\end{aligned}$$

for $t \geq 0$. It just remains to show that as $t \rightarrow \infty$, the left-hand side of (B.85) converges to the left-hand side of (B.82) and the right-hand side of (B.85) converges to the right-hand side of (B.82). Since f is bounded, the first term on the left-hand side of (B.85) converges to zero. Similarly, since $\mathcal{L}f - \alpha f$ is bounded, it follows from the dominated convergence theorem that the first term on the right-hand side of (B.85) converges to the first term on the right-hand side of (B.82). Finally, since ∇f is bounded, we can write

$$\begin{aligned}
&\left| \sum_{l=1}^d \int_0^t e^{-\alpha s} \nabla f(\hat{W}(s))' \hat{G}^l d\hat{U}_l^c(s) \right| + \left| \sum_{s \in \mathcal{D}(\hat{U}), s \leq t} e^{-\alpha s} (f(\hat{W}(s^-) + \hat{G}\Delta\hat{U}(s)) - f(\hat{W}(s^-))) \right| \\
&\leq K \left| \sum_{l=1}^d \int_0^t e^{-\alpha s} d\hat{U}_l^c(s) \right| + \left| \sum_{s \in \mathcal{D}(\hat{U}), s \leq t} e^{-\alpha s} e' \Delta\hat{U}(s) \right| \\
&\leq K \left| \int_0^t e^{-\alpha s} d(e' \hat{U})(s) \right|, \quad t \geq 0 \quad (\text{B.86})
\end{aligned}$$

for some finite K . Then, it follows from (2.51), (B.86), and the dominated convergence theorem that the second and third terms on the right-hand side of (B.85) converge to the second and third terms on the right-hand side of (B.82). \square

APPENDIX C

APPENDIX OF CHAPTER 3

C.1 Riccati Equation

This section discusses a special case of the Riccati equation and its solution. First, we introduce the Riccati equation. Then, we provide its general solution in terms of the Bessel functions. Finally, we use this characterization and the relationship between the Bessel functions and the Airy functions to write the solution in terms of the Airy functions. Fix $c_0, c_1, c_2, \hat{c}, y_0 \in \mathbb{R}$ such that $c_2 \neq 0$ and consider the differential equation

$$y'(x) = c_2 y^2(x) + c_1 y(x) + \hat{c}x + c_0 \quad (\text{C.1})$$

for $x \in (0, \infty)$ subject to the boundary condition $y(0) = y_0$. This equation falls in the class of Riccati equations; see e.g., Zaitsev and Polyanin [221, Section 1.2.2.24]. The next lemma provides a closed-form expression for the general solution to Equation (C.1) in terms of the Bessel functions of order $1/3$, i.e., $J_{1/3}$ and $Y_{1/3}$. As a preliminary to this lemma, let C_1 and C_2 be fixed constants and define

$$z(x) = x + \frac{4c_0c_2 - c_1^2}{4\hat{c}c_2}$$

$$u(x) = C_1 \exp\left(\frac{c_1 x}{2}\right) \sqrt{z(x)} J_{1/3}\left(\frac{2}{3}\sqrt{\hat{c}c_2} z^{3/2}(x)\right) + C_2 \exp\left(\frac{c_1 x}{2}\right) \sqrt{z(x)} Y_{1/3}\left(\frac{2}{3}\sqrt{\hat{c}c_2} z^{3/2}(x)\right)$$

for $x \in [0, \infty)$, where $J_{1/3}$ and $Y_{1/3}$ are the order $1/3$ Bessel functions of the first and second kind.

Lemma 19. *Zaitsev and Polyanin [221, Sections 1.2.1.2 and 2.1.2.12]. The (general) solution to the differential equation (C.1) is*

$$y(x) = -\frac{1}{c_2} \frac{u'(x)}{u(x)}, \quad x \in [0, \infty).$$

Proof. First, we transform the Riccati equation into a second-order linear differential equation. Then, we write the solution to the second-order equation in terms of the Bessel functions of order 1/3. Following Zaitsev and Polyanin [221, Section 1.2.1.2], we substitute

$$u(x) = \exp \left(-c_2 \int_0^x y(w) dw \right), \quad x \in [0, \infty) \quad (\text{C.2})$$

into Equation (C.1) to obtain

$$u''(x) - c_1 u'(x) + c_2(\hat{c}x + c_0)u(x) = 0, \quad x \in (0, \infty). \quad (\text{C.3})$$

By Zaitsev and Polyanin [221, Section 2.1.2.12], the (general) solution to the differential equation (C.3) is of the form

$$u(x) = C_1 \exp\left(\frac{c_1 x}{2}\right) \sqrt{z(x)} J_{1/3}\left(\frac{2}{3}\sqrt{\hat{c}c_2} z^{3/2}(x)\right) + C_2 \exp\left(\frac{c_1 x}{2}\right) \sqrt{z(x)} Y_{1/3}\left(\frac{2}{3}\sqrt{\hat{c}c_2} z^{3/2}(x)\right)$$

for $x \in [0, \infty)$, where C_1 and C_2 are constants, $J_{1/3}$ and $Y_{1/3}$ are the order 1/3 Bessel functions of the first and second kind, respectively, and

$$z(x) = x + \frac{4c_0c_2 - c_1^2}{4\hat{c}c_2}, \quad x \in [0, \infty).$$

Note that Equation (C.2) can be equivalently written as

$$y(x) = -\frac{1}{c_2} \frac{u'(x)}{u(x)}, \quad x \in [0, \infty).$$

This completes the proof. \square

Next, we use the relationship between the Airy functions and the Bessel functions to write the solution to the differential equation (C.1) subject to the boundary condition $y(0) = y_0$ in terms of the Airy functions. The Airy function of the first kind $\text{Ai} : \mathbb{R} \rightarrow \mathbb{R}$ and the second

kind $\text{Bi} : \mathbb{R} \rightarrow \mathbb{R}$ are defined as follows:

$$\begin{aligned}\text{Ai}(x) &= \frac{1}{\pi} \int_0^\infty \cos\left(\frac{1}{3}z^3 + zx\right) dz, \quad x \in \mathbb{R} \\ \text{Bi}(x) &= \frac{1}{\pi} \int_0^\infty \left(\cos\left(-\frac{1}{3}z^3 + zx\right) + \sin\left(\frac{1}{3}z^3 + zx\right) \right) dz, \quad x \in \mathbb{R}.\end{aligned}$$

The Airy functions and the Bessel functions of order $1/3$ can be expressed in terms of each other; see e.g., Abramowitz and Stegun [4, Section 10.4]. The following lemma uses this relationship to provide a closed-form solution for the differential equation (C.1) in terms of the Airy functions. This result is crucially used in solving the Bellman equation in Section 3.6.3. In particular, Lemma 11 solves Equations (3.49)-(3.50) and (3.51)-(3.52) by showing that they belong to the class of Riccati equations studied here. As a preliminary to the next lemma, define the function

$$\tilde{z}(x) = \frac{c_1^2 - 4c_2(\hat{c}x + c_0)}{4(\hat{c}c_2)^{2/3}}, \quad x \in [0, \infty)$$

define the constants

$$\begin{aligned}\tilde{C}_1 &= \frac{\text{Bi}'(\tilde{z}(0)) - \text{Bi}(\tilde{z}(0))(\hat{c}c_2)^{-1/3}(c_2y_0 + \frac{c_1}{2})}{\text{Ai}(\tilde{z}(0))\text{Bi}'(\tilde{z}(0)) - \text{Ai}'(\tilde{z}(0))\text{Bi}(\tilde{z}(0))} \\ \tilde{C}_2 &= \frac{-\text{Ai}'(\tilde{z}(0)) + \text{Ai}(\tilde{z}(0))(\hat{c}c_2)^{-1/3}(c_2y_0 + \frac{c_1}{2})}{\text{Ai}(\tilde{z}(0))\text{Bi}'(\tilde{z}(0)) - \text{Ai}'(\tilde{z}(0))\text{Bi}(\tilde{z}(0))},\end{aligned}$$

and the function

$$\tilde{u}(x) = \tilde{C}_1 \exp\left(\frac{c_1 x}{2}\right) \text{Ai}(\tilde{z}(x)) + \tilde{C}_2 \exp\left(\frac{c_1 x}{2}\right) \text{Bi}(\tilde{z}(x)), \quad x \in [0, \infty). \quad (\text{C.4})$$

Lemma 20. *The (unique) solution to the differential equation (C.1) subject to the boundary*

condition $y(0) = y_0$ is

$$y(x) = -\frac{1}{c_2} \frac{\tilde{u}'(x)}{\tilde{u}(x)}, \quad x \in [0, \infty).$$

Proof. By Lemma 19, the (general) solution to the differential equation (C.1) is of the form

$$y(x) = -\frac{1}{c_2} \frac{u'(x)}{u(x)}, \quad x \in [0, \infty) \quad (\text{C.5})$$

where

$$\begin{aligned} u(x) = & C_1 \exp\left(\frac{c_1 x}{2}\right) \sqrt{z(x)} J_{1/3}\left(\frac{2}{3} \sqrt{\hat{c} c_2} z^{3/2}(x)\right) \\ & + C_2 \exp\left(\frac{c_1 x}{2}\right) \sqrt{z(x)} Y_{1/3}\left(\frac{2}{3} \sqrt{\hat{c} c_2} z^{3/2}(x)\right), \quad x \in [0, \infty) \end{aligned} \quad (\text{C.6})$$

the functions $J_{1/3}$ and $Y_{1/3}$ are the order $1/3$ Bessel functions of the first and second kind, respectively, and

$$z(x) = x + \frac{4c_0 c_2 - c_1^2}{4\hat{c} c_2}, \quad x \in [0, \infty). \quad (\text{C.7})$$

First, we use the relationship between the Bessel functions of the first and second kind to write u solely in terms of the Bessel functions of the first kind. To do so, note that by Abramowitz and Stegun [4, Equation (9.1.2)],

$$Y_{1/3}(x) = \frac{J_{1/3}(x) \cos(\pi/3) - J_{-1/3}(x)}{\sin(\pi/3)}, \quad x \in \mathbb{R}. \quad (\text{C.8})$$

It follows from Equations (C.6) and (C.8) that u can be written as

$$\begin{aligned} u(x) = & \hat{C}_1 \exp\left(\frac{c_1 x}{2}\right) \sqrt{z(x)} J_{1/3}\left(\frac{2}{3} \sqrt{\hat{c} c_2} z^{3/2}(x)\right) \\ & + \hat{C}_2 \exp\left(\frac{c_1 x}{2}\right) \sqrt{z(x)} J_{-1/3}\left(\frac{2}{3} \sqrt{\hat{c} c_2} z^{3/2}(x)\right), \quad x \in [0, \infty) \end{aligned} \quad (\text{C.9})$$

for some $\hat{C}_1, \hat{C}_2 \in \mathbb{R}$. Next, we use the relationship between the Airy functions and the Bessel functions to write u in terms of the Airy functions. By Abramowitz and Stegun [4, Equation (10.4.22)],

$$\sqrt{\hat{z}} J_{\pm 1/3}(\frac{2}{3}\hat{z}^{3/2}) = \frac{3}{2} \text{Ai}(-\hat{z}) \mp \frac{\sqrt{3}}{2} \text{Bi}(-\hat{z}), \quad \hat{z} \in \mathbb{R}, \quad (\text{C.10})$$

where Ai and Bi are the Airy functions of the first and second kind, respectively. It follows from Equations (C.9)-(C.10) and the change of variable $\hat{z} = \sqrt[3]{\hat{c}c_2}z$ that $u = \tilde{u}$, where

$$\tilde{u}(x) = \tilde{C}_1 \exp\left(\frac{c_1 x}{2}\right) \text{Ai}(\tilde{z}(x)) + \tilde{C}_2 \exp\left(\frac{c_1 x}{2}\right) \text{Bi}(\tilde{z}(x)), \quad (\text{C.11})$$

for $x \in [0, \infty)$, \tilde{C}_1 and \tilde{C}_2 are constants and for $x \in [0, \infty)$,

$$\begin{aligned} \tilde{z}(x) &= -\sqrt[3]{\hat{c}c_2}z(x) \\ &= -\sqrt[3]{\hat{c}c_2}\left(x + \frac{4c_0c_2 - c_1^2}{4\hat{c}c_2}\right) \\ &= -\frac{4(\hat{c}c_2)^{2/3}}{4(\hat{c}c_2)^{2/3}} \sqrt[3]{\hat{c}c_2} \left(x + \frac{4c_0c_2 - c_1^2}{4\hat{c}c_2}\right) \\ &= \frac{c_1^2 - 4c_2(\hat{c}x + c_0)}{4(\hat{c}c_2)^{2/3}}. \end{aligned} \quad (\text{C.12})$$

Moreover, it follows from $u = \tilde{u}$, Equation (C.5), and boundary condition $y(0) = y_0$ that

$$\tilde{u}(0) = 1 \quad \text{and} \quad \tilde{u}'(0) = -c_2 y_0. \quad (\text{C.13})$$

It follows from Equations (C.11) and (C.13) that the constants \tilde{C}_1 and \tilde{C}_2 must satisfy

$$\begin{aligned} \left. \left(\tilde{C}_1 \text{Ai}(\tilde{z}(x)) + \tilde{C}_2 \text{Bi}(\tilde{z}(x)) \right) \right|_{x=0} &= \left. \exp\left(-\frac{c_1 x}{2}\right) \tilde{u}(x) \right|_{x=0} \\ \left. \left(\tilde{C}_1 \text{Ai}(\tilde{z}(x)) + \tilde{C}_2 \text{Bi}(\tilde{z}(x)) \right)' \right|_{x=0} &= \left. \left(\exp\left(-\frac{c_1 x}{2}\right) \tilde{u}(x) \right)' \right|_{x=0}, \end{aligned}$$

which can be simplified to

$$\begin{aligned}\tilde{C}_1 \text{Ai}(\tilde{z}(0)) + \tilde{C}_2 \text{Bi}(\tilde{z}(0)) &= 1 \\ \tilde{C}_1 \text{Ai}'(\tilde{z}(0)) + \tilde{C}_2 \text{Bi}'(\tilde{z}(0)) &= (\hat{c}c_2)^{-1/3}(c_2 y_0 + \frac{c_1}{2}).\end{aligned}$$

Thus,

$$\begin{bmatrix} \tilde{C}_1 \\ \tilde{C}_2 \end{bmatrix} = \begin{bmatrix} \text{Ai}(\tilde{z}(0)) & \text{Bi}(\tilde{z}(0)) \\ \text{Ai}'(\tilde{z}(0)) & \text{Bi}'(\tilde{z}(0)) \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ (\hat{c}c_2)^{-1/3}(c_2 y_0 + \frac{c_1}{2}) \end{bmatrix},$$

which gives

$$\begin{aligned}\begin{bmatrix} \tilde{C}_1 \\ \tilde{C}_2 \end{bmatrix} &= \beta \begin{bmatrix} \text{Bi}'(\tilde{z}(0)) & -\text{Bi}(\tilde{z}(0)) \\ -\text{Ai}'(\tilde{z}(0)) & \text{Ai}(\tilde{z}(0)) \end{bmatrix} \begin{bmatrix} 1 \\ (\hat{c}c_2)^{-1/3}(c_2 y_0 + \frac{c_1}{2}) \end{bmatrix} \\ &= \beta \begin{bmatrix} \text{Bi}'(\tilde{z}(0)) - \text{Bi}(\tilde{z}(0))(\hat{c}c_2)^{-1/3}(c_2 y_0 + \frac{c_1}{2}) \\ -\text{Ai}'(\tilde{z}(0)) + \text{Ai}(\tilde{z}(0))(\hat{c}c_2)^{-1/3}(c_2 y_0 + \frac{c_1}{2}) \end{bmatrix}. \quad (\text{C.14})\end{aligned}$$

where

$$\beta = \frac{1}{\text{Ai}(\tilde{z}(0))\text{Bi}'(\tilde{z}(0)) - \text{Ai}'(\tilde{z}(0))\text{Bi}(\tilde{z}(0))}.$$

In conclusion,

$$y(x) = -\frac{1}{c_2} \frac{\tilde{u}'(x)}{\tilde{u}(x)}, \quad x \in [0, \infty),$$

where \tilde{u} is given by Equation (C.11), \tilde{z} is given by Equation (C.12), and \tilde{C}_1 and \tilde{C}_2 are given by Equation (C.14). This completes the proof. \square

C.2 Derivations

C.2.1 Formal Derivation of the Approximating Brownian System

We start by writing the scaled inventory and the cumulative cost processes in terms of other scaled processes. Then, we formally take the limit to obtain the approximating Brownian system.

Scaled Inventory Process. By Equation (3.1), for $k = 1, \dots, K$ and $t \geq 0$, we have

$$\begin{aligned} Q_k^n(t) &= S_k^n(T_k^n(t)) - N_k \left(\int_0^t \lambda_k^n(s) ds \right) + O_k^n(t) \\ &= S_k^n(T_k^n(t)) + \left[\mu_k^n T_k^n(t) - \mu_k^n T_k^n(t) \right] + \left[\mu_k^n \rho_k t - \mu_k^n \rho_k t \right] \\ &\quad + \left[\int_0^t \lambda_k^n(s) ds - \int_0^t \lambda_k^n(s) ds \right] - N_k^n \left(\int_0^t \lambda_k^n(s) ds \right) + O_k^n(t) \end{aligned} \quad (\text{C.15})$$

$$\begin{aligned} &= \left[S_k^n(T_k^n(t)) - \mu_k^n T_k^n(t) \right] - \mu_k^n \left[\rho_k t - T_k^n(t) \right] + \mu_k^n \rho_k t - \int_0^t \lambda_k^n(s) ds \\ &\quad - \left[N_k^n \left(\int_0^t \lambda_k^n(s) ds \right) - \int_0^t \lambda_k^n(s) ds \right] + O_k^n(t), \end{aligned} \quad (\text{C.16})$$

where (C.15) is obtained by adding and subtracting the terms insides the brackets and (C.16) is obtained by rearranging the terms. It follows from the functional strong approximations (Ethier and Kurtz [86, Section 7.5]) that for $t \geq 0$,

$$S_k^n(t) = \mu_k^n t + \sqrt{n \mu_k} \nu_{sk} \hat{\chi}_k^n(t) + o(\sqrt{n}) \quad (\text{C.17})$$

$$N_k^n(nt) = nt + \sqrt{n} \tilde{\chi}_k^n(t) + o(\sqrt{n}), \quad (\text{C.18})$$

almost surely, where $\hat{\chi}_k^n$ and $\tilde{\chi}_k^n$ are independent standard Brownian motions and the notation $f(n) = o(g(n))$ denotes that $f(n)/g(n) \rightarrow 0$ as $n \rightarrow \infty$. Plugging (C.17)-(C.18) into (C.16)

gives

$$\begin{aligned} Q_k^n(t) &= \sqrt{n\mu_k}\nu_{sk}\hat{\chi}_k^n(T_k^n(t)) - \mu_k^n[\rho_k t - T_k^n(t)] + \mu_k^n\rho_k t - \int_0^t \lambda_k^n(s)ds \\ &\quad - \sqrt{n}\tilde{\chi}_k^n\left(\frac{1}{n}\int_0^t \lambda_k^n(s)ds\right) + O_k^n(t) + o(\sqrt{n}), \quad t \geq 0. \end{aligned} \quad (\text{C.19})$$

Dividing all terms in (C.19) by \sqrt{n} and plugging in (3.9)-(3.13) gives

$$\begin{aligned} \frac{Q_k^n(t)}{\sqrt{n}} &= \sqrt{\mu_k}\nu_{sk}\hat{\chi}_k^n(T_k^n(t)) - \mu_k Y_k^n(t) + \sqrt{n}\rho_k\mu_k t + \rho_k\eta_k t - \sqrt{n}\lambda_k^\star t - \int_0^t \zeta_k(s)ds \\ &\quad - \tilde{\chi}_k^n\left(\lambda_k^\star t + \frac{1}{\sqrt{n}}\int_0^t \zeta_k(s)ds\right) + R_k^n(t) + o(1), \quad t \geq 0. \end{aligned} \quad (\text{C.20})$$

Then, plugging $\rho_k\mu_k = \lambda_k^\star$ into (C.20) gives

$$\begin{aligned} \frac{Q_k^n(t)}{\sqrt{n}} &= \sqrt{\mu_k}\nu_{sk}\hat{\chi}_k^n(T_k^n(t)) - \mu_k Y_k^n(t) + \rho_k\eta_k t - \int_0^t \zeta_k(s)ds \\ &\quad - \tilde{\chi}_k^n\left(\lambda_k^\star t + \frac{1}{\sqrt{n}}\int_0^t \zeta_k(s)ds\right) + R_k^n(t) + o(1), \quad t \geq 0. \end{aligned}$$

Next, using a straightforward application of the random time change theorem (on the fifth term), we obtain

$$\frac{Q_k^n(t)}{\sqrt{n}} = \sqrt{\mu_k}\nu_{sk}\hat{\chi}_k^n(T_k^n(t)) - \mu_k Y_k^n(t) + \rho_k\eta_k t - \int_0^t \zeta_k(s)ds - \tilde{\chi}_k^n(\lambda_k^\star t) + R_k^n(t) + o(1)$$

for $t \geq 0$. Finally, letting

$$\chi_k^n(t) = \sqrt{\mu_k}\nu_{sk}\hat{\chi}_k^n(T_k^n(t)) + \rho_k\eta_k t - \tilde{\chi}_k^n(\lambda_k^\star t),$$

for $k = 1, \dots, K$ and $t \geq 0$, we arrive at

$$Z_k^n(t) = \chi_k^n(t) - \mu_k Y_k^n(t) - \int_0^t \zeta_k(s)ds + R_k^n(t) + o(1)$$

for $k = 1, \dots, K$ and $t \geq 0$.

Cost Process. By Equation (3.6), we have

$$V^n(t) = \int_0^t \pi^n(\lambda^n(s)) ds - \sum_{k=1}^K \int_0^t q_k^n(Q_k^n(s)) ds - \sum_{k=1}^K \nu_k^n O_k^n(t), \quad t \geq 0. \quad (\text{C.21})$$

Let us start by focusing on the first term:

$$\pi^n(\lambda^n(s)) = \pi^n(n\lambda^\star + \sqrt{n}\zeta(s)) = n\pi(\lambda^\star + \frac{\zeta(s)}{\sqrt{n}}), \quad (\text{C.22})$$

where (C.22) follows from (3.10). By Taylor's expansion, we have

$$\begin{aligned} \pi(\lambda^\star + \frac{\zeta(s)}{\sqrt{n}}) &= \pi(\lambda^\star) + \frac{1}{\sqrt{n}} \nabla \pi(\lambda^\star) \zeta(s) + \frac{1}{2n} \zeta(s)' \nabla^2 \pi(\lambda^\star) \zeta(s) + o(\frac{1}{n}) \\ &= \pi(\lambda^\star) + \frac{1}{2n} \zeta(s)' \nabla^2 \pi(\lambda^\star) \zeta(s) + o(\frac{1}{n}), \end{aligned} \quad (\text{C.23})$$

where (C.23) follows the fact that $\nabla \pi(\lambda^\star) = 0$ by Assumption 7. Therefore,

$$\int_0^t \pi^n(\lambda^n(s)) ds = \pi(\lambda^\star) nt + \frac{1}{2} \int_0^t \zeta(s)' \nabla^2 \pi(\lambda^\star) \zeta(s) ds + o(1), \quad t \geq 0. \quad (\text{C.24})$$

Then, by (3.16), (C.21), and (C.24), we have for $t \geq 0$,

$$\begin{aligned} \xi^n(t) &= -\frac{1}{2} \int_0^t \zeta(s)' \nabla^2 \pi(\lambda^\star) \zeta(s) ds + \sum_{k=1}^K \int_0^t q_k^n(Q_k^n(s)) ds + \sum_{k=1}^K \nu_k^n O_k^n(t) + o(1) \\ &= -\frac{1}{2} \int_0^t \zeta(s)' \nabla^2 \pi(\lambda^\star) \zeta(s) ds + \sum_{k=1}^K \int_0^t g_k(Z_k^n(s)) ds + \sum_{k=1}^K r_k R_k^n(t) + o(1), \end{aligned} \quad (\text{C.25})$$

where (C.25) follows from (3.12) and (3.14)-(3.15).

Approximating Brownian System. So far, we have established that for $k = 1, \dots, K$ and $t \geq 0$,

$$Z_k^n(t) = \chi_k^n(t) - \mu_k Y_k^n(t) - \int_0^t \zeta_k(s) ds + R_k^n(t) + o(1), \quad (\text{C.26})$$

where

$$\chi_k^n(t) = \sqrt{\mu_k} \nu_{sk} \hat{\chi}_k^n(T_k^n(t)) + \rho_k \eta_k t - \tilde{\chi}_k^n(\lambda_k^\star t), \quad t \geq 0.$$

Moreover, for $k = 1, \dots, K$ and $t \geq 0$,

$$\xi^n(t) = -\frac{1}{2} \int_0^t \zeta(s)' \nabla^2 \pi(\lambda^\star) \zeta(s) ds + \sum_{k=1}^K \int_0^t g_k(Z_k^n(s)) ds + \sum_{k=1}^K r_k R_k^n(t) + o(1). \quad (\text{C.27})$$

The key step in the development of our Brownian approximation is the following claim: If n is large, the only interesting allocation policies are those for which

$$T_k^n(t) \approx \rho_k t \quad (\text{C.28})$$

for $k = 1, \dots, K$ and $t \geq 0$. For a defense of this claim, let us first restrict attention to policies that do not idle the manufacturing system as long as there are backordered requests. If n is large and $t > 0$ is moderate, it can be shown to follow from Assumption 7 and Equations (3.9) and (3.11) that $I^n(t) \approx 0$. Moreover, the relative amount of time the manufacturing system manufactures different products over the interval $[0, t]$ must be approximately equal to ρ_k . Thus, we conclude that the difference between $T_k^n(t)$ and $\rho_k t$ must be small, i.e., (C.28). The same conclusion holds for policies that allow idleness while there are backordered requests provided $I^n(t) \approx 0$. Under plausible assumptions, it can be shown that no greater amount of idleness could ever be desirable. The heart of the matter is that backorders of order n result inevitably if $I^n(t) \not\approx 0$ but can be otherwise avoided; see Harrison [110, Sections 5

and 11] for a similar informal argument. Next, using a straightforward application of the random time change theorem and Equation (C.28), we obtain

$$\chi_k^n \Rightarrow X_k \quad (\text{C.29})$$

for $k = 1, \dots, K$, where X_1, \dots, X_K are independent Brownian motions with infinitesimal drift $\rho_k \eta_k$, infinitesimal variance $\sigma_k^2 = \lambda_k^*(1 + \nu_{sk}^2)$, and initial value zero.

Motivated by (C.26), (C.27), and (C.29), we argue similar to Harrison [110] that as the system gets large, the scaled processes defined above converge weakly to ξ , Z , Y , U , and R that satisfy the following for $t \geq 0$,

$$\begin{aligned} \xi(t) &= \int_0^t \zeta(s)' H \zeta(s) ds + \sum_{k=1}^K \int_0^t g_k(Z_k(s)) ds + \sum_{k=1}^K r_k R_k(t), \\ Z_k(t) &= X_k(t) - \mu_k Y_k(t) - \int_0^t \zeta_k(s) ds + R_k(t), \quad k = 1, \dots, K, \\ U(t) &= \sum_{k=1}^K Y_k(t), \end{aligned}$$

U, R are nondecreasing with $U(0) = R(0) = 0$,

where $H = -\nabla^2 \pi(\lambda^*)/2$ and $X_k = \{X_k(t), t \geq 0\}$ are independent $(\eta_k \rho_k, \sigma_k^2)$ Brownian motions with initial value zero.

C.2.2 Derivation of Equation (3.23)

Recall that the instantaneous demand rate in the n -th system is of the form

$$\lambda^n(t) = n\lambda^* + \sqrt{n}\zeta(t), \quad t \geq 0.$$

It follows from the definition of the inverse demand function that for $t \geq 0$,

$$\begin{aligned}
p^n(t) &= (\Lambda^n)^{-1}(\lambda^n(t)) \\
&= (\Lambda^n)^{-1}(n\lambda^* + \sqrt{n}\zeta(t)) \\
&= \Lambda^{-1}(\lambda^* + \frac{\zeta(t)}{\sqrt{n}}),
\end{aligned} \tag{C.30}$$

where (C.30) follows from (3.10). Then, by Taylor's expansion, we have

$$p^n(t) = \Lambda^{-1}(\lambda^*) + \frac{\nabla \Lambda^{-1}(\lambda^*) \zeta(t)}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right),$$

for $t \geq 0$, where $\nabla \Lambda^{-1}(\lambda^*)$ denotes the Jacobian of Λ^{-1} evaluated at λ^* .

C.3 Miscellaneous Proofs

Proof of Proposition 5. The proof is almost identical to the proof of Rubino and Ata [181, Proposition 3] and closely follows the proof of Harrison and Van Mieghem [120, Propositions 3-4 and Theorem 2] using the approach outlined in the appendix of Harrison and Van Mieghem [120]. The only minor modification needed is the generalization of the arguments in Harrison and Van Mieghem [120] to allow a state-dependent drift term and an outsourcing term. However, the analysis of Harrison and Van Mieghem [120] is not sensitive to the drift term or the inclusion of an outsourcing process, and therefore, a state-dependent drift term and an outsourcing term can be accommodated easily. \square

Proof of Proposition 6. Consider the barrier policy (L, U, θ) and assume $\gamma \in \mathbb{R}$ and

$f \in C^2([l, u])$ jointly satisfy (3.34)-(3.35). A routine application of Ito's lemma gives

$$\begin{aligned}\mathbb{E}[f(W(t))] - f(W(0)) &= \mathbb{E}\left[\int_0^t \Gamma_\theta f(W(s))ds\right] + \mathbb{E}\left[\int_0^t f'(W(s))dL(s)\right] \\ &\quad - \mathbb{E}\left[\int_0^t f'(W(s))dU(s)\right] \\ &= \mathbb{E}\left[\int_0^t \Gamma_\theta f(W(s))ds\right] + \mathbb{E}\left[\int_0^t -\kappa dL(s)\right] \tag{C.31}\end{aligned}$$

$$= \gamma t - \mathbb{E}\left[\int_0^t c(\theta(W(s)))ds + \int_0^t h(W(s))ds\right] - \mathbb{E}[\kappa L(t)], \tag{C.32}$$

where (C.31) follows from (3.33) and (3.35), and (C.32) follows from (3.34). Divide both sides of (C.32) by t and take the limit as $t \rightarrow \infty$. The first term on the left hand side vanishes, because f' is bounded and $W(t) \in [l, u]$ for $t \geq 0$. The second term on the left hand side vanishes since $W(0) = 0$. Consequently, we arrive at

$$\lim_{t \rightarrow \infty} \frac{1}{t} \mathbb{E}\left[\int_0^t c(\theta(W(s)))ds + \int_0^t h(W(s))ds + \kappa L(t)\right] = \gamma.$$

This completes the proof. \square

Proof of Lemma 9. It is straightforward to write l and u in terms of γ using Equations (3.45)-(3.46). Here, we prove that $l < 0$ and $u > 0$, from which $\gamma > 0$ follows.

Lower barrier is negative. Recall that κ is non-negative by definition, μ is non-negative since $\eta \geq 0$, and H is positive definite by Assumption 6. Plugging these into Equation (3.47) gives $h(u) \leq h(l)$. Next, we prove $l < 0$ by contradiction. Assume $l \geq 0$. Since h is strictly increasing on $[0, \infty)$, $u \geq l$ by definition, and $h(u) \leq h(l)$, we must have $l = u$. This, however, contradicts $v(l) = -\kappa$ and $v(u) = 0$. Therefore, the assumption that $l \geq 0$ is incorrect.

Upper barrier is positive. We prove $u > 0$ by contradiction. Assume $u \leq 0$. Since $v \in C^1([l, u])$, the quadratic function is smooth, and h is continuously differentiable on (l, u) , it follows from (3.42) that $v \in C^2([l, u])$; i.e., v'' is well-defined on $[l, 0]$. Therefore, we

can differentiate both sides of (3.42) to get

$$v''(w) = \frac{m'H^{-1}m}{\sigma^2} v(w) v'(w) - \frac{2\mu}{\sigma^2} v'(w) - \frac{2h'(w)}{\sigma^2}, \quad w \in (l, u). \quad (\text{C.33})$$

Since v and v' are continuous on $[l, u]$, $v(u) = v'(u) = 0$, and $h'(w) = -b^* < 0$ for $w \in [l, u]$, it follows from (C.33) that there exists $w_0 \in (l, u)$ such that $v'' > 0$ on (w_0, u) . On the other hand, since $l, u \leq 0$, by Lemma 22, $v' > 0$ on (l, u) . Putting these two together, we arrive at the following contradiction: $v', v'' > 0$ on (w_0, u) imply $v'(u) > 0$ while we know $v'(u) = 0$ by (3.44). Therefore, the assumption that $u \leq 0$ is incorrect. This completes the proof. \square

Proof of Lemma 10. Recall that $\gamma \geq 0$.

Part 1A: $v_0^-(0) < 0$. We prove $v_0^-(0) < 0$ by contradiction. Assume that $v_0^-(0) \geq 0$. It follows from the continuity of v_0^- , the fact that $v_0^-(l_\gamma) = -\kappa < 0$, and the assumption that $v_0^-(0) \geq 0$ that there exists $w_0 \in [l_\gamma, 0]$ such that $v_0^-(w_0) = 0$. By Equation (3.49) and the continuity of v' , for $w \in (l_\gamma, 0]$, we have

$$\begin{aligned} \frac{\partial}{\partial w} v_0^-(w_0) &= \frac{m'H^{-1}m}{2\sigma^2} ((v_0^-)^2(w_0) - \kappa^2) - \frac{2\mu}{\sigma^2} (v_0^-(w_0) + \kappa) + \frac{2(h(l_\gamma) - h(w_0))}{\sigma^2} \\ &= -\frac{m'H^{-1}m}{2\sigma^2} \kappa^2 - \frac{2\mu}{\sigma^2} \kappa + \frac{2(h(l_\gamma) - h(w_0))}{\sigma^2} \end{aligned} \quad (\text{C.34})$$

$$\leq -\frac{m'H^{-1}m}{2\sigma^2} \kappa^2 - \frac{2\mu}{\sigma^2} \kappa + \frac{2h(l_\gamma)}{\sigma^2} \quad (\text{C.35})$$

$$= 0, \quad (\text{C.36})$$

where (C.34) follows $v_0^-(w_0) = 0$, (C.35) follows from the non-negativity of h , and (C.36) follows from (3.45). However, by Lemma 22, we must have $\partial v_0^-(w)/\partial w > 0$ for $w \in (l_\gamma, 0]$, which contradicts (C.36). Therefore, the initial assumption that $v_0^-(0) \geq 0$ is incorrect.

Part 1B: $v_\gamma^-(0)$ is strictly increasing in γ and $v_\gamma^-(0) \rightarrow \infty$ as $\gamma \rightarrow \infty$. Consider the differential equation

$$\hat{v}'(w) = \frac{m'H^{-1}m}{2\sigma^2} (\hat{v}^2(w) - \kappa^2) - \frac{2\mu}{\sigma^2} (\hat{v}(w) + \kappa) + \frac{2b^*}{\sigma^2} w, \quad w \in (0, \infty) \quad (\text{C.37})$$

subject to the boundary condition $\hat{v}(0) = -\kappa$. It is straightforward to see that

$$v_\gamma^-(w) = \hat{v}(w - l_\gamma), \quad w \in [l_\gamma, 0]. \quad (\text{C.38})$$

Then, it follows from Lemma 22 that $\hat{v}'(w) = \partial v_\gamma^-(w + l_\gamma)/\partial w > 0$ for $w \in (0, \infty)$. Thus, for $\gamma > 0$,

$$\frac{\partial}{\partial \gamma} v_\gamma^-(0) = \frac{\partial}{\partial \gamma} \hat{v}(-l_\gamma) = \frac{1}{b^*} \hat{v}'(-l_\gamma) > 0, \quad (\text{C.39})$$

where the second equality in (C.39) follows from (3.48). Thus, $v_\gamma^-(0)$ is strictly increasing. It remains to show that $v_\gamma^-(0) \rightarrow \infty$ as $\gamma \rightarrow \infty$. By (3.48) and (C.38), it suffices to show that $\hat{v}(w) \rightarrow \infty$ as $w \rightarrow \infty$. First, we show that $\hat{v}'(w) \rightarrow \infty$ as $w \rightarrow \infty$. To see this, note that

$$\frac{m' H^{-1} m}{2\sigma^2} (\hat{v}^2(w) - \kappa^2) - \frac{2\mu}{\sigma^2} (\hat{v}(w) + \kappa)$$

is lower bounded since $\frac{m' H^{-1} m}{2\sigma^2} > 0$. Then, since $\frac{2b^*}{\sigma^2} w \rightarrow \infty$ as $w \rightarrow \infty$, it follows from (C.37) that $\hat{v}(w) \rightarrow \infty$ as $w \rightarrow \infty$. Then, by the fundamental theorem of calculus,

$$\hat{v}(w) = \int_0^w \hat{v}'(w') dw' + \hat{v}(0) \rightarrow \infty. \quad (\text{C.40})$$

Part 2: $v_0^+(0) = 0$, $v_\gamma^+(0)$ is strictly decreasing in γ and $v_\gamma^+(0) \rightarrow -\infty$ as $\gamma \rightarrow \infty$.

The proof of Part 2 resembles the proof of Part 1. The equality $v_0^+(0) = 0$ follows from the boundary condition (3.52). Consider the differential equation

$$\tilde{v}'(w) = -\frac{m' H^{-1} m}{2\sigma^2} \tilde{v}^2(w) + \frac{2\mu}{\sigma^2} \tilde{v}(w) - \frac{2h^*}{\sigma^2} w, \quad w \in (0, \infty) \quad (\text{C.41})$$

subject to the boundary condition $\tilde{v}(0) = 0$. It is straightforward to see that

$$v_\gamma^+(w) = \tilde{v}(u_\gamma - w), \quad w \in [0, u_\gamma]. \quad (\text{C.42})$$

By Lemma 23, $\tilde{v}, \tilde{v}' < 0$ on $(0, \infty)$. Therefore, for $\gamma > 0$,

$$\frac{\partial}{\partial \gamma} v_\gamma^+(0) = \frac{\partial}{\partial \gamma} \hat{v}(u_\gamma) = \frac{1}{h^\star} \tilde{v}'(u_\gamma) < 0.$$

Thus, $v_\gamma^+(0)$ is strictly decreasing in γ . It remains to show that $v_\gamma^+(0) \rightarrow -\infty$ as $\gamma \rightarrow \infty$. By (C.42), it suffices to show that $\tilde{v}(w) \rightarrow -\infty$ as $w \rightarrow \infty$. Note that

$$-\frac{m' H^{-1} m}{2\sigma^2} \tilde{v}^2(w) + \frac{2\mu}{\sigma^2} \tilde{v}(w)$$

is upper bounded and $-\frac{2h^\star}{\sigma^2} w \rightarrow -\infty$ as $w \rightarrow \infty$. Thus, by (C.41), $\tilde{v}(w)' \rightarrow -\infty$ as $w \rightarrow \infty$.

Therefore,

$$\tilde{v}(w) = \int_0^w \tilde{v}'(w') dw' + \tilde{v}(0) \rightarrow -\infty. \quad (\text{C.43})$$

This completes the proof. \square

Proof of Lemma 11. We prove the statements for v_γ^- and v_γ^+ , separately.

Part 1. We start by focusing on v_γ^- . Consider the differential equation

$$\hat{v}'(w) = \frac{m' H^{-1} m}{2\sigma^2} (\hat{v}^2(w) - \kappa^2) - \frac{2\mu}{\sigma^2} (\hat{v}'(w) + \kappa) + \frac{2b^\star}{\sigma^2} w, \quad w \in (0, \infty) \quad (\text{C.44})$$

subject to the boundary condition $\hat{v}(0) = -\kappa$. It is straightforward to see that

$$v_\gamma^-(w) = \hat{v}(w - l_\gamma), \quad w \in [l_\gamma, 0]. \quad (\text{C.45})$$

Equation (C.44) is a Riccati equation of the form (C.1) with

$$c_0 = -\frac{m' H^{-1} m}{2\sigma^2} \kappa^2 - \frac{2\mu}{\sigma^2} \kappa, \quad c_1 = -\frac{2\mu}{\sigma^2}, \quad c_2 = \frac{m' H^{-1} m}{2\sigma^2}, \quad \text{and} \quad \hat{c} = \frac{2b^\star}{\sigma^2}.$$

The closed-form expression for v_γ^- follows from Lemma 20 and the fact that by (3.45), $c_0 =$

$$\gamma + b^* l_\gamma.$$

Part 2. Next, we focus on v_γ^+ . Consider the differential equation

$$\tilde{v}'(w) = -\frac{m'H^{-1}m}{2\sigma^2} \tilde{v}^2(w) + \frac{2\mu}{\sigma^2} \tilde{v}'(w) - \frac{2h^*}{\sigma^2} w, \quad w \in (0, \infty) \quad (\text{C.46})$$

subject to the boundary condition $\tilde{v}(0) = 0$. It is straightforward to see that

$$v_\gamma^+(w) = \tilde{v}(u_\gamma - w), \quad w \in [0, u_\gamma]. \quad (\text{C.47})$$

Equation (C.46) is a Riccati equation of the form (C.1) with

$$c_0 = 0, \quad c_1 = \frac{2\mu}{\sigma^2}, \quad c_2 = -\frac{m'H^{-1}m}{2\sigma^2}, \quad \text{and} \quad \hat{c} = -\frac{2h^*}{\sigma^2}.$$

The closed-form expression for v_γ^+ follows from Lemma 20. \square

Proof of Proposition 7. We prove the result in several steps.

Step 1: v satisfies (3.42) on $(l_{\gamma^*}, 0) \cup (0, u_{\gamma^*})$, $v(l_{\gamma^*}) = -\kappa$, and $v(u_{\gamma^*}) = 0$. It follows from the definition of l_{γ^*} and u_{γ^*} , the fact that $v_{\gamma^*}^-$ solves (3.49)-(3.50), and the definition of v in (3.53) that v satisfies (3.42) on $(l_{\gamma^*}, 0)$ and $v(l_{\gamma^*}) = -\kappa$. Similarly, it follows from the definition of v in (3.53) and the fact that $v_{\gamma^*}^+$ satisfies (3.51)-(3.52) that v satisfies (3.42) on $(0, u_{\gamma^*})$ and $v(u_{\gamma^*}) = 0$. To prove that $(l_{\gamma^*}, u_{\gamma^*}, \gamma^*, v)$ is a solution to (3.42)-(3.44) it remains to show that v and v' are continuous at the origin and $v'(l_{\gamma^*}) = v'(u_{\gamma^*}) = 0$.

Step 2: v and v' are continuous at the origin. The continuity of v at the origin follows from Corollary 3 and the definition of v in (3.53). To show the continuity of v' , we calculate the left and right limits of $v'(w)$ at the origin and show that the two limits coincide. By

(3.49)-(3.50) and (3.53),

$$\lim_{w \rightarrow 0^-} v'(w) = \lim_{w \rightarrow 0^-} \frac{\partial}{\partial w} v_{\gamma^*}^-(w) = \frac{m' H^{-1} m}{2\sigma^2} ((v_{\gamma^*}^-)^2(0) - \kappa^2) - \frac{2\mu}{\sigma^2} (v_{\gamma^*}^-(0) + \kappa) + \frac{2h(l_{\gamma^*})}{\sigma^2}. \quad (\text{C.48})$$

Similarly, by (3.51)-(3.52) and (3.53),

$$\lim_{w \rightarrow 0^+} v'(w) = \lim_{w \rightarrow 0^+} \frac{\partial}{\partial w} v_{\gamma^*}^+(w) = \frac{m' H^{-1} m}{2\sigma^2} (v_{\gamma^*}^+)^2(w) - \frac{2\mu}{\sigma^2} v_{\gamma^*}^+(0) + \frac{2h(u_{\gamma^*})}{\sigma^2}. \quad (\text{C.49})$$

Then, it follows from Corollary 3, Equations (C.48)-(C.49) and the definition of l_{γ^*} and u_{γ^*} that

$$\lim_{w \rightarrow 0^-} v'(w) = \lim_{w \rightarrow 0^+} v'(w). \quad (\text{C.50})$$

It follows from (C.50), the definition of v in (3.53), and the fact that $v_{\gamma^*}^- \in C^1([l_{\gamma^*}, 0])$ and $v_{\gamma^*}^+ \in C^1([0, u_{\gamma^*}])$ that $v \in C^1([l_{\gamma^*}, u_{\gamma^*}])$.

Step 3: $v'(l_{\gamma^*}) = v'(u_{\gamma^*}) = 0$. It follows from (3.49)-(3.50) that $\partial v_{\gamma^*}^-(l_{\gamma^*})/\partial w = 0$. Similarly, it follows from (3.51)-(3.52) that $\partial v_{\gamma^*}^+(l_{\gamma^*})/\partial w = 0$. Plugging $\partial v_{\gamma^*}^-(l_{\gamma^*})/\partial w = \partial v_{\gamma^*}^+(l_{\gamma^*})/\partial w = 0$ into the definition of v in (3.53) gives $v'(l_{\gamma^*}) = v'(u_{\gamma^*}) = 0$.

Step 4: v is strictly increasing on $[l_{\gamma^*}, u_{\gamma^*}]$. It follows from the definition of v in (3.53) and Lemma 22 that $v'(w) > 0$ for $w \in (l_{\gamma^*}, 0]$. Moreover, note that $v_{\gamma^*}^+(w) = \tilde{v}(u_{\gamma^*} - w)$ for $w \in [0, u_{\gamma^*}]$, where \tilde{v} solves

$$\tilde{v}'(w) = -\frac{m' H^{-1} m}{2\sigma^2} \tilde{v}^2(w) + \frac{2\mu}{\sigma^2} \tilde{v}(w) - \frac{2h^*}{\sigma^2} w, \quad w \in (0, \infty)$$

subject to the boundary condition $\tilde{v}(0) = 0$. Then, it follows from Lemma 23 that $v'(w) > 0$ for $w \in [0, u_{\gamma^*}]$. By combining $v'(w) > 0$ for $w \in (l_{\gamma^*}, 0]$ and $v'(w) > 0$ for $w \in [0, u_{\gamma^*}]$, we arrive at $v'(w) > 0$ for $w \in (l_{\gamma^*}, u_{\gamma^*})$.

Step 5: The solution to (3.42)-(3.44) is unique. The uniqueness of the solution follows

from Corollary 3 and the fact that v is uniquely determined given γ^* . This completes the proof. \square

Proof of Theorem 2. Consider v given by (3.53) and define $f : \mathbb{R} \rightarrow \mathbb{R}$ as follows:

$$f(w) = \begin{cases} \kappa(l_{\gamma^*} - w), & w \in (-\infty, l_{\gamma^*}), \\ \int_{l_{\gamma^*}}^w v(x)dx, & w \in [l_{\gamma^*}, u_{\gamma^*}], \\ \int_{l_{\gamma^*}}^{u_{\gamma^*}} v(x)dx, & w \in (u_{\gamma^*}, \infty). \end{cases} \quad (\text{C.51})$$

Note that (C.51) extends the domain of f in (3.54) from $[l_{\gamma^*}, u_{\gamma^*}]$ to the entire real line. In Step 1, we show that $-\kappa \leq f'(w) \leq 0$ for $w \in \mathbb{R}$ and

$$\min_{x \in \mathbb{R}} \left\{ \frac{1}{2}\sigma^2 f''(w) + \mu f'(w) - xf'(w) + c(x) + h(w) \right\} \geq \gamma^*, \quad w \in \mathbb{R}. \quad (\text{C.52})$$

Then, in Step 2, we show that the average cost of any arbitrarily chosen policy (L, U, θ) is greater than or equal to the average cost of our candidate policy.

Step 1. Since v is strictly increasing on $(l_{\gamma^*}, u_{\gamma^*})$ (by Lemma 7), $v(l_{\gamma^*}) = -\kappa$, and $v(u_{\gamma^*}) = 0$, we have $-\kappa \leq v(w) \leq 0$ for $w \in [l_{\gamma^*}, u_{\gamma^*}]$, which gives $-\kappa \leq f'(w) \leq 0$ for $w \in \mathbb{R}$. Moreover, by Corollary 4, $(l_{\gamma^*}, u_{\gamma^*}, \gamma^*, f)$ solves (3.36)-(3.37). Equation (C.52), then, follows from (3.36) and the fact that h is strictly decreasing on $(-\infty, l_{\gamma^*})$ and strictly increasing on (u_{γ^*}, ∞) .

Step 2. Let (L, U, θ) be an arbitrary admissible policy and denote its state process with W . A routine application of Ito's lemma gives

$$\begin{aligned} \mathbb{E}[f(W(t))] - f(W(0)) &= \mathbb{E} \left[\int_0^t \frac{1}{2}\sigma^2 f''(W(s))ds + \int_0^t \mu f'(W(s))ds - \int_0^t \theta(s)f'(W(s))ds \right] \\ &\quad + \mathbb{E} \left[\int_0^t f'(W(s))dL(s) \right] - \mathbb{E} \left[\int_0^t f'(W(s))dU(s) \right] \end{aligned} \quad (\text{C.53})$$

for ≥ 0 . The second term on the right hand side of (C.53) can be written out as

$$\begin{aligned}\mathbb{E}\left[\int_0^t f'(W(s))dL(s)\right] &= \mathbb{E}\left[\int_0^t f'(W(s))dL^c(s)\right. \\ &\quad \left. + \sum_{s \leq t} \left(f(W(s^-) + \Delta L(s)) - f(W(s^-))\right)\right],\end{aligned}\quad (\text{C.54})$$

where L^c denotes the continuous part of L and ΔL denotes its jumps. To be specific, for $t \geq 0$,

$$\Delta L(t) = L(t) - L(t^-) \quad \text{and} \quad L^c(t) = L(t) - \sum_{\substack{L(s) \neq L(s^-) \\ 0 \leq s \leq t}} \Delta L(s).$$

Plugging $f'(w) \geq -\kappa$ for $w \in \mathbb{R}$ into (C.54) gives

$$\begin{aligned}\mathbb{E}\left[\int_0^t f'(W(s))dL(s)\right] &\geq -\mathbb{E}\left[\int_0^t \kappa dL^c(s) + \sum_{s \leq t} \kappa \Delta L(s)\right] \\ &= -\kappa \mathbb{E}[L(t)], \quad t \geq 0.\end{aligned}\quad (\text{C.55})$$

Using a similar argument, we obtain

$$\mathbb{E}\left[\int_0^t f'(W(s))dU(s)\right] \leq 0, \quad t \geq 0. \quad (\text{C.56})$$

Moreover, by (C.52),

$$\frac{1}{2}\sigma^2 f''(W(t)) + \mu f'(W(t)) - \theta(t)f'(W(t)) + c(\theta(t)) + h(W(t)) \geq \gamma^*, \quad t \geq 0. \quad (\text{C.57})$$

It follows from (C.57) that for $t \geq 0$,

$$\begin{aligned} \mathbb{E} \left[\int_0^t \frac{1}{2} \sigma^2 f''(W(s)) ds + \int_0^t \mu f'(W(s)) ds - \int_0^t \theta(s) f'(W(s)) ds \right] &\geq \gamma^* t - \mathbb{E} \left[\int_0^t c(\theta(s)) ds \right] \\ &\quad - \mathbb{E} \left[\int_0^t h(W(s)) ds \right]. \end{aligned} \quad (\text{C.58})$$

Plugging (C.55)-(C.56) and (C.58) into (C.53) gives

$$\mathbb{E}[f(W(t))] - f(W(0)) \geq \gamma^* t - \mathbb{E} \left[\int_0^t c(\theta(s)) ds - \int_0^t h(W(s)) ds - \kappa L(t) \right] \quad (\text{C.59})$$

for $t \geq 0$. Divide both sides of (C.59) by t and take the \liminf as $t \rightarrow \infty$. The first term on the left hand side vanishes, because f' is bounded and W satisfies (3.31). The second term on the left hand side vanishes since $W(0) = 0$. Consequently, we arrive at

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \mathbb{E} \left[\int_0^t c(\theta(s)) ds + \int_0^t h(W(s)) ds + \kappa L(t) \right] \geq \gamma^*.$$

This completes the proof. \square

C.4 Auxiliary Results

This section discusses a number of auxiliary results that are crucially used in proving various lemmas and propositions in Appendix C.3.

Lemma 21. *Given $l < 0$, let $v \in C^1([l, 0])$ be the solution to*

$$v'(w) = \frac{m' H^{-1} m}{2\sigma^2} (v^2(w) - \kappa^2) - \frac{2\mu}{\sigma^2} (v(w) + \kappa) + \frac{2(h(l) - h(w))}{\sigma^2}, \quad w \in (l, 0) \quad (\text{C.60})$$

subject to the boundary condition $v(l) = -\kappa$. Then, $v(w) \geq -\kappa$ for $w \in [l, 0]$.

Proof. We prove the lemma by contradiction. Assume that there exists $w_0 \in (l, 0)$ such that $v(w_0) < -\kappa$. Then, it follows from the continuity of v and v' that there exists $w_1 \in (l, w_0)$

such that $v(w_1) < -\kappa$ and $v'(w_1) < 0$. Plugging $v(w_1) < -\kappa$ into (C.60) gives

$$v'(w_1) \geq \frac{2(h(l) - h(w_1))}{\sigma^2} > 0, \quad (\text{C.61})$$

which contradicts $v'(w_1) < 0$. Therefore, the assumption that there exists $w_0 \in (l, 0]$ with $v(w_0) < -\kappa$ is incorrect. This completes the proof. \square

Lemma 22. *Given $l < 0$, let $v \in C^1([l, 0])$ be the solution to*

$$v'(w) = \frac{m'H^{-1}m}{2\sigma^2} (v^2(w) - \kappa^2) - \frac{2\mu}{\sigma^2} (v(w) + \kappa) + \frac{2(h(l) - h(w))}{\sigma^2}, \quad w \in (l, 0) \quad (\text{C.62})$$

subject to the boundary condition $v(l) = -\kappa$. Then, $v'(w) > 0$ for $w \in (l, 0]$.

Proof. Since $v \in C^1([l, 0])$, the quadratic function is smooth, and h is continuously differentiable on $[l, 0]$, it follows from (C.62) that $v \in C^2([l, 0])$, i.e., v'' is well-defined on $[l, 0]$. Therefore, we can differentiate both sides of (C.62) to get

$$v''(w) = \frac{m'H^{-1}m}{\sigma^2} v(w) v'(w) - \frac{2\mu}{\sigma^2} v'(w) - \frac{2h'(w)}{\sigma^2}, \quad w \in (l, 0). \quad (\text{C.63})$$

By continuous differentiability of v , Equation (C.62), and boundary condition $v(l) = -\kappa$, we have $v'(l) = 0$. Then, since v and v' are continuous on $[l, 0]$, $v'(l) = 0$, and $h'(w) = -b^* < 0$ for $w \in [l, 0]$, it follows from (C.63) that there exists $w_0 \in (l, 0]$ such that $v''(w) > 0$ for $w \in (l, w_0)$. Then, since $v'(l) = 0$, we conclude that $v'(w) > 0$ for $w \in (l, w_0]$. It remains to show that $v'(w) > 0$ for $w \in (w_0, 0]$. To do so, let

$$w_1 = \inf \left\{ w \in (l, 0) : v(w) = \frac{2\mu}{m'H^{-1}m} \right\}. \quad (\text{C.64})$$

Without loss of generality, we assume $w_0 < w_1$. We prove $v'(w) > 0$ for $w \in (w_0, w_1)$ and $[w_1, 0]$, separately.

Case 1: $w \in (w_0, w_1)$. By (C.63), if $v''(w) < 0$,

$$\frac{m'H^{-1}m}{\sigma^2} v(w) v'(w) - \frac{2\mu}{\sigma^2} v'(w) - \frac{2h'(w)}{\sigma^2} < 0,$$

which is equivalent to

$$2b^* < (2\mu - m'H^{-1}m v(w)) v'(w). \quad (\text{C.65})$$

By (C.64), the fact that $w < w_1$, $v(l) = -\kappa < 0$, and the continuity of v , we have $2\mu - m'H^{-1}m v(w) > 0$. Thus, dividing both sides of (C.65) by $2\mu - m'H^{-1}m v(w)$ gives

$$v'(w) > \frac{2b^*}{2\mu - m'H^{-1}m v(w)}. \quad (\text{C.66})$$

It follows from (C.66) and the fact that $v(w) \geq -\kappa$ (by Lemma 21) that if $v''(w) < 0$,

$$v'(w) > \frac{2b^*}{2\mu + m'H^{-1}m \kappa}. \quad (\text{C.67})$$

In other words, letting

$$a_2 = \frac{2b^*}{2\mu + m'H^{-1}m \kappa} > 0, \quad (\text{C.68})$$

if $v'(w) \leq a_2$, we have $v''(w) > 0$. Then, since $a_1 = v'(w_0) > 0$, we conclude that $v'(w) \geq \min(a_1, a_2) > 0$ for $w \in (w_0, w_1)$.

Case 2: $w \in [w_1, 0]$. It follows from (C.63) that $v''(w_1) > 0$. Then, by the continuity of v' and the fact that $v'(w) \geq \min(a_1, a_2) > 0$ for $w \in (a, w_1)$, we have $a_3 = v'(w_1) \geq \min(a_1, a_2) > 0$. It remains to show that $v'(w) > 0$ for $w \in (w_1, 0]$. We prove this by contradiction. Assume that there exists $w_2 \in (w_1, 0]$ such that $v'(w_2) \leq 0$. Since $v'(w_1) > 0$ and $v'(w_2) \leq 0$, there exists $w_3 \in (w_1, w_2)$ such that $v'(w_3) > 0$ and $v''(w_3) \leq 0$. Without loss of generality, assume w_3 is the smallest such number, i.e., $v'(w) > 0$ for $w \in (w_1, w_3)$.

Then, it follows from $v(w_1) = 0$ (by (C.64)) and $v'(w) > 0$ for $w \in (w_1, w_3]$ that $v(w) > 2\mu/(m'H^{-1}m)$. Plugging this into (C.63) gives $v''(w_3) > 0$. This, however, contradicts the assumption that $v''(w_3) < 0$. Therefore, the initial assumption that there exists $w_2 \in (w_1, 0]$ such that $v'(w_2) \leq 0$ is incorrect. This completes the proof. \square

Lemma 23. *Let $v \in C^1([0, \infty))$ be the solution to*

$$v'(w) = -\frac{m'H^{-1}m}{2\sigma^2}v^2(w) + \frac{2\mu}{\sigma^2}v(w) - \frac{2h^*}{\sigma^2}w, \quad w \in (0, \infty) \quad (\text{C.69})$$

subject to the boundary condition $v(0) = 0$. Then, $v, v' < 0$ on $(0, \infty)$.

Proof. Since $v \in C^1([0, \infty))$ and the quadratic function is smooth, it follows from (C.69) that $v \in C^2([0, \infty))$, i.e., v'' is well-defined on $[0, \infty)$. Therefore, we can differentiate both sides of (C.69) to get

$$v''(w) = -\frac{m'H^{-1}m}{\sigma^2}v(w)v'(w) + \frac{2\mu}{\sigma^2}v'(w) - \frac{2h^*}{\sigma^2}, \quad w \in (0, \infty). \quad (\text{C.70})$$

By continuous differentiability of v , Equation (C.69), and boundary condition $v(0) = 0$, we have $v'(0) = 0$. Then, since v and v' are continuous on $[0, \infty)$, $v'(0) = 0$, and $-2h^*/\sigma^2 < 0$, it follows from (C.70) that there exists $w_0 \in (0, \infty)$ such that $v''(w) < 0$ for $w \in (0, w_0)$. Then, since $v(0) = 0$, we have $v'(w) < 0$ for $w \in (0, w_0]$.

Next, we show that $v'(w) < 0$ for $w \in (w_0, \infty)$. Let us prove this statement by contradiction. Assume that there exists $w_1 \in (w_0, \infty)$ such that $v(w_1) \geq 0$. Then, there must exist $w_2 \in (w_0, w_1)$ such that $v''(w_2) \geq 0$. Without loss of generality, assume that w_2 is the smallest value such that $v''(w_2) \geq 0$. Then, by plugging $v(w_0), v'(w_0) < 0$ into (C.70), we arrive at the contradiction $v''(w_2) < 0$. Thus, the assumption that there exists $w_1 \in (w_0, \infty)$ such that $v(w_1) \geq 0$ is incorrect, i.e., $v' < 0$ on $(0, \infty)$. Consequently, $v < 0$ on $(0, \infty)$. This completes the proof. \square

REFERENCES

- [1] (2014). 2014 taxicab factbook. Technical report, NYC Taxi & Limousine Commission.
- [2] (2016). 2016 tlc factbook. Technical report, NYC Taxi & Limousine Commission.
- [3] Abhishek, V., Dogan, M., and Jacquillat, A. (2018). Strategic timing and dynamic pricing in on-demand platforms. Working paper.
- [4] Abramowitz, M. and Stegun, I. A. (1965). *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*, volume 55. Courier Corporation.
- [5] Adlakha, S. and Johari, R. (2013). Mean field equilibrium in dynamic games with strategic complementarities. *Oper. Res.*, 61(4):971–989.
- [6] Adlakha, S., Johari, R., and Weintraub, G. Y. (2015). Equilibria of dynamic games with many players: Existence, approximation, and market structure. *J. Econ. Theory*, 156:269–316.
- [7] Adusumilli, K. M. and Hasenbein, J. J. (2010). Dynamic admission and service rate control of a queue. *Queueing Systems*, 66(2):131–154.
- [8] Afèche, P. (2013). Incentive-compatible revenue management in queueing systems: Optimal strategic delay. *M&SOM*, 15(3):423–443.
- [9] Afèche, P., Baron, O., and Kerner, Y. (2013). Pricing time-sensitive services based on realized performance. *M&SOM*, 15(3):492–506.
- [10] Afèche, P., Liu, Z., and Maglaras, C. (2018). Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. Working paper.
- [11] Afèche, P., Liu, Z., and Maglaras, C. (2020). Surge pricing and dynamic matching for hotspot demand shock in ride-hailing networks. Working paper.

[12] Aherne, F. J., Thacker, N. A., and Rockett, P. I. (1998). The Bhattacharyya metric as an absolute similarity measure for frequency coded data. *Kybernetika*, 34(4):363–368.

[13] Akşin, Z., Ata, B., Emadi, S. M., and Su, C.-L. (2013). Structural estimation of callers' delay sensitivity in call centers. *Manage. Sci.*, 59(12):2727–2746.

[14] Akşin, Z., Ata, B., Emadi, S. M., and Su, C.-L. (2016). Impact of delay announcements in call centers: An empirical approach. *Oper. Res.*, 65(1):242–265.

[15] Alwan, A. and Ata, B. (2020). A diffusion approximation framework for ride-hailing with travel delays. Working paper.

[16] Anderson, S. P., De Palma, A., and Thisse, J.-F. (1992). *Discrete choice theory of product differentiation*. MIT press.

[17] Aouad, A. and Saritac, O. (2019). Dynamic stochastic matching under limited time. Working paper.

[18] Ata, B. (2005). Dynamic power control in a wireless static channel subject to a quality-of-service constraint. *Oper. Res.*, 53(5):842–851.

[19] Ata, B. (2006). Dynamic control of a multiclass queue with thin arrival streams. *Oper. Res.*, 54(5):876–892.

[20] Ata, B. and Barjesteh, N. (2019). Dynamic pricing of a multiclass make-to-stock queue. Working paper.

[21] Ata, B., Barjesteh, N., and Kumar, S. (2019a). Dynamic dispatch and centralized relocation of cars in ride-hailing platforms. Working paper.

[22] Ata, B., Barjesteh, N., and Kumar, S. (2019b). Spatial pricing: An empirical analysis of taxi rides in New York City. Working paper.

[23] Ata, B., Friedewald, J., and Randa, A. C. (2020a). An equilibrium framework for assessing the impact of policy changes under endogenous patient choice. Working paper.

[24] Ata, B., Friedewald, J., and Randa, A. C. (2020b). Structural estimation of kidney transplant candidates' quality of life scores: Improving national kidney allocation policy under endogenous patient choice and geographical sharing. Working paper.

[25] Ata, B., Glynn, P. W., and Peng, X. (2017). An equilibrium analysis of a discrete-time markovian queue with endogenous abandonments. *Queueing Systems*, 86(1-2):141–212.

[26] Ata, B., Harrison, J. M., and Shepp, L. A. (2005). Drift rate control of a Brownian processing system. *Ann. Appl. Probab.*, 15(2):1145–1160.

[27] Ata, B. and Kumar, S. (2005). Heavy traffic analysis of open processing networks with complete resource pooling: Asymptotic optimality of discrete review policies. *Ann. Appl. Probab.*, 15(1A):331–391.

[28] Ata, B., Lee, D., and Sönmez, E. (2019c). Dynamic volunteer staffing in multicrop gleaning operations. *Oper. Res.*, 67(2):295–314.

[29] Ata, B. and Lin, W. (2008). Heavy traffic analysis of maximum pressure policies for stochastic processing networks with multiple bottlenecks. *Queueing Systems*, 59(3-4):191–235.

[30] Ata, B. and Olsen, T. L. (2009). Near-optimal dynamic lead-time quotation and scheduling under convex-concave customer delay costs. *Oper. Res.*, 57(3):753–768.

[31] Ata, B. and Olsen, T. L. (2013). Congestion-based leadtime quotation and pricing for revenue maximization with heterogeneous customers. *Queueing Systems*, 73(1):35–78.

[32] Ata, B. and Peng, X. (2017). An equilibrium analysis of a multiclass queue with endogenous abandonments in heavy traffic. *Oper. Res.*, 66(1):163–183.

[33] Ata, B. and Shneorson, S. (2006). Dynamic control of an M/M/1 service system with adjustable arrival and service rates. *Manage. Sci.*, 52(11):1778–1791.

[34] Ata, B. and Tongarlak, M. H. (2013). On scheduling a multiclass queue with abandonments under general delay costs. *Queueing Systems*, 74(1):65–104.

[35] Ata, B. and Zachariadis, K. E. (2007). Dynamic power control in a fading downlink channel subject to an energy constraint. *Queueing Systems*, 55(1):41–69.

[36] Bai, J., So, K. C., Tang, C. S., Chen, X., and Wang, H. (2019). Coordinating supply and demand on an on-demand service platform with impatient customers. *M&SOM*, 21(3):556–570.

[37] Balseiro, S. R., Besbes, O., and Weintraub, G. Y. (2015). Repeated auctions with budgets in ad exchanges: Approximations and design. *Manage. Sci.*, 61(4):864–884.

[38] Balseiro, S. R., Brown, D. B., and Chen, C. (2020). Dynamic pricing of relocating resources in large networks. *Manage. Sci.*

[39] Banerjee, S., Freund, D., and Lykouris, T. (2016). Pricing and optimization in shared vehicle systems: An approximation framework. Working paper.

[40] Banerjee, S., Kanoria, Y., and Qian, P. (2018). State dependent control of closed queueing networks with application to ride-hailing. Working paper.

[41] Banerjee, S., Riquelme, C., and Johari, R. (2015). Pricing in ride-share platforms: A queueing-theoretic approach.

[42] Barber, C. B., Dobkin, D. P., and Huhdanpaa, H. (1996). The Quickhull algorithm for convex hulls. *ACM T. Math. Software*, 22(4):469–483.

[43] Basseville, M. (1989). Distance measures for signal processing and pattern recognition. *Signal processing*, 18(4):349–369.

[44] Bell, S. L. and Williams, R. J. (2001). Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: Asymptotic optimality of a threshold policy. *Ann. Appl. Probab.*, 11(3):608–649.

[45] Bell, S. L. and Williams, R. J. (2005). Dynamic scheduling of a parallel server system in heavy traffic with complete resource pooling: Asymptotic optimality of a threshold policy. *Electron. J. Probab.*, 10:1044–1115.

[46] Ben-Akiva, M. E., Lerman, S. R., and Lerman, S. R. (1985). *Discrete choice analysis: Theory and application to travel demand*, volume 9. MIT press.

[47] Benjaafar, S., Jiang, D., Li, X., and Li, X. (2018). Inventory repositioning in on-demand product rental networks. Working paper.

[48] Bertsimas, D., Jaillet, P., and Martin, S. (2019). Online vehicle routing: The edge of optimization in large-scale applications. *Oper. Res.*, 67(1):143–162.

[49] Besbes, O., Castro, F., and Lobel, I. (2018). Spatial capacity planning. Working paper.

[50] Besbes, O., Castro, F., and Lobel, I. (2020). Surge pricing and its spatial supply response. *Manage. Sci.*

[51] Bhattacharyya, A. (1946). On a measure of divergence between two multinomial populations. *Sankhyā: the Indian journal of statistics*, pages 401–406.

[52] Bimpikis, K., Candogan, O., and Saban, D. (2019). Spatial pricing in ride-sharing networks. *Oper. Res.*, 67(3):744–769.

[53] Blanchet, J. H., Reiman, M. I., Shah, V., and Wein, L. M. (2020). Asymptotically optimal control of a centralized dynamic matching market with general utilities. Working paper.

[54] Bramson, M. and Williams, R. (2003). Two workload properties for Brownian networks. *Queueing Systems*, 45(3):191–221.

[55] Braverman, A., Dai, J. G., Liu, X., and Ying, L. (2019). Empty-car routing in ridesharing systems. *Oper. Res.*, 67(5):1437–1452.

[56] Buchholz, N. (2018). Spatial equilibrium, search frictions and efficient regulation in the taxi industry. Working paper.

[57] Budhiraja, A., Ghosh, A. P., and Lee, C. (2011). Ergodic rate control problem for single class queueing networks. *SIAM J. Control Optim.*, 49(4):1570–1606.

[58] Budhiraja, A. and Ross, K. (2006). Existence of optimal controls for singular control problems with state constraints. *Ann. Appl. Probab.*, 16(4):2235–2255.

[59] Burdett, K., Shi, S., and Wright, R. (2001). Pricing and matching with frictions. *J. Polit. Econ.*, 109(5):1060–1085.

[60] Byrd, R. H., Nocedal, J., and Waltz, R. A. (2006). An integrated package for nonlinear optimization. In *Large-scale Nonlinear Optimization*, pages 35–59. Springer.

[61] Cachon, G. P., Daniels, K. M., and Lobel, R. (2017). The role of surge pricing on a service platform with self-scheduling capacity. *M&SOM*.

[62] Caldentey, R. and Wein, L. M. (2006). Revenue management of a make-to-stock queue. *Oper. Res.*, 54(5):859–875.

[63] Cao, P. and Yao, D. (2018). Optimal drift rate control and impulse control for a stochastic inventory/production system. *SIAM J. Control Optim.*, 56(3):1856–1883.

[64] Carlsson, J. G., Behroozi, M., and Mihic, K. (2018). Wasserstein distance and the distributionally robust tsp. *Oper. Res.*, 66(6):1603–1624.

[65] Castillo, J. C. (2018). Who benefits from surge pricing? Working paper.

[66] Castillo, J. C., Knoepfle, D., and Weyl, G. (2018). Surge pricing solves the wild goose chase. Working paper.

[67] Çelik, S. and Maglaras, C. (2008). Dynamic pricing and lead-time quotation for a multiclass make-to-order queue. *Manage. Sci.*, 54(6):1132–1146.

[68] Chao, X. and Zhou, S. X. (2006). Joint inventory-and-pricing strategy for a stochastic continuous-review system. *IIE Transactions*, 38(5):401–408.

[69] Chen, M., Dai, J., Sun, P., and Wan, Z. (2019). Matching supply and demand with mismatch-sensitive players. Working paper.

[70] Chen, M. K. and Sheldon, M. (2016). Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform. Working paper.

[71] Chen, Q., Lei, Y., and Jasin, S. (2020). Real-time spatial-intertemporal dynamic pricing for balancing supply and demand in a network. Working paper.

[72] Chevalier, P. B. and Wein, L. M. (1993). Scheduling networks of queues: Heavy traffic analysis of a multistation closed network. *Oper. Res.*, 41(4):743–758.

[73] Chu, L. Y., Wan, Z., and Zhan, D. (2018). Harnessing the double-edged sword via routing: Information provision on ride-hailing platforms. Working paper.

[74] Clark, C. E. (1961). The greatest of a finite set of random variables. *Oper. Res.*, 9(2):145–162.

[75] Coffman, R. B. and Shreiber, C. (1977). The economic reasons for price and entry regulation of taxicabs (comment and rejoinder). *J. Transp. Econ. Policy*, pages 288–304.

[76] Cohen, P., Hahn, R., Hall, J., Levitt, S., and Metcalfe, R. (2016). Using big data to estimate consumer surplus: The case of Uber. Working paper.

[77] Crabbill, T. B. (1972). Optimal control of a service facility with variable exponential service times and constant arrival rate. *Manage. Sci.*, 18(9):560–566.

[78] Crabill, T. B. (1974). Optimal control of a maintenance system with variable service rates. *Oper. Res.*, 22(4):736–745.

[79] Cramer, J. and Krueger, A. B. (2016). Disruptive change in the taxi business: The case of Uber. *Am. Econ. Rev.*, 106(5):177–182.

[80] Dai, J. and Tezcan, T. (2011). State space collapse in many-server diffusion limits of parallel server systems. *Math. Oper. Res.*, 36(2):271–320.

[81] Dai, J. and Yao, D. (2013a). Brownian inventory models with convex holding cost, part 1: Average-optimal controls. *Stochastic Systems*, 3(2):442–499.

[82] Dai, J. and Yao, D. (2013b). Brownian inventory models with convex holding cost, part 2: Discount-optimal controls. *Stochastic Systems*, 3(2):500–573.

[83] Dai, J. G. and Lin, W. (2008). Asymptotic optimality of maximum pressure policies in stochastic processing networks. *Ann. Appl. Probab.*, 18(6):2239–2299.

[84] Duque, J. C., Anselin, L., and Rey, S. J. (2012). The max-p-regions problem. *Journal of Regional Science*, 52(3):397–419.

[85] Elmaghraby, W. and Keskinocak, P. (2003). Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Manage. Sci.*, 49(10):1287–1309.

[86] Ethier, S. N. and Kurtz, T. G. (2009). *Markov processes: characterization and convergence*, volume 282. John Wiley & Sons.

[87] Farber, H. S. (2008). Reference-dependent preferences and labor supply: The case of New York City taxi drivers. *Am. Econ. Rev.*, 98(3):1069–82.

[88] Feng, G., Kong, G., and Wang, Z. (2020a). We are on the way: Analysis of on-demand ride-hailing systems. *M&SOM*.

[89] Feng, Y., Niazadeh, R., and Saberi, A. (2020b). Two-stage matching and pricing with applications to ride hailing. Working paper.

[90] Fleming, W. H. and Soner, H. M. (2006). *Controlled Markov processes and viscosity solutions*, volume 25. Springer Science & Business Media.

[91] Foerster, J. F. and Gilbert, G. (1979). Taxicab deregulation: Economic consequences and regulatory choices. *Transportation*, 8(4):371–387.

[92] Frechette, G. R., Lizzeri, A., and Salz, T. (2019). Frictions in a competitive, regulated market: Evidence from taxis. *Am. Econ. Rev.*, 109(8):2954–92.

[93] Gallego, G. and Topaloglu, H. (2018). *Revenue management and pricing analytics*. Springer.

[94] Garg, N. and Nazerzadeh, H. (2019). Driver surge pricing. Working paper.

[95] Gayon, J.-P. and Dallery, Y. (2007). Dynamic vs static pricing in a make-to-stock queue with partially controlled production. *OR Spectrum*, 29(2):193–205.

[96] Gentry, S., Chow, E., Massie, A., and Segev, D. (2015). Gerrymandering for justice: Redistricting US liver allocation. *Interfaces*, 45(5):462–480.

[97] George, J. M. and Harrison, J. M. (2001). Dynamic control of a queue with adjustable service rate. *Oper. Res.*, 49(5):720–731.

[98] Ghamami, S. and Ward, A. R. (2013). Dynamic scheduling of a two-server parallel server system with complete resource pooling and reneging in heavy traffic: Asymptotic optimality of a two-threshold policy. *Math. Oper. Res.*, 38(4):761–824.

[99] Ghosh, A. P. and Weerasinghe, A. P. (2007). Optimal buffer size for a stochastic processing network in heavy traffic. *Queueing Systems*, 55(3):147–159.

[100] Ghosh, A. P. and Weerasinghe, A. P. (2010). Optimal buffer size and dynamic rate control for a queueing system with impatient customers in heavy traffic. *Stoch. Process. Their Appl.*, 120(11):2103–2141.

[101] Gilbarg, D. and Trudinger, N. S. (2015). *Elliptic partial differential equations of second order*. Springer.

[102] Guda, H. and Subramanian, U. (2019). Your Uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. *Manage. Sci.*, 65(5):1995–2014.

[103] Gummadi, R., Johari, R., Schmit, S., and Yu, J. Y. (2013). Mean field analysis of multi-armed bandit games. Working paper.

[104] Ha, A. Y. (1997). Optimal dynamic scheduling policy for a make-to-stock production system. *Oper. Res.*, 45(1):42–53.

[105] Häckner, J. and Nyberg, S. (1995). Deregulating taxi services: A word of caution. *J. Transp. Econ. Policy*, pages 195–207.

[106] Hall, J., Kendrick, C., and Nosko, C. (2015). The effects of Uber’s surge pricing: A case study. *The University of Chicago Booth School of Business*.

[107] Hall, R. E. (1979). A theory of the natural unemployment rate and the duration of employment. *J. Monetary Econ.*, 5(2):153–169.

[108] Hao, Z., He, L., Hu, Z., and Jiang, J. (2018). Robust vehicle allocation with uncertain covariates. Working paper.

[109] Hardin, J. W., Hilbe, J. M., and Hilbe, J. (2007). *Generalized linear models and extensions*. Stata Press.

[110] Harrison, J. M. (1988). Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic differential systems, stochastic control theory and applications*, pages 147–186. Springer.

[111] Harrison, J. M. (1996). The BIGSTEP approach to flow management in stochastic processing networks. *Stochastic Networks: Theory and Applications*, 4:147–186.

[112] Harrison, J. M. (1998). Heavy traffic analysis of a system with parallel servers: Asymptotic optimality of discrete-review policies. *Ann. Appl. Probab.*, pages 822–848.

[113] Harrison, J. M. (2000). Brownian models of open processing networks: Canonical representation of workload. *Ann. Appl. Probab.*, 10(1):75–103.

[114] Harrison, J. M. (2003). A broader view of Brownian networks. *Ann. Appl. Probab.*, 13(3):1119–1150.

[115] Harrison, J. M. (2013). *Brownian models of performance and control*. Cambridge University Press.

[116] Harrison, J. M. and López, M. J. (1999). Heavy traffic resource pooling in parallel-server systems. *Queueing systems*, 33(4):339–368.

[117] Harrison, J. M., Sellke, T. M., and Taylor, A. J. (1983). Impulse control of Brownian motion. *Math. Oper. Res.*, 8(3):454–466.

[118] Harrison, J. M. and Taksar, M. I. (1983). Instantaneous control of Brownian motion. *Math. Oper. Res.*, 8(3):439–453.

[119] Harrison, J. M. and Taylor, A. J. (1978). Optimal control of a Brownian storage system. *Stoch. Process. Their Appl.*, 6(2):179–194.

[120] Harrison, J. M. and Van Mieghem, J. A. (1997). Dynamic control of Brownian networks: State space collapse and equivalent workload formulations. *Ann. Appl. Probab.*, 7(3):747–771.

[121] Harrison, J. M. and Wein, L. M. (1989). Scheduling networks of queues: Heavy traffic analysis of a simple open network. *Queueing Systems*, 5(4):265–279.

[122] Harrison, J. M. and Wein, L. M. (1990). Scheduling networks of queues: Heavy traffic analysis of a two-station closed network. *Oper. Res.*, 38(6):1052–1064.

[123] Harrison, J. M. and Williams, R. J. (2005). Workload reduction of a generalized Brownian network. *Ann. Appl. Probab.*, 15(4):2255–2295.

[124] Harrison, J. M. and Williams, R. J. (2007). Workload interpretation for Brownian models of stochastic processing networks. *Math. Oper. Res.*, 32(4):808–820.

[125] He, L., Hu, Z., and Zhang, M. (2020). Robust repositioning for vehicle sharing. *M&SOM*, 22(2):241–256.

[126] Hilbe, J. M. (2011). *Negative binomial regression*. Cambridge University Press.

[127] Horowitz, J. L. (2001). The bootstrap. In *Handbook of econometrics*, volume 5, pages 3159–3228. Elsevier.

[128] Hosseini, M., Milner, J., and Romero, G. (2020). Efficient repositioning in car-sharing networks. Working paper.

[129] Hu, B., Hu, M., and Zhu, H. (2019). Surge pricing and two-sided temporal responses in ride-hailing. Working paper.

[130] Hu, M. and Zhou, Y. (2019). Price, wage and fixed commission in on-demand matching. Working paper.

[131] Huang, M., Caines, P. E., and Malhamé, R. P. (2003). Individual and mass behaviour in large population stochastic wireless power control problems: Centralized and nash equilibrium solutions. In *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, volume 1, pages 98–103. IEEE.

[132] Hughes, T. J. (2012). *The finite element method: linear static and dynamic finite element analysis*. Courier Corporation.

[133] Iyer, K., Johari, R., and Sundararajan, M. (2014). Mean field equilibria of dynamic auctions with learning. *Manage. Sci.*, 60(12):2949–2970.

[134] Jacob, J. and Roet-Green, R. (2018). Ride solo or pool: designing price-service menus for a ride-sharing platform. Working paper.

[135] Kailath, T. (1967). The divergence and Bhattacharyya distance measures in signal selection. *IEEE T. Commun. Techn.*, 15(1):52–60.

[136] Kanoria, Y. and Qian, P. (2019). Blind dynamic resource allocation in closed networks via mirror backpressure. Working paper.

[137] Karaesmen, F., Buzacott, J. A., and Dallery, Y. (2002). Integrating advance order information in make-to-stock production systems. *IIE Transactions*, 34(8):649–662.

[138] Kim, E. and Van Oyen, M. P. (1998). Dynamic scheduling to minimize lost sales subject to set-up costs. *Queueing Systems*, 29(2-4):193–229.

[139] Kohavi, R. et al. (1995). A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Ijcai*, volume 14, pages 1137–1145. Montreal, Canada.

[140] Krichagina, E. V., Lou, S. X., Sethi, S. P., and Taksar, M. I. (1993). Production control in a failure-prone manufacturing system: Diffusion approximation and asymptotic optimality. *Ann. Appl. Probab.*, 3(2):421–453.

[141] Krichagina, E. V. and Taksar, M. I. (1992). Diffusion approximation for GI/G/1 controlled queues. *Queueing systems*, 12(3-4):333–367.

[142] Krichagina, E. V. and Wein, L. M. (1994). Heavy traffic analysis of a production-distribution system. Working paper.

[143] Kumar, R., Lewis, M. E., and Topaloglu, H. (2013). Dynamic service rate control for a single-server queue with Markov-modulated arrivals. *Nav. Res. Log.*, 60(8):661–677.

[144] Kumar, S. (2000). Two-server closed networks in heavy traffic: Diffusion limits and asymptotic optimality. *Ann. Appl. Probab.*, 10(3):930–961.

[145] Kumar, S. and Muthuraman, K. (2004). A numerical method for solving singular stochastic control problems. *Oper. Res.*, 52(4):563–582.

[146] Kushner, H. J. (1977). *Probability Methods for Approximations in Stochastic Control and for Elliptic Equations*. Academic Press, Inc.

[147] Kushner, H. J. and Dupuis, P. G. (2013). *Numerical methods for stochastic control problems in continuous time*, volume 24. Springer Science & Business Media.

[148] Kushner, H. J. and Martins, L. F. (1991). Numerical methods for stochastic singular control problems. *SIAM J. Control Optim.*, 29(6):1443–1475.

[149] Lagos, R. (2000). An alternative approach to search frictions. *J. Polit. Econ.*, 108(5):851–873.

[150] Lam, C. T. and Liu, M. (2017). Demand and consumer surplus in the on-demand economy: The case of ride sharing. Working paper.

[151] Larson, M. G. and Bengzon, F. (2013). *The finite element method: theory, implementation, and applications*, volume 10. Springer Science & Business Media.

[152] Lasry, J.-M. and Lions, P.-L. (2007). Mean field games. *Jap. J. Math.*, 2(1):229–260.

[153] Lawless, J. F. (1987). Negative binomial and mixed poisson regression. *Canadian Journal of Statistics*, 15(3):209–225.

[154] Le Dret, H. and Lucquin, B. (2016). *Partial differential equations: modeling, analysis and numerical approximation*, volume 168. Springer.

[155] Li, J., Granados, N., and Netessine, S. (2014). Are consumers strategic? structural estimation from the air-travel industry. *Manage. Sci.*, 60(9):2114–2137.

[156] Lieberman, G. M. (2013). *Oblique derivative problems for elliptic equations*. World Scientific.

[157] Lu, A., Frazier, P., and Kislev, O. (2018). Surge pricing moves Uber’s driver partners. Working paper.

[158] Luo, Q. and Saigal, R. (2017). Dynamic pricing for on-demand ride-sharing: A continuous approach. Working paper.

[159] Lyu, G., Cheung, W. C., Teo, C.-P., and Wang, H. (2019). Multi-objective online ride-matching. Working paper.

[160] Maglaras, C. (2000). Discrete-review policies for scheduling stochastic networks: Trajectory tracking and fluid-scale asymptotic optimality. *Ann. Appl. Probab.*, 10(3):897–929.

[161] Maglaras, C. and Zeevi, A. (2003). Pricing and capacity sizing for systems with shared resources: Approximate solutions and scaling relations. *Manage. Sci.*, 49(8):1018–1038.

[162] Maglaras, C. and Zeevi, A. (2005). Pricing and design of differentiated services: Approximate analysis and structural insights. *Oper. Res.*, 53(2):242–262.

[163] Markowitz, D. M. and Wein, L. M. (2001). Heavy traffic analysis of dynamic cyclic policies: A unified treatment of the single machine scheduling problem. *Oper. Res.*, 49(2):246–270.

[164] Mendelson, H. (1985). Pricing computer services: Queueing effects. *Commun. ACM*, 28(3):312–321.

[165] Mendelson, H. and Whang, S. (1990). Optimal incentive-compatible priority pricing for the M/M/1 queue. *Oper. Res.*, 38(5):870–883.

[166] Ming, L., Tunca, T., Xu, Y., and Zhu, W. (2017). An empirical analysis of price formation, utilization, and value generation in ride sharing services. Working paper.

[167] Muthuraman, K. and Kumar, S. (2006). Multidimensional portfolio optimization with proportional transaction costs. *Mathematical Finance*, 16(2):301–335.

[168] Nahmias, S. and Olsen, T. L. (2015). *Production and operations analysis*. Waveland Press.

[169] Nair, H. (2007). Intertemporal price discrimination with forward-looking consumers: Application to the US market for console video-games. *Quant. Mark. Econ.*, 5(3):239–292.

[170] Øksendal, B. and Sulem, A. (2007). *Applied stochastic control of jump diffusions*. Springer Science & Business Media.

[171] Özkan, E. (2020). Joint pricing and matching in ride-sharing systems. *Eur. J. Oper. Res.*

[172] Özkan, E. and Ward, A. R. (2020). Dynamic matching for real-time ride sharing. *Stochastic Systems*, 10(1):29–70.

[173] Pele, O. and Werman, M. (2009). Fast and robust earth mover’s distances. In *2009 IEEE 12th International Conference on Computer Vision*, pages 460–467. IEEE.

[174] Pesic, V. and Williams, R. J. (2016). Dynamic scheduling for parallel server systems in heavy traffic: Graphical structure, decoupled workload matrix and some sufficient conditions for solvability of the brownian control problem. *Stochastic Systems*, 6(1):26–89.

[175] Plambeck, E., Kumar, S., and Harrison, J. M. (2001). A multiclass queue in heavy traffic with throughput time constraints: Asymptotically optimal dynamic controls. *Queueing Systems*, 39(1):23–54.

[176] Plambeck, E. L. (2004). Optimal leadtime differentiation via diffusion approximations. *Oper. Res.*, 52(2):213–228.

[177] Reed, J. and Ward, A. R. (2008). Approximating the GI/GI/1+GI queue with a nonlinear drift diffusion: Hazard rate scaling in heavy traffic. *Math. Oper. Res.*, 33(3):606–644.

[178] Reiman, M. I. and Wein, L. M. (1998). Dynamic scheduling of a two-class queue with setups. *Oper. Res.*, 46(4):532–547.

[179] Rheingans-Yoo, D., Kominers, S. D., Ma, H., and Parkes, D. C. (2019). Ridesharing with driver location preferences. Working paper.

[180] Royden, H. L. and Fitzpatrick, P. (1968). *Real analysis*. Macmillan New York, 4 edition.

[181] Rubino, M. and Ata, B. (2009). Dynamic control of a make-to-order, parallel-server system with cancellations. *Oper. Res.*, 57(1):94–108.

[182] Rubio, R. and Wein, L. M. (1996). Setting base stock levels using product-form queueing networks. *Manage. Sci.*, 42(2):259–268.

[183] Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, pages 999–1033.

[184] Sanajian, N., Abouee-Mehrizi, H., and Balcio glu, B. (2010). Scheduling policies in the M/G/1 make-to-stock queue. *J. Oper. Res. Soc.*, 61(1):115–123.

[185] Sayler, B. (2017). IBISWorld industry report 48533 taxi & limousine services in the US.

[186] Schneider, T. W. (2018). Analyzing 1.1 billion NYC taxi and Uber trips, with a vengeance.

[187] Schroeter, J. R. (1983). A model of taxi service under fare structure and fleet size regulation. *Bell J. Econ.*, pages 81–96.

[188] Selcuk, C. and Gokpinar, B. (2019). The selection of prices and commissions in a spatial model of ride-hailing. Working paper.

[189] Shapiro, M. H. (2018). Density of demand and the benefit of Uber. Working paper.

[190] Stidham, S. (1988). Scheduling, routing, and flow control in stochastic networks. In *Stochastic Differential Systems, Stochastic Control Theory and Applications*, pages 529–561. Springer.

[191] Stidham, S. (1992). Pricing and capacity decisions for a service facility: Stability and multiple local optima. *Manage. Sci.*, 38(8):1121–1139.

[192] Stidham, S. (2002). Analysis, design, and control of queueing systems. *Oper. Res.*, 50(1):197–216.

[193] Stidham, S. and Weber, R. R. (1989). Monotonic and insensitive optimal policies for control of queues with undiscounted costs. *Oper. Res.*, 37(4):611–625.

[194] Su, C.-L. and Judd, K. L. (2012). Constrained optimization approaches to estimation of structural models. *Econometrica*, 80(5):2213–2230.

[195] Taksar, M. I. (1985). Average optimal singular control and a related stopping problem. *Math. Oper. Res.*, 10(1):63–81.

[196] Taksar, M. I. (1997). Infinite-dimensional linear programming approach to singular stochastic control. *SIAM J. Control Optim.*, 35(2):604–625.

[197] Talluri, K. T. and Van Ryzin, G. J. (2006). *The theory and practice of revenue management*, volume 68. Springer Science & Business Media.

[198] Tezcan, T. and Dai, J. (2010). Dynamic control of N-systems with many servers: Asymptotic optimality of a static priority policy in heavy traffic. *Oper. Res.*, 58(1):94–110.

[199] Van Zandt, T. (2012). *Firms, Prices, and Markets*.

[200] Varma, S. M., Bumpensanti, P., Maguluri, S. T., and Wang, H. (2019). Dynamic pricing and matching for two-sided queues. Working paper.

[201] Veatch, M. H. and Véricourt, F. D. (2003). Zero-inventory conditions for a two-part-type make-to-stock production system. *Queueing systems*, 43(3):251–266.

[202] Veatch, M. H. and Wein, L. M. (1996). Scheduling a make-to-stock queue: Index policies and hedging points. *Oper. Res.*, 44(4):634–647.

[203] Véricourt, F. D., Karaesmen, F., and Dallery, Y. (2000). Dynamic scheduling in a make-to-stock system: A partial characterization of optimal policies. *Oper. Res.*, 48(5):811–819.

[204] Vulcano, G. (2008). Dynamic pricing for an M/M/1 make-to-stock system with controllable backlog. Working paper.

[205] Wang, G., Zhang, H., and Zhang, J. (2019). On-demand ride-matching in a spatial model with abandonment and cancellation. Working paper.

[206] Wang, X., Agatz, N., and Erera, A. (2017). Stable matching for dynamic ride-sharing systems. *Transport. Sci.*, 52(4):850–867.

[207] Ward, A. R. and Armony, M. (2013). Blind fair routing in large-scale service systems with heterogeneous customers and servers. *Oper. Res.*, 61(1):228–243.

[208] Ward, A. R. and Kumar, S. (2008). Asymptotically optimal admission control of a queue with impatient customers. *Math. Oper. Res.*, 33(1):167–202.

[209] Weerasinghe, A. (2015). Optimal service rate perturbations of many server queues in heavy traffic. *Queueing Systems*, 79(3-4):321–363.

[210] Weerasinghe, A. P. (2002). Stationary stochastic control for itô processes. *Adv. Appl. Probab.*, 34(1):128–140.

[211] Wein, L. M. (1990). Scheduling networks of queues: Heavy traffic analysis of a two-station network with controllable inputs. *Oper. Res.*, 38(6):1065–1078.

[212] Wein, L. M. (1991). Due-date setting and priority sequencing in a multiclass M/G/1 queue. *Manage. Sci.*, 37(7):834–850.

[213] Wein, L. M. (1992a). Dynamic scheduling of a multiclass make-to-stock queue. *Oper. Res.*, 40(4):724–735.

[214] Wein, L. M. (1992b). Scheduling networks of queues: Heavy traffic analysis of a multistation network with controllable inputs. *Oper. Res.*, 40(3-supplement-2):S312–S334.

[215] Xu, J. and Hajek, B. (2013). The supermarket game. *Stochastic Systems*, 3(2):405–441.

[216] Xu, Y. and Chao, X. (2009). Dynamic pricing and inventory control for a production system with average profit criterion. *Probab. Eng. Inform. Sc.*, 23(3):489–513.

[217] Yan, C., Zhu, H., Korolko, N., and Woodard, D. (2019). Dynamic pricing and matching in ride-hailing platforms. *Nav. Res. Log.*

[218] Yang, P., Iyer, K., and Frazier, P. (2018). Mean field equilibria for resource competition in spatial settings. *Stochastic Systems*, 8(4):307–334.

[219] Yoon, S. and Lewis, M. E. (2004). Optimal pricing and admission control in a queueing system with periodically varying parameters. *Queueing Systems*, 47(3):177–199.

[220] Yu, J. J., Tang, C. S., Max Shen, Z.-J., and Chen, X. M. (2020). A balancing act of regulating on-demand ride services. *Manage. Sci.*, 66(7):2975–2992.

[221] Zaitsev, V. F. and Polyanin, A. D. (2002). *Handbook of exact solutions for ordinary differential equations*. Chapman and Hall/CRC.

[222] Zheng, F. (2016). Spatial competition and preemptive entry in the discount retail industry. Working paper.

[223] Zheng, F., He, P., Belavina, E., and Girotra, K. (2020). Customer preference and station network in the London bike share system. *Manage. Sci.*

[224] Zheng, Y.-S. and Zipkin, P. (1990). A queueing model to analyze the value of centralized inventory information. *Oper. Res.*, 38(2):296–307.

[225] Zipkin, P. H. (1995). Performance analysis of a multi-item production-inventory system under alternative policies. *Manage. Sci.*, 41(4):690–703.

[226] Ziya, S., Ayhan, H., and Foley, R. D. (2006). Optimal prices for finite capacity queueing systems. *Oper. Res. Lett.*, 34(2):214–218.

[227] Ziya, S., Ayhan, H., and Foley, R. D. (2008). A note on optimal pricing for finite capacity queueing systems with multiple customer classes. *Nav. Res. Log.*, 55(5):412–418.