

THE UNIVERSITY OF CHICAGO

ESSAYS ON STOCHASTIC MODELS FOR RIDESHARING AND ONLINE 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

AMIR ANASTASIOS ALWAN

CHICAGO, ILLINOIS

JUNE 2024

Copyright © 2024 by Amir Anastasios Alwan  
All Rights Reserved

To Jed

# TABLE OF CONTENTS

LIST OF FIGURES . . . . .	vi
LIST OF TABLES . . . . .	vii
ACKNOWLEDGMENTS . . . . .	viii
ABSTRACT . . . . .	ix
1 A DIFFUSION APPROXIMATION FRAMEWORK FOR RIDE-HAILING WITH TRAVEL DELAYS . . . . .	1
1.1 Introduction . . . . .	1
1.2 Literature Review . . . . .	3
1.3 Ridesharing Model . . . . .	6
1.3.1 Model Primitives . . . . .	7
1.3.2 State Dynamics . . . . .	9
1.4 Heavy Traffic Limit Theorem . . . . .	11
1.5 Auxiliary Results . . . . .	16
1.6 Proof of the Main Theorem . . . . .	18
1.6.1 Convergence of Fluid Scaled Processes . . . . .	18
1.6.2 Convergence of Diffusion Scaled Processes . . . . .	21
1.7 Concluding Remarks . . . . .	24
1.8 Appendix . . . . .	24
1.8.1 Notation and Terminology . . . . .	24
1.8.2 Proofs for Section 1.5 . . . . .	27
1.8.3 Proof of Lemma 4 . . . . .	36
1.8.4 Miscellaneous Proofs . . . . .	44
1.8.5 Miscellaneous Derivations . . . . .	48
2 A QUEUEING MODEL OF DYNAMIC PRICING AND DISPATCH CONTROL FOR RIDE-HAILING SYSTEMS INCORPORATING TRAVEL TIMES . . . . .	51
2.1 Introduction . . . . .	51
2.2 Literature Review . . . . .	54
2.3 Ridesharing Model . . . . .	57
2.4 Brownian Control Problem . . . . .	64
2.5 Equivalent Workload Formulation . . . . .	71
2.6 Solution to the Equivalent Workload Formulation . . . . .	77
2.7 Solution to the Bellman Equation . . . . .	83
2.8 Proposed Policy . . . . .	104
2.9 Simulation Study . . . . .	106
2.10 Concluding Remarks . . . . .	114
2.11 Appendix . . . . .	114
2.11.1 Formal Derivations . . . . .	114

2.11.2	Miscellaneous Proofs . . . . .	119
3	A STOCHASTIC MODELING FRAMEWORK FOR ONLINE CONTENT MOD- ERATION: MANAGING CONTENT VIA MACHINE LEARNING AND HUMAN LABELING . . . . .	131
3.1	Introduction . . . . .	131
3.2	Literature Review . . . . .	133
3.3	Stochastic Model for Online Content Moderation . . . . .	135
3.3.1	User-Generated Content . . . . .	136
3.3.2	Content Classifiers . . . . .	138
3.4	Simulation Study . . . . .	144
3.5	A Related Model via a Bayesian Learning Framework . . . . .	149
3.5.1	User-Generated Content and Content Scores . . . . .	150
3.5.2	Content Arrival Process . . . . .	151
3.5.3	Automated Content Moderation . . . . .	152
3.5.4	Human-Based Content Moderation . . . . .	153
3.5.5	The Stochastic Control Problem . . . . .	156
3.5.6	Weak Convergence of the Belief Process . . . . .	158
3.6	Concluding Remarks . . . . .	163
	REFERENCES . . . . .	165

## LIST OF FIGURES

1.1	A map of New York City partitioned into nine geographical regions. . . . .	2
2.1	A network with four regions and ten dispatch activities. The open rectangles are the buffers, the circles are the single servers, and the oval is an infinite-server node. The ten activities are represented by the arrows between the buffers and servers. The numbers on the arrows indicate their index. Activities 1, 2, 3, and 4 are local dispatch activities while activities 5 through 10 are non-local dispatch activities. The arrows from the infinite-server to the buffers represent probabilistic rerouting of jobs in the network. . . . .	58
2.2	Manhattan area that is partitioned into four regions. . . . .	106
2.3	Unique solution $x^* \in \mathbb{R}^{10}$ to the static problem from (2.28)–(2.30). We see that Activities 6, 7, and 10 are nonbasic while the rest are basic. . . . .	108
2.4	Average cost with respect to varying holding cost. The shaded area along each line shows the 95% confidence interval based on 10 macro-replications. . . . .	113
2.5	Average cost with respect to varying idleness cost. The shaded area along each line shows the 95% confidence interval based on 10 macro-replications. . . . .	113
3.1	A process flow diagram of an online content moderation system featuring a machine learning classifier and human moderators. The dashed line symbolizes the feedback mechanism where human-generated labels refine the machine learning classifier’s accuracy over time via supervised learning. . . . .	136
3.2	A graph illustrating the relationship between learning improvement and system congestion in a system with three content classes and two virality levels, utilizing static policies that label content exclusively from one class. . . . .	147
3.3	A graph illustrating the relationship between learning improvement and system congestion in a system with three content classes and two virality levels, but with a faster learning improvement rate. . . . .	148
3.4	A graph illustrating the relationship between learning improvement and system congestion in a system with three content classes and two virality levels, utilizing static policies that label content exclusively from two classes. . . . .	149

## LIST OF TABLES

2.1	Estimated average cost along with the 95% confidence interval based on 10 macro-replications. . . . .	112
-----	---	-----

## ACKNOWLEDGMENTS

I would first like to acknowledge and thank the coauthors of the papers that comprise this dissertation. Chapter 1 is coauthored with Dr. Barış Ata of The University of Chicago Booth School of Business. Chapter 2 is coauthored with Dr. Barış Ata and Dr. Yuwei Zhou, also from The University of Chicago Booth School of Business. Finally, Chapter 3 is coauthored with Dr. René A. Caldentey, Dr. Amy R. Ward, and Tingrui Shi, all from of The University of Chicago Booth School of Business.

I would like to extend my deepest gratitude to my Committee Chair, Dr. Barış Ata, for his unwavering guidance and invaluable expertise throughout the course of my doctoral journey. Dr. Barış Ata's insightful feedback, rigorous academic standards, and encouragement were pivotal in navigating the complex challenges of my dissertation work. I am equally grateful to Dr. René Caldentey and Dr. Amy Ward for their guidance and support, particularly in the latter stages of my doctoral journey. My sincere thanks also go to Dr. Yuan Zhong and Dr. Yuwei Zhou for their support and for agreeing to serve on my dissertation committee.

Last but not least, I would like to thank the other people in my life that helped make this possible. I would like to thank my father, Dr. Layth C. Alwan, my mother, Athena M. Alwan, my sister, Vanessa A. Alwan, and my brother, Gabriel T. Alwan. This dissertation is dedicated to my paternal grandfather, Dr. Abdul Jabbar Alwan. He was the most wonderful grandfather anyone could wish for and I miss him. I am also grateful to the lifelong friends and classmates I had the fortune to meet during my doctoral studies. In alphabetical order, they are: Dr. T. Deniz Akturk, Dr. Nasser Barjesteh, Farbod Ekbatani, Dr. Zuguang Gao, Dr. R. Charlie Hannigan, Dr. Lisa A. Hillas, Dr. Çağla Keçeli, Dr. Robert J. Montgomery, Fabricio J. Previgliano, Dr. A. Cem Randa, Maximiliano Stock, Dr. Gizem Yilmaz, and Dr. Yanyang "Alex" Zhao. Finally, my heartfelt thanks go to my extended family members and close friends who have made a significant impact on my life. You know who you are.



## ABSTRACT

This dissertation presents work on the stochastic modeling and control of service systems, within the broader fields of management science and operations management. Chapters 1 and 2 focus on the ridesharing industry. They contain the contents of the papers Alwan and Ata [2020] and Alwan et al. [2024a], respectively. Chapter 3 studies an online content moderation system. It contains the contents of the working paper Alwan et al. [2024b].

In the first chapter, we develop a closed queueing network model of a ride-hailing system, where cars are conceptualized as jobs circulating through various nodes representing different city regions. By incorporating travel times between these nodes, our model achieves a high degree of generality, and allows for a detailed analysis of the dynamics and flow of cars within the system. We rigorously prove a novel heavy traffic limit theorem for this queueing network, providing an approximation for the original ride-hailing system that becomes increasingly accurate as the system approaches heavy traffic.

In the second chapter, we study a stochastic control problem based on the model presented in the first chapter, this time with a focus on a simplified structure featuring only a singular travel time node. To be more specific, we study a system where a system manager makes dynamic pricing and dispatch control decisions in a queueing network model motivated by ride-hailing applications. As in the first chapter, a novel feature of the model is that it incorporates travel times. The objective is to maximize the long-run average profit by making dynamic pricing and dispatch control decisions. Since this problem appears analytically intractable, we consider an approximating Brownian control problem in the heavy traffic regime. Under the assumptions of complete resource pooling and common travel time and routing distributions, we exploit an equivalence between the Brownian control problem and a one-dimensional workload formulation. We then solve the workload formulation in closed form by analyzing the corresponding Bellman equation. Using this solution, we propose a policy for the original queueing system and illustrate its effectiveness in a simulation study.

In the third chapter, we study the moderation of user-generated content by online platforms, with a particular focus on social media companies. The increased popularity of social media in recent years has led to an explosion of user-generated content. Although a substantial portion of this content is harmless, social media companies bear the crucial responsibility of protecting their users from harmful material. This has prompted the need for robust content moderation systems capable of accurate content classification. We propose a stochastic modeling framework for an online content moderation system that integrates a machine learning classifier with human moderators. Our focus is on determining which content should be reviewed by human moderators. On the one hand, manual review by human moderators improves the classification accuracy of content and better trains the machine learning classifier through supervised learning. On the other hand, sending content for human review results in delayed posting and removal decisions, and leaves the human reviewers at a higher risk for viewing disturbing content. Through a simulation study, we aim to explore the trade-off between learning and system congestion that arises when the platform sends multiple copies of a piece of content for human review. We also present a related stochastic model for online content moderation via a Bayesian learning framework. We study this model in an asymptotic regime where the volume of arriving content is high and where the informativeness of each human-assigned label is low. In that regime, we derive a diffusion approximation for a process that describes the platform's belief on the underlying violation status of the content in the system.

# CHAPTER 1

## A DIFFUSION APPROXIMATION FRAMEWORK FOR RIDE-HAILING WITH TRAVEL DELAYS

### 1.1 Introduction

This paper formulates and studies a closed queueing network motivated by ride-hailing systems such as Uber, Lyft, and Didi Chuxing. In these systems, customers request rides via a mobile application to travel from their current (pick-up) location to their final (drop-off) destination. Simultaneously, drivers wait for the mobile application to match them with the customer requests. Consequently, cars circulate through different regions of the city by picking up customers and delivering them to their desired destinations.

We consider a city partitioned into a finite set of geographical regions. Each such region constitutes a pick-up and a drop-off location. Crucially, our model incorporates the travel times between different regions of the city. We also allow customer heterogeneity by allowing customers in a region to have different destinations. Our model tracks the movements of cars between the various city regions as they pick up arriving customers and deliver them to their destinations. However, before a customer is dropped off, there is a time delay due to the travel times between the pick-up and drop-off locations. This is an important feature of our model. Figure 1.1 shows an illustrative example. Namely, it partitions the map of New York City into nine geographical regions. For this example, our model tracks the movements of cars between the nine regions.

To be specific, the queueing network model has single-server and infinite-server nodes. The single-server nodes correspond to different regions of the city, whereas the infinite-server nodes model the travel times between different regions. Cars move between different server nodes according to a probabilistic routing structure; see Section 1.3 for details.

The ultimate goal of a ride-hailing system is to maximize its profit. As such, it is important



Figure 1.1: A map of New York City partitioned into nine geographical regions.

for a ridesharing company to exert some control over the movement of cars on their platform. However, to effectively exert control requires first understanding how cars on their platform organically move throughout the city. This paper provides a first step in this important direction. In particular, we develop a diffusion approximation framework for ride-hailing systems with travel times. This approximation framework is justified via a weak convergence result for the queue length processes corresponding to the cars in the network; see Theorem 1.

The rest of the paper is organized as follows. Section 1.2 reviews the relevant literature. Section 1.3 formalizes our ridesharing model by introducing the model primitives and making a sample-path construction of the queue-length processes describing our ridesharing network. Section 1.4 puts forth our heavy traffic assumption and diffusion scaling regime, and states the main result of the paper (Theorem 1). Section 1.5 outlines the main tools needed to prove Theorem 1. Section 1.6 is devoted to a proof of Theorem 1, which involves proving convergence of the fluid-scaled processes. Section 1.7 concludes and discusses future research. Relevant notation and technical preliminaries are given in Appendix 1.8.1. Proofs of the results in Section 1.5, as well as all auxiliary results and derivations, are given in Appendices 1.8.2, 1.8.3, 1.8.4, and 1.8.5.

## 1.2 Literature Review

Our paper is related to two streams of literature. The first stream is on the modeling and analysis of ride-hailing networks, while the second stream is on the heavy traffic analysis of queueing network models.

The literature on ride-hailing has greatly expanded in recent years. A majority of the operational work on ride-hailing is on how pricing, matching, and car repositioning can impact system performance. The effect of pricing in ridesharing has been studied in Banerjee et al. [2016], Banerjee et al. [2022], Cachon et al. [2017], Besbes et al. [2021], Bimpikis et al. [2019], and Ata et al. [2020b], among others. Banerjee et al. [2016] studies dynamic pricing for a single-region system. They show that system performance under any dynamic pricing policy cannot exceed the performance under the optimal static pricing policy. Banerjee et al. [2022] designs pricing policies through a general approximation framework and show that the approximation ratio of the resulting pricing policy improves as the number of cars in each region grows. Despite the negative publicity of surge pricing, Cachon et al. [2017] demonstrate that surge pricing in a ridesharing platform can actually make both the platform and the customers better off. Besbes et al. [2021] also studies the problem of pricing in ridesharing systems with price sensitive customers and drivers.

Both Bimpikis et al. [2019] and Ata et al. [2020b] use spatial search models to study spatial pricing in ride-hailing networks with strategic drivers. Predating these papers, Lagos [2000] is one of the first papers to study spatial search frictions in the taxi industry. This paper highlights how search frictions endogenously arise as a result of the strategic movement of taxi drivers. Building on this paper, Bimpikis et al. [2019] studies spacial pricing of ridesharing systems analytically. They show that when demand patterns in the network are not “balanced,” spatial pricing can be very beneficial by increasing consumer surplus. On the other hand, Ata et al. [2020b] take an empirical approach to study how spatial

pricing and search frictions can impact the taxi market in New York City. To be specific, they use a mean field model to represent the taxis in the system and show how origin and origin-destination based pricing can reduce search frictions and increase consumer surplus. They do so empirically by performing a counterfactual analysis using data of taxi trips in New York City. Other related papers that study pricing in ridesharing include Chen and Sheldon [2016], Castillo et al. [2018], Lu et al. [2018], Hu and Zhou [2020], Gokpinar and Selcuk [2019], Garg and Nazerzadeh [2021], Hu et al. [2022], and Afèche et al. [2021].

Özkan and Ward [2020], Özkan [2020], and Banerjee et al. [2018] study matching (or dispatch control) in ridesharing platforms. In particular, Özkan and Ward [2020] models a ridesharing system as a non-stationary open queueing network where both drivers and customers exogenously arrive to the system. They propose a matching policy based on the solution to a continuous linear program and demonstrate that this policy is asymptotically optimal (in the fluid scale) in a large market regime where the arrival rates of drivers and customers are large. Özkan [2020] studies both the pricing and matching in ridesharing systems. The author demonstrates that joint pricing and matching provides significant performance increase over only optimizing over pricing or only over matching decisions. Finally, Banerjee et al. [2018] studies dispatch control in ridesharing systems. They model ridesharing systems as a closed queueing network model and proposes a family of state-dependent dispatch policies called Scaled MaxWeight policies. Under a complete resource pooling assumption, they show that the proportion of dropped demand under any Scaled MaxWeight policy decays exponentially. Other related papers include Wang et al. [2017], Lam and Liu [2017], Yan et al. [2020], Guda and Subramanian [2019b], Kanoria and Qian [2019], and Bertsimas et al. [2019].

Other work considers repositioning control, including that of Braverman et al. [2019], Afèche et al. [2023], and Ata et al. [2020a], among others. Braverman et al. [2019] models a ridesharing system as a closed queueing network and study it under fluid scaling. They

focus on “empty-car routing” where a car without a customer can be repositioned to another region of the city. This paper proves that the solution to a suitable linear program serves as an upper bound for the utility obtained under any state-dependent routing policy in the finite-car system. From a modeling perspective, Braverman et al. [2019] is closely related to ours because both explicitly model the travel times between city regions. However, an important difference is that our work allows for stochastic variability, while Braverman et al. [2019] studies a fluid-based model. On the other hand, Afèche et al. [2023] studies both demand-side admission control of customers as well as supply-side repositioning control of cars. In particular, they develop several insights on the interplay between centralized and de-centralized admission and repositioning control on the effect they have on the system. Contrary to the previous two papers, Ata et al. [2020a] considers a ride-hailing system with both repositioning and matching control. In particular, they model a ridesharing system as a stochastic processing network where the activities in the network correspond to repositioning cars from one area to another and dynamically matching customers with cars. By employing the general approach outlined in Harrison [2003], Ata et al. [2020a] proposes a solution to the original stochastic control problem by studying a the related Brownian control problem, which arises as a heavy traffic approximation to the original system. Other related papers include He et al. [2020] and Hosseini et al. [2021].

On the other hand, from a methodological perspective our paper involves the analysis of queueing network models, but more specifically on the asymptotic analysis of closed queueing networks containing infinite-server nodes. Related work includes that of Krichagina and Puhalskii [1997], Kogan and Lipster [1993], Kogan et al. [1986], and Smorodinskii [1986]. For example, Krichagina and Puhalskii [1997] studies a closed queueing model containing one single-server queue and one infinite-server queue (with a general service time distribution), and derive heavy traffic limits. On the other hand, Kogan and Lipster [1993] study a closed queueing network with many single-server queues and one infinite-server queue. In that

paper, all the single-server queues except one are in a “light-usage regime,” while the other single-server queue is studied in both a “moderate-usage regime” and a “heavy-usage regime.” They prove a state-space collapse result for the single-server queues in the light-usage regime, but prove diffusion approximation results for the single-server queue in the moderate-usage and heavy-usage regimes. Limit results for the queue length processes in closed queueing results are also established in the latter two papers Kogan et al. [1986] and Smorodinskii [1986]. Other papers that employ a similar type of analysis to ours include Ata and Kumar [2005], Ata and Lin [2008], Ata and Olsen [2009, 2013], and Reed and Ward [2008], among others.

To the best of our knowledge, this is the first paper to develop a diffusion approximation for a ride-hailing system while incorporating travel times between regions. Ata et al. [2020a] also proposes a diffusion approximation for ride-sharing systems, but does not incorporate travel times. Travel times are important from a practical perspective. However, with the exception of Braverman et al. [2019], most of the ride-hailing literature appears to ignore them effectively by assuming instantaneous pick-up and drop-off of customers. Our work builds on but differs from Braverman et al. [2019]. Namely, we model the uncertainty through the second moments of stochastic primitives, whereas Braverman et al. [2019] work with a deterministic model.

### 1.3 Ridesharing Model

The model has three key components: (i)  $J$  city regions where customers are picked up from and dropped off to, (ii)  $K$  travel times that correspond to the time that rides from one city region to another take, and (iii) a probabilistic routing structure that governs movement of cars. To be more specific, our model contains two levels of probabilistic routing. The first level is from the city regions to travel time nodes, while the second level is from travel time nodes to city regions. In the first level, a customer arriving to region  $j \in [J]$  requests a ride



requiring travel time  $k \in [K]$  with probability  $p_{jk}$ . In the second level, a car taking travel node  $k$  will, with probability  $q_{kj}$ , deliver the customer to region  $j$ . It should be noted that the arbitrary nature of the routing structure and the arbitrary number of travel nodes makes our model almost completely general. In particular, our model subsumes the  $K = J^2$  case, where an arriving customer to region  $j$  will get routed to region  $j' \in [J]$  with probability  $p_{jj'}$ . This occurs when the travel times between any two regions of the city are different. The two-level probabilistic routing with  $K$  travel times not only allows for arbitrarily general routing structures but also provides modeling flexibility. It also can help lower the dimension of the state space.

Each city region is modeled by a single-server queue, where a service completion corresponds to a customer arrival to the region who is subsequently picked up by a car. On the other hand, the  $K$  travel time nodes are modeled by infinite-server queues. A service completion at an infinite-server queue corresponds to a car finishing its travel with a customer from the pick-up location to the drop-off location.

In what follows, we consider a sequence of systems indexed by the total number of cars  $n \in \mathbb{N} := \{1, 2, \dots\}$ . Each system is a closed queueing network with  $J$  single-server queues,  $K$  infinite-server queues, and the probabilistic routing structure mentioned above. We study this sequence of systems in the heavy traffic asymptotic regime as  $n \rightarrow \infty$ . A superscript of  $n$  will be attached to various quantities of interest to indicate they correspond to the  $n$ th system.

### 1.3.1 *Model Primitives*

The service rate at a single-server queue reflects the demand rate at the corresponding city region. We consider a regime in which both the number of cars and the demand gets large.

Thus, the service rate  $\mu_j^n$  at the single-server queue  $j \in [J]$  in the  $n$ th system is given by

$$\mu_j^n := n\mu_j, \quad n \in \mathbb{N}, \quad (1.1)$$

where  $\mu_j > 0$  is a fixed parameter for  $j \in [J]$ . The service rates at the infinite-server queues do not vary with  $n$ . In particular, the service rate of infinite server  $k \in [K]$  is denoted by  $\eta_k > 0$ .

To facilitate the description of the system dynamics in Section 1.3.2, we define the following Poisson processes (one for each queue): For  $j \in [J]$  and  $k \in [K]$  we let

$$N_j = \{N_j(t) : t \geq 0\} \quad \text{and} \quad M_k = \{M_k(t) : t \geq 0\} \quad (1.2)$$

be independent rate-one Poisson processes.

To model the probabilistic routing of cars, we take as given the stochastic matrices

$$P = (p_{jk}) \in \mathbb{R}^{J \times K} \quad \text{and} \quad Q = (q_{kj}) \in \mathbb{R}^{K \times J} \quad (1.3)$$

representing routing from single-server queues to infinite-server queues and from infinite-server queues to single-server queues, respectively. To be more specific, for  $j \in [J]$  and  $k \in [K]$  let

$$\phi_j = \{\phi_j(l) : l \geq 1\} \quad \text{and} \quad \psi_k = \{\psi_k(l) : l \geq 1\} \quad (1.4)$$

denote independent sequences of i.i.d. random (routing) vectors. Their probability distributions are given by

$$p_{jk} := P(\phi_j(1) = e_k) \quad \text{and} \quad q_{kj} := P(\psi_k(1) = e_j), \quad (j, k) \in [J] \times [K], \quad (1.5)$$

where  $e_k$  and  $e_j$  are the  $k$ th and  $j$ th standard unit basis vector in  $\mathbb{R}^K$  and  $\mathbb{R}^J$ , respectively.

We assume that these random routing vectors are independent of all other stochastic model

primitives. The associated cumulative routing processes  $\Phi_{jk} = \{\Phi_{jk}(m) : m \geq 1\}$  and  $\Psi_{kj} = \{\Psi_{kj}(m) : m \geq 1\}$  are defined by the following:

$$\Phi_{jk}(m) := \sum_{l=1}^m \phi_{jk}(l) \quad \text{and} \quad \Psi_{kj}(m) := \sum_{l=1}^m \psi_{kj}(l), \quad m \in \mathbb{N}, \quad (1.6)$$

where  $\phi_{jk}(l)$  and  $\psi_{kj}(l)$  are the  $k$ th and  $j$ th components of  $\phi_j(l)$  and  $\psi_k(l)$ , respectively. In particular,  $\Phi_{jk}(m)$  represents the total number of customers that are routed from single-server queue  $j$  to infinite-server queue  $k$  among the first  $m$  customers arriving to region  $j$ . The expression  $\Psi_{kj}(m)$  can be interpreted similarly.

### 1.3.2 State Dynamics

For  $j \in [J]$  we denote by  $Q_j^n(t)$  the number of jobs in the  $j$ th single-server queue at time  $t$  in the  $n$ th system. Similarly, for  $k \in [K]$  we denote by  $V_k^n(t)$  the number of jobs in the  $k$ th infinite-server queue at time  $t$  in the  $n$ th system. Equations (1.8)–(1.9) below define these processes  $Q_j^n$  and  $V_k^n$  in the natural way. Also, we take as given random variables  $Q_j^n(0)$  and  $V_k^n(0)$  representing the initial number of cars in the network. We assume that the collection of these random variables is independent of all other stochastic primitives and that

$$\sum_{j=1}^J Q_j^n(0) + \sum_{k=1}^K V_k^n(0) = n, \quad \text{a.s.}, \quad (1.7)$$

as required for any closed system.

By conservation of flow, we define the system dynamics via the following equations:

$$Q_j^n(t) = Q_j^n(0) + A_j^n(t) - D_j^n(t), \quad j \in [J], \quad t \geq 0, \quad (1.8)$$

$$V_k^n(t) = V_k^n(0) + E_k^n(t) - F_k^n(t), \quad k \in [K], \quad t \geq 0, \quad (1.9)$$

where  $A_j^n = \{A_j^n(t) : t \geq 0\}$  and  $D_j^n = \{D_j^n(t) : t \geq 0\}$  denote the arrival and departure

processes for single-server queue  $j$ . Similarly,  $E_k^n = \{E_k^n(t) : t \geq 0\}$  and  $F_k^n = \{F_k^n(t) : t \geq 0\}$  denote the arrival and departure processes for infinite-server queue  $k$ . That is,  $A_j^n$  tracks the total number of cars that have arrived to region  $j$ , while  $D_j^n$  tracks the total number of cars that have departed from region  $j$  with a customer. We interpret the processes  $E_k^n$  and  $F_k^n$  similarly. The following equations define these processes:

$$A_j^n(t) := \sum_{k=1}^K \Psi_{kj}(F_k^n(t)), \quad t \geq 0, \quad (1.10)$$

$$E_k^n(t) := \sum_{j=1}^J \Phi_{jk}(D_j^n(t)), \quad t \geq 0, \quad (1.11)$$

$$D_j^n(t) := N_j(n\mu_j T_j^n(t)), \quad t \geq 0, \quad (1.12)$$

$$F_k^n(t) := M_k\left(\eta_k \int_0^t V_k^n(s) ds\right), \quad t \geq 0, \quad (1.13)$$

where  $\Phi_{jk}$  and  $\Psi_{kj}$  are the cumulative routing processes defined in (1.6) and

$$T_j^n = \{T_j^n(t) : t \geq 0\} \quad (1.14)$$

is the busy time process for the  $j$ th single-server queue. To be specific,  $T_j^n(t)$  is the cumulative amount of time the server is busy over  $[0, t]$  at single-server queue  $j \in [J]$ . The corresponding idleness process  $I_j^n = \{I_j^n(t) : t \geq 0\}$  at single-server queue  $j$  is defined by

$$I_j^n(t) := t - T_j^n(t), \quad t \geq 0. \quad (1.15)$$

By (1.7) and (1.10)–(1.13), it is straightforward to verify for  $t \geq 0$  that

$$\sum_{j=1}^J Q_j^n(t) + \sum_{k=1}^K V_k^n(t) = n, \quad \text{a.s.} \quad (1.16)$$

Finally, we assume that the following hold for each  $j \in [J]$ :

$$I_j^n \text{ is continuous and nondecreasing with } I_j^n(0) = 0, \quad (1.17)$$

$$\int_0^\infty \mathbb{1}_{\{Q_j^n(s) > 0\}} dI_j^n(s) = 0, \quad (1.18)$$

$$Q_j^n(t) \geq 0 \text{ for all } t \geq 0, \quad (1.19)$$

$$I_j^n(t) - I_j^n(s) \leq t - s \text{ for all } 0 \leq s \leq t. \quad (1.20)$$

Clearly, we must have  $Q_j(t) \geq 0$  for  $j \in [J]$ . Also, we restrict attention to work-conserving policies. That is, the server idleness does not increase as long as the queue is not empty. These are reflected in (1.18)–(1.19). Lastly, (1.17) and (1.20) are natural consequences of the interpretation of  $I_j^n$  as server idleness.

## 1.4 Heavy Traffic Limit Theorem

As a preliminary to stating our main result, we first introduce the heavy traffic assumption.

**Assumption 1** (Heavy Traffic Assumption). The following conditions hold. First,

$$\sum_{k=1}^K \sum_{i=1}^J \mu_i p_{ik} q_{kj} = \mu_j \quad \text{for } j \in [J]. \quad (1.21)$$

Second, letting  $m_k := \eta_k^{-1} \sum_{j=1}^J \mu_j p_{jk}$  for  $k \in [K]$ , we assume

$$\sum_{k=1}^K m_k = 1. \quad (1.22)$$

Roughly speaking, the first condition assumes that every single-server is fully utilized, whereas the second condition corresponds to assuming almost all cars are in transit.<sup>1</sup> To

---

1. These heavy traffic conditions can easily be generalized to limit conditions. To state it, suppose that  $\mu_j$  depends on  $n$ , for which we write  $\mu_j^n$ , and that  $\mu_j^n \rightarrow \mu_j \in \mathbb{R}$  as  $n \rightarrow \infty$  for each  $j \in [J]$ . Then (1.21) can

elaborate on the first condition, note that  $n\mu_i p_{ik}$  is the rate of jobs from single-server queue  $i$  to infinite-server queue  $k$ , i.e., the rate of customer arrivals into region  $i$  that require travel time  $k$ . The sum  $\sum_{i=1}^J n\mu_i p_{ik}$  therefore represents the total rate of jobs leaving the city regions that require travel time  $k$ . Then, the expression  $q_{kj} \sum_{i=1}^J n\mu_i p_{ik}$  represents the rate of cars entering region  $j$  from travel node  $k$ , and thus the entire sum  $\sum_{k=1}^K q_{kj} \sum_{i=1}^J n\mu_i p_{ik}$  represents the total rate of cars entering region  $j$  from travel nodes. Because  $n\mu_j$  is the customer arrival rate to region  $j$ , the first condition states the arrival rate of cars is balanced out by the customer demand in each region of the city.

We expect the first condition to hold in large cities under optimal spatial pricing provided the demand is elastic. To be more specific, if it was the case that for some region  $j$  that  $\sum_{k=1}^K \sum_{i=1}^J \mu_i p_{ik} q_{kj} \gg \mu_j$ , then the supply of cars in region  $j$  greatly exceeds customer demand. By decreasing the price for rides requested from region  $j$ , more people in region  $j$  will switch from their current form of transportation to ridesharing as their means for travel. This will increase the customer arrival rate  $\mu_j$  and would eventually balance out supply. It would also increase the profit because demand is elastic. On the other hand, if it was the case that for some region  $j$  that  $\sum_{k=1}^K \sum_{i=1}^J \mu_i p_{ik} q_{kj} \ll \mu_j$ , then customer demand greatly exceeds the supply of cars in region  $j$ . In this case, the ridesharing platform could increase prices for rides requested in region  $j$ , while maintaining the current prices in all other regions  $i \in [J] \setminus \{j\}$ .

To shed light on the second condition, note that the arrival rate to the  $k$ th infinite-server queue is  $n \sum_{j=1}^J \mu_j p_{jk}$  whereas its service rate is  $\eta_k$ . Based on intuition from the classical  $M/M/\infty$  queue, we expect the steady-state average queue length at the  $k$ th infinite-server to be  $n \sum_{j=1}^J \mu_j p_{jk} / \eta_k$ . Thus, the expected fraction of jobs at the  $k$ th infinite-server is  $m_k$ . So the condition  $\sum_{k=1}^K m_k = 1$  means (almost) all jobs are at the infinite-server queues, i.e., (almost) all cars are in transit. This is consistent with the first condition because of the

---

be replaced with the following condition:  $\sqrt{n} \left[ \sum_{k=1}^K \sum_{i=1}^J \mu_i^n p_{ik} q_{kj} - \mu_j^n \right] \rightarrow c_j \in \mathbb{R}$  for all  $j \in [J]$ . The main convergence result (Theorem 1) remains unchanged, except that we have an additional drift term  $c_j$ . A similar limit condition can be used in place of (1.22).

following: For each single-server queue, the arrival and service rates are equal and of order  $n$ . Thus, by the intuition from the central limit theorem, we expect the queue lengths to be of order  $\sqrt{n}$  for each single-server queue. Consequently, the total number of jobs in the single-server queues is of order  $\sqrt{n}$ , meaning (almost) all jobs are at the infinite-server queues because this is a closed network.

To facilitate the analysis to follow, we next define the following diffusion and fluid-scaled processes:

**Diffusion Scaled Processes:** For  $j \in [J]$  and  $k \in [K]$ , we define the following processes:

$$\widehat{Q}_j^n(t) := \frac{Q_j^n(t)}{\sqrt{n}}, \quad t \geq 0, \quad (1.23)$$

$$\widehat{V}_k^n(t) := \frac{V_k^n(t) - nm_k}{\sqrt{n}}, \quad t \geq 0, \quad (1.24)$$

$$\widehat{I}_j^n(t) := \sqrt{n}I_j^n(t), \quad t \geq 0, \quad (1.25)$$

$$\widehat{T}_j^n(t) := \sqrt{n}T_j^n(t), \quad t \geq 0, \quad (1.26)$$

$$\widehat{\Phi}_{jk}^n(t) := \frac{\Phi_{jk}(\lfloor nt \rfloor) - p_{jk}nt}{\sqrt{n}}, \quad t \geq 0, \quad (1.27)$$

$$\widehat{\Psi}_{kj}^n(t) := \frac{\Psi_{kj}(\lfloor nt \rfloor) - q_{kj}nt}{\sqrt{n}}, \quad t \geq 0, \quad (1.28)$$

$$\widehat{N}_j^n(t) := \frac{N_j(nt) - nt}{\sqrt{n}}, \quad t \geq 0, \quad (1.29)$$

$$\widehat{M}_k^n(t) := \frac{M_k(nt) - nt}{\sqrt{n}}, \quad t \geq 0. \quad (1.30)$$

**Fluid Scaled Processes:** For  $j \in [J]$  and  $k \in [K]$ , we define the following processes:

$$\bar{Q}_j^n(t) := \frac{\widehat{Q}_j^n(t)}{\sqrt{n}}, \quad t \geq 0, \quad (1.31)$$

$$\bar{V}_k^n(t) := \frac{\widehat{V}_k^n(t)}{\sqrt{n}}, \quad t \geq 0, \quad (1.32)$$

$$\bar{\bar{V}}_k^n(t) := \frac{V_k^n(t)}{n}, \quad t \geq 0, \quad (1.33)$$

$$\bar{N}_j^n(t) := \frac{N_j(nt)}{n}, \quad t \geq 0, \quad (1.34)$$

$$\bar{M}_k^n(t) := \frac{M_k(nt)}{n}, \quad t \geq 0. \quad (1.35)$$

By applying the scaling in (1.23)–(1.35) and using the heavy traffic conditions, it is straightforward to show that the following diffusion-scaled system equations hold for  $t \geq 0$ :

$$\begin{aligned} \widehat{Q}_j^n(t) &= \widehat{Q}_j^n(0) + \sum_{k=1}^K \widehat{\Psi}_{kj}^n \left( \bar{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) \right) - \widehat{N}_j^n \left( \mu_j T_j^n(t) \right) \\ &\quad + \sum_{k=1}^K q_{kj} \widehat{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) + \sum_{k=1}^K q_{kj} \eta_k \int_0^t \widehat{V}_k^n(s) ds + \mu_j \widehat{I}_j^n(t), \end{aligned} \quad (1.36)$$

$$\begin{aligned} \widehat{V}_k^n(t) &= \widehat{V}_k^n(0) + \sum_{j=1}^J \widehat{\Phi}_{jk}^n \left( \bar{N}_j^n \left( \mu_j T_j^n(t) \right) \right) - \widehat{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) \\ &\quad + \sum_{j=1}^J p_{jk} \widehat{N}_j^n \left( \mu_j T_j^n(t) \right) - \eta_k \int_0^t \widehat{V}_k^n(s) ds - \sum_{j=1}^J p_{jk} \mu_j \widehat{I}_j^n(t). \end{aligned} \quad (1.37)$$

See Appendix 1.8.5 for a straightforward, albeit tedious, derivation of the diffusion-scaled system equations (1.36) and (1.37). Defining

$$\begin{aligned} \widehat{\xi}_j^n(t) &:= \widehat{Q}_j^n(0) + \sum_{k=1}^K \widehat{\Psi}_{kj}^n \left( \bar{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) \right) - \widehat{N}_j^n \left( \mu_j T_j^n(t) \right) \\ &\quad + \sum_{k=1}^K q_{kj} \widehat{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right), \quad t \geq 0, \end{aligned} \quad (1.38)$$

$$\begin{aligned} \widehat{\zeta}_k^n(t) &:= \widehat{V}_k^n(0) + \sum_{j=1}^J \widehat{\Phi}_{jk}^n \left( \bar{N}_j^n \left( \mu_j T_j^n(t) \right) \right) - \widehat{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) \\ &\quad + \sum_{j=1}^J p_{jk} \widehat{N}_j^n \left( \mu_j T_j^n(t) \right), \quad t \geq 0, \end{aligned} \quad (1.39)$$

we see that (1.36)–(1.37) can be rewritten as follows:

$$\widehat{Q}_j^n(t) = \widehat{\xi}_j^n(t) + \sum_{k=1}^K q_{kj} \eta_k \int_0^t \widehat{V}_k^n(s) ds + \mu_j \widehat{I}_j^n(t), \quad t \geq 0, \quad (1.40)$$



$$\widehat{V}_k^n(t) = \widehat{\zeta}_k^n(t) - \eta_k \int_0^t \widehat{V}_k^n(s) ds - \sum_{j=1}^J p_{jk} \mu_j \widehat{I}_j^n(t), \quad t \geq 0. \quad (1.41)$$

Furthermore, by applying the fluid scaling in (1.31)–(1.32) to (1.40)–(1.41), we obtain the following fluid-scaled system equations:

$$\bar{Q}_j^n(t) = \bar{\xi}_j^n(t) + \sum_{k=1}^K q_{kj} \eta_k \int_0^t \bar{V}_k^n(s) ds + \mu_j I_j^n(t), \quad t \geq 0, \quad (1.42)$$

$$\bar{V}_k^n(t) = \bar{\zeta}_k^n(t) - \eta_k \int_0^t \bar{V}_k^n(s) ds - \sum_{j=1}^J p_{jk} \mu_j I_j^n(t), \quad t \geq 0, \quad (1.43)$$

where

$$\bar{\xi}_j^n(t) := n^{-1/2} \widehat{\xi}_j^n(t), \quad t \geq 0, \quad (1.44)$$

$$\bar{\zeta}_k^n(t) := n^{-1/2} \widehat{\zeta}_k^n(t), \quad t \geq 0. \quad (1.45)$$

We make the following assumptions on the initial conditions:

**Assumption 2** (Joint Convergence of the Initial Conditions). As  $n \rightarrow \infty$ ,

$$\begin{aligned} & (\widehat{Q}_1^n(0), \dots, \widehat{Q}_J^n(0), \widehat{V}_1^n(0), \dots, \widehat{V}_K^n(0)) \\ & \Rightarrow (Q_1(0), \dots, Q_J(0), V_1(0), \dots, V_K(0)) \in D^{J+K}. \end{aligned} \quad (1.46)$$

The result below establishes the joint convergence of the diffusion-scaled queue length processes and the idleness processes at the single-server stations to a multidimensional diffusion process. To facilitate the statement of the result, let  $(\xi^*, \zeta^*) = (\xi_1^*, \dots, \xi_J^*, \zeta_1^*, \dots, \zeta_K^*)$  be a  $(0, \Sigma)$  Brownian motion with initial state  $(\xi^*(0), \zeta^*(0)) = (Q(0), V(0))$ , where  $(Q(0), V(0)) = (Q_1(0), \dots, Q_J(0), V_1(0), \dots, V_K(0))$ . By construction and (1.46), observe that  $\xi^*(0) \geq 0$

and  $\sum_{j=1}^J \xi_j^*(0) + \sum_{k=1}^K \zeta_k^*(0) = 0$ . The covariance matrix  $\Sigma \in \mathbb{R}^{(J+K) \times (J+K)}$  is given by

$$\Sigma_{j,j} = \sum_{k=1}^K q_{kj}(1 - q_{kj})\eta_k m_k + \mu_j + \sum_{k=1}^K q_{kj}^2 \eta_k m_k, \quad j \in [J], \quad (1.47)$$

$$\Sigma_{J+k, J+k} = \sum_{j=1}^J p_{jk}(1 - p_{jk})\mu_j + \eta_k m_k + \sum_{j=1}^J p_{jk}^2 \mu_j, \quad k \in [K], \quad (1.48)$$

$$\Sigma_{i,j} = \sum_{k=1}^K q_{ki} q_{kj} \eta_k m_k, \quad i, j \in [J], i \neq j, \quad (1.49)$$

$$\Sigma_{j, J+k} = -p_{jk} \mu_j - q_{kj} \eta_k m_k, \quad j \in [J], k \in [K], \quad (1.50)$$

$$\Sigma_{J+l, J+k} = \sum_{j=1}^J p_{jl} p_{jk} \mu_j, \quad l, k \in [K], l \neq k. \quad (1.51)$$

**Theorem 1** (Joint Convergence of Diffusion Scaled Processes). *We have that as  $n \rightarrow \infty$ ,*

$$(\hat{Q}^n, \hat{I}^n, \hat{V}^n) \Rightarrow (Q^*, I^*, V^*) \quad \text{in } D^{2J+K}, \quad (1.52)$$

where  $Q^*$ ,  $I^*$ , and  $V^*$  are multidimensional diffusion processes satisfying the following for all  $j \in [J]$ ,  $k \in [K]$ , and  $t \geq 0$ :

$$Q_j^*(t) = \xi_j^*(t) + \sum_{k=1}^K q_{kj} \eta_k \int_0^t V_k^*(s) ds + \mu_j I_j^*(t) \geq 0, \quad (1.53)$$

$$V_k^*(t) = \zeta_k^*(t) - \eta_k \int_0^t V_k^*(s) ds - \sum_{j=1}^J p_{jk} \mu_j I_j^*(t), \quad (1.54)$$

$$\int_0^\infty \mathbf{1}_{\{Q_j^*(t) > 0\}} dI_j^*(t) = 0. \quad (1.55)$$

## 1.5 Auxiliary Results

This section establishes the existence of suitably defined continuous functions that will aid in the proof of Theorem 1 via a continuous mapping argument; see Appendix 1.8.2 for the proofs of the results in this section. To that end, fix  $\xi = (\xi_1, \dots, \xi_J) \in D^J$  and  $\zeta = (\zeta_1, \dots, \zeta_K) \in D^K$

such that

$$\sum_{j=1}^J \xi_j(t) + \sum_{k=1}^K \zeta_k(t) = 0, \quad \text{for all } t \geq 0, \quad (1.56)$$

$$\xi_j(0) \geq 0, \quad \text{for all } j \in [J]. \quad (1.57)$$

The following collection of equations corresponds to our closed ridesharing network with  $J$  single-server stations and  $K$  infinite-server stations. That is, for  $j \in [J]$ ,  $k \in [K]$ , and  $t \geq 0$  we consider the following set of equations:

$$x_j(t) = \xi_j(t) + \sum_{k=1}^K q_{kj} \eta_k \int_0^t y_k(s) ds + \mu_j u_j(t) \in [0, \infty), \quad (1.58)$$

$$y_k(t) = \zeta_k(t) - \eta_k \int_0^t y_k(s) ds - \sum_{j=1}^J p_{jk} \mu_j u_j(t), \quad (1.59)$$

$$\sum_{j=1}^J x_j(t) + \sum_{k=1}^K y_k(t) = 0, \quad (1.60)$$

$$u_j \text{ is nondecreasing with } u_j(0) = 0, \quad (1.61)$$

$$\int_0^\infty \mathbb{1}_{\{x_j(t) > 0\}} du_j(t) = 0. \quad (1.62)$$

The next result establishes the existence and uniqueness of solutions to these equations.

**Proposition 1.** *For every  $(\xi, \zeta) \in D^{J+K}$  satisfying (1.56)–(1.57), there exists a unique  $(x, u, y) \in D^{2J+K}$  satisfying (1.58)–(1.62).*

The following result is now immediate from Proposition 1.

**Corollary 1.** *There exists a unique triple  $(f, g, h) : D^{J+K} \rightarrow D^{2J+K}$  such that whenever  $(\xi, \zeta) \in D^{J+K}$  satisfies (1.56)–(1.57),  $(f(\xi, \zeta), g(\xi, \zeta), h(\xi, \zeta))$  satisfies (1.58)–(1.62).*

See Appendix 1.8.2 for a characterization of these functions  $f$ ,  $g$ , and  $h$ . The final result of this section establishes the continuity of the mappings defined in Corollary 1.

**Proposition 2.** *The mapping  $(f, g, h) : D^{J+K} \rightarrow D^{2J+K}$  from Corollary 1 is continuous when both the domain and range spaces are endowed with the Skorokhod  $J_1$  topology.*

## 1.6 Proof of the Main Theorem

This section contains the main convergence results of this paper, culminating with a proof of Theorem 1. To begin, Section 1.6.1 proves convergence of the fluid-scaled processes. These results are necessary because several of the fluid-scaled processes serve as random time changes in the diffusion-scaled equations. Then, in Section 1.6.2, convergence of the process  $(\hat{\xi}^n, \hat{\zeta}^n)$  given by (1.38)–(1.39) is proven. This, combined with a continuous mapping argument, allows us to complete the proof of Theorem 1.

### 1.6.1 Convergence of Fluid Scaled Processes

We begin by establishing convergence of the fluid-scaled processes.

**Lemma 1.** *As  $n \rightarrow \infty$ ,  $(\bar{\xi}^n, \bar{\zeta}^n) \Rightarrow \mathbf{0} \in D^{J+K}$ .*

*Proof.* To prove that  $(\bar{\xi}^n, \bar{\zeta}^n) \Rightarrow \mathbf{0}$ , it suffices to show [by, for example, Theorem 11.4.5 in Whitt [2002]] that  $\bar{\xi}_j^n \Rightarrow 0$  and  $\bar{\zeta}_k^n \Rightarrow 0$  for all  $j \in [J]$  and  $k \in [K]$ . On the other hand, to prove that  $\bar{\xi}_j^n \Rightarrow 0$  and  $\bar{\zeta}_k^n \Rightarrow 0$ , it is sufficient to show that for all  $T > 0$ ,

$$\|\bar{\xi}_j^n\|_T \Rightarrow 0 \quad \text{and} \quad \|\bar{\zeta}_k^n\|_T \Rightarrow 0 \tag{1.63}$$

as random variables; see Lemma 5 in Appendix 1.8.4 for a proof of this claim. By (1.38)–(1.39), the triangle inequality, and the facts that  $\int_0^t \bar{V}_k^n(s) ds \leq t$  and  $T_j^n(t) \leq t$  for all  $t \geq 0$ , one readily checks that for all  $T > 0$ ,

$$\|\widehat{\xi}_j^n\|_T \leq \|\widehat{Q}_j^n(0)\|_T + \sum_{k=1}^K \|\widehat{\Psi}_{kj}^n(\bar{M}_k^n(\eta_k \cdot))\|_T + \|\widehat{N}_j^n(\mu_j \cdot)\|_T + \sum_{k=1}^K \|\widehat{M}_k^n(\eta_k \cdot)\|_T, \tag{1.64}$$

$$\|\widehat{\zeta}_k^n\|_T \leq \|\widehat{V}_k^n(0)\|_T + \sum_{j=1}^J \|\widehat{\Phi}_{jk}^n(\bar{N}_j^n(\mu_j \cdot))\|_T + \|\widehat{M}_k^n(\eta_k \cdot)\|_T + \sum_{j=1}^J \|\widehat{N}_j^n(\mu_j \cdot)\|_T. \quad (1.65)$$

Then an application of Donsker's theorem, the functional central limit theorem for renewal processes, the random time change theorem, and the continuous mapping theorem can be used to show that the right-hand sides of (1.64) and (1.65) converge weakly to a nondegenerate limit. By this and the fact that  $\bar{\xi}_j^n = n^{-1/2}\widehat{\xi}_j^n$  and  $\bar{\zeta}_k^n = n^{-1/2}\widehat{\zeta}_k^n$ , we obtain (1.63). We refer the reader to Appendix 1.8.4 for additional details.  $\square$

**Lemma 2.** *As  $n \rightarrow \infty$ ,  $(\bar{Q}^n, I^n, \bar{V}^n) \Rightarrow \mathbf{0} \in D^{J+K}$*

*Proof.* Recall from (1.42)–(1.43) that we have

$$\bar{Q}_j^n(t) = \bar{\xi}_j^n(t) + \sum_{k=1}^K q_{kj} \eta_k \int_0^t \bar{V}_k^n(s) ds + \mu_j I_j^n(t), \quad t \geq 0, \quad (1.66)$$

$$\bar{V}_k^n(t) = \bar{\zeta}_k^n(t) - \eta_k \int_0^t \bar{V}_k^n(s) ds - \sum_{j=1}^J p_{jk} \mu_j I_j^n(t), \quad t \geq 0. \quad (1.67)$$

Note that the process  $(\bar{\xi}^n, \bar{\zeta}^n)$  satisfies

$$\xi_j^n(0) = \bar{Q}_j^n(0) \geq 0, \quad \text{a.s. for all } j \in [J]. \quad (1.68)$$

Recall by (1.17) that the process  $I_j^n$  is nondecreasing with  $I_j^n(0) = 0$ . Furthermore, by (1.18),

$$\int_0^\infty \mathbb{1}_{\{\bar{Q}_j^n(t) > 0\}} dI_j^n(t) = \int_0^\infty \mathbb{1}_{\{n^{-1}Q_j^n(t) > 0\}} dI_j^n(t) = \int_0^\infty \mathbb{1}_{\{Q_j^n(t) > 0\}} dI_j^n(t) = 0. \quad (1.69)$$

By the uniqueness in Proposition 1 and (1.66)–(1.69), it follows that

$$(\bar{Q}^n, I^n, \bar{V}^n) = (f(\bar{\xi}^n, \bar{\zeta}^n), g(\bar{\xi}^n, \bar{\zeta}^n), h(\bar{\xi}^n, \bar{\zeta}^n)). \quad (1.70)$$

Then, by Proposition 2, Lemma 1, the continuous mapping theorem, and (1.70), we have that

$$(\bar{Q}^n, I^n, \bar{V}^n) = (f(\bar{\xi}^n, \bar{\zeta}^n), g(\bar{\xi}^n, \bar{\zeta}^n), h(\bar{\xi}^n, \bar{\zeta}^n)) \Rightarrow (f(\mathbf{0}), g(\mathbf{0}), h(\mathbf{0})). \quad (1.71)$$

To complete the proof, we must show that  $f(\mathbf{0}) = \mathbf{0}$ ,  $g(\mathbf{0}) = \mathbf{0}$ , and  $h(\mathbf{0}) = \mathbf{0}$ . If we can show that  $\bar{V} := h(\mathbf{0}) = \mathbf{0}$ , then by definition of  $\bar{Q} := f(\mathbf{0})$  and  $I := g(\mathbf{0})$ , we would have

$$\begin{aligned} \bar{Q}_j &= \phi \left( \pi_j \circ \mathbf{0} + \sum_{k=1}^K q_{kj} \eta_k \int_0^\cdot \bar{V}_k(s) ds \right) = \phi \left( 0 + \sum_{k=1}^K q_{kj} \eta_k \int_0^\cdot 0 ds \right) = 0, \\ I_j &= \mu_j^{-1} \psi \left( \pi_j \circ \mathbf{0} + \sum_{k=1}^K q_{kj} \eta_k \int_0^\cdot \bar{V}_k(s) ds \right) = \mu_j^{-1} \psi \left( 0 + \sum_{k=1}^K q_{kj} \eta_k \int_0^\cdot 0 ds \right) = 0, \end{aligned}$$

and the proof would be complete. By definition of  $\bar{V} := h(\mathbf{0})$ , we have

$$\bar{V}_k(t) = -\eta_k \int_0^t \bar{V}_k(s) ds - \sum_{j=1}^J p_{jk} \psi \left( \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot \bar{V}_l(s) ds \right) (t), \quad t \geq 0, \quad (1.72)$$

for all  $k = 1, \dots, K$ . But, for a fixed  $T > 0$  and any  $0 \leq t \leq T$ , we have

$$\begin{aligned} \|\bar{V}_k\|_t &\leq \eta_k \left\| \int_0^\cdot \bar{V}_k(s) ds \right\|_t + \sum_{j=1}^J \left\| \psi \left( \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot \bar{V}_l(s) ds \right) \right\|_t \\ &\leq \eta \int_0^t \|\bar{V}_k\|_s ds + \sum_{j=1}^J \sum_{l=1}^K \eta_l \left\| \int_0^\cdot \bar{V}_l(s) ds \right\|_t \\ &\leq \eta \int_0^t \max_{1 \leq k \leq K} \|\bar{V}_k\|_s ds + \eta JK \int_0^t \max_{1 \leq k \leq K} \|\bar{V}_k\|_s ds \\ &\leq 2\eta JK \int_0^t \max_{1 \leq k \leq K} \|\bar{V}_k\|_s ds, \end{aligned} \quad (1.73)$$

where  $\eta := \max_{1 \leq k \leq K} \eta_k$ . It then follows from (1.73) that

$$\max_{1 \leq k \leq K} \|\bar{V}_k\|_t \leq 2\eta JK \int_0^t \max_{1 \leq k \leq K} \|\bar{V}_k\|_s ds, \quad t \geq 0. \quad (1.74)$$

Therefore, by Gronwall's inequality (see, e.g., Lemma 4.1 in Pang et al. [2007]) and (1.74), it

follows that

$$\max_{1 \leq k \leq K} \|\bar{V}_k\|_T = 0. \quad (1.75)$$

Since  $T$  was arbitrary, (1.75) implies that  $\bar{V} \equiv 0$ . The proof is complete.  $\square$

**Corollary 2.** *As  $n \rightarrow \infty$ ,  $T^n \Rightarrow e \in D^J$ , where  $e(t) = (t, \dots, t)$  for all  $t \geq 0$ .*

*Proof.* By definition,  $T^n = e - I^n$ . The result then follows from Lemma 2 since  $I^n \Rightarrow \mathbf{0}$ .  $\square$

### 1.6.2 Convergence of Diffusion Scaled Processes

**Lemma 3.** *As  $n \rightarrow \infty$ ,  $(\widehat{\xi}^n, \widehat{\zeta}^n) \Rightarrow (\xi^*, \zeta^*) \in D^{J+K}$ , where  $(\xi^*, \zeta^*)$  is a  $(0, \Sigma)$  Brownian motion with initial state  $(Q(0), V(0))$  and covariance matrix  $\Sigma$  given by (1.47)–(1.51).*

*Proof.* Let  $e : [0, \infty) \rightarrow [0, \infty)$  denote the one-dimensional identity map  $e(t) = t$  for  $t \geq 0$ .

Furthermore, recall from (1.38)–(1.39) that we have

$$\begin{aligned} \widehat{\xi}_j^n(t) &:= \widehat{Q}_j^n(0) + \sum_{k=1}^K \widehat{\Psi}_{kj}^n \left( \bar{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) \right) - \widehat{N}_j^n \left( \mu_j T_j^n(t) \right) \\ &\quad + \sum_{k=1}^K q_{kj} \bar{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right), \\ \widehat{\zeta}_k^n(t) &:= \widehat{V}_k^n(0) + \sum_{j=1}^J \widehat{\Phi}_{jk}^n \left( \bar{N}_j^n \left( \mu_j T_j^n(t) \right) \right) - \bar{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) + \sum_{j=1}^J p_{jk} \widehat{N}_j^n \left( \mu_j T_j^n(t) \right). \end{aligned}$$

By Lemma 2, Corollary 2, and the derivations in Appendix 1.8.5, we have the following convergence for the fluid-scaled processes:

$$\bar{M}_k^n(\eta_k \cdot) \Rightarrow \eta_k e, \quad \bar{N}_j^n(\mu_j \cdot) \Rightarrow \mu_j e, \quad T_j^n \Rightarrow e, \quad \text{and} \quad \bar{V}_k^n \Rightarrow m_k.$$

We also have the following convergence for the diffusion-scaled processes:

$$\widehat{\Psi}_{kj}^n \Rightarrow \sqrt{q_{kj}(1 - q_{kj})} B_{kj}, \quad \widehat{\Phi}_{kj}^n \Rightarrow \sqrt{p_{jk}(1 - p_{jk})} B_{jk},$$

$$\widehat{M}_k^n(\eta_k \cdot) \Rightarrow \eta_k^{1/2} B_k, \quad \widehat{N}_j^n(\mu_j \cdot) \Rightarrow \mu_j^{1/2} B_j,$$

where  $B_{kj}$ ,  $B_{jk}$ ,  $B_k$ , and  $B_j$  are independent standard Brownian motions. Moreover, the function  $H : D \rightarrow D$  defined by

$$H(x)(t) = \int_0^t x(s) ds, \quad t \geq 0,$$

is continuous in the Skorokhod topology (see, e.g., page 229 in Pang et al. [2007]). Therefore, it follows that

$$H(\bar{V}_k^n) \Rightarrow H(m_k) = m_k e.$$

By the above convergence results and (1.46), Theorems 11.4.4 and 11.4.5 in Whitt [2002], and independence of all stochastic primitives, it follows that the (joint) processes

$$\begin{aligned} & \left( \widehat{Q}_1^n(0), \dots, \widehat{Q}_J^n(0), \widehat{V}_1^n(0), \dots, \widehat{V}_K^n(0), \widehat{\Psi}_{11}^n, \dots, \widehat{\Psi}_{1J}^n, \dots, \widehat{\Psi}_{K1}^n, \dots, \widehat{\Psi}_{KJ}^n, \right. \\ & \left. \widehat{\Phi}_{11}^n, \dots, \widehat{\Phi}_{1K}^n, \dots, \widehat{\Phi}_{J1}^n, \dots, \widehat{\Phi}_{JK}^n, \widehat{N}_1^n, \dots, \widehat{N}_J^n, \widehat{M}_1^n, \dots, \widehat{M}_K^n \right) \end{aligned} \quad (1.76)$$

and

$$(T_1^n, \dots, T_J^n, \bar{V}_1^n, \dots, \bar{V}_K^n, \bar{N}_1^n, \dots, \bar{N}_J^n, \bar{M}_1^n, \dots, \bar{M}_K^n) \quad (1.77)$$

converge weakly to their appropriate limits. By the convergence of (1.76)–(1.77), the random time change theorem, and the continuous mapping theorem, we get following convergence:

$$(\widehat{\xi}_1^n, \dots, \widehat{\xi}_J^n, \widehat{\zeta}_1^n, \dots, \widehat{\zeta}_K^n) \Rightarrow (\xi_1^*, \dots, \xi_J^*, \zeta_1^*, \dots, \zeta_K^*). \quad (1.78)$$

It is straightforward, albeit tedious, to derive the covariance matrix  $\Sigma$ , so we omit those details here. □



Finally, we conclude with a proof of Theorem 1.

*Proof of Theorem 1.* It is straightforward to show that the process  $(\widehat{\xi}^n, \widehat{\zeta}^n)$  satisfies

$$\widehat{\xi}_j^n(0) = \widehat{Q}_j^n(0) \geq 0, \quad \text{a.s. for all } j \in [J], \quad (1.79)$$

$$\sum_{j=1}^J \widehat{\xi}_j^n(t) + \sum_{k=1}^K \widehat{\zeta}_k^n(t) = 0, \quad \text{a.s. for all } t \geq 0. \quad (1.80)$$

Similar to the proof of Lemma 2, one can show that  $\widehat{I}_j^n$  is nondecreasing,  $\widehat{I}_j^n(0) = 0$ , and

$$\int_0^\infty \mathbb{1}_{\{\widehat{Q}_j^n(t) > 0\}} d\widehat{I}_j^n(t) = 0.$$

This, combined with the uniqueness in Proposition 1, implies that

$$(\widehat{Q}^n, \widehat{I}^n, \widehat{V}^n) = (f(\widehat{\xi}^n, \widehat{\zeta}^n), g(\widehat{\xi}^n, \widehat{\zeta}^n), h(\widehat{\xi}^n, \widehat{\zeta}^n)). \quad (1.81)$$

By Lemma 3, Proposition 2, and the continuous mapping theorem, we then have that

$$\begin{aligned} (\widehat{Q}^n, \widehat{I}^n, \widehat{V}^n) &= (f(\widehat{\xi}^n, \widehat{\zeta}^n), g(\widehat{\xi}^n, \widehat{\zeta}^n), h(\widehat{\xi}^n, \widehat{\zeta}^n)) \\ &\Rightarrow (f(\xi^*, \zeta^*), g(\xi^*, \zeta^*), h(\xi^*, \zeta^*)). \end{aligned}$$

Since inequalities are preserved under weak convergence, (1.46) and (1.79) imply that

$$\xi_j^*(0) = Q_j(0) \geq 0, \quad \text{a.s. for all } j \in [J]. \quad (1.82)$$

Furthermore, by Lemma 3 and (1.80) we have that

$$\sum_{j=1}^J \xi_j^*(t) + \sum_{k=1}^K \zeta_k^*(t) = 0, \quad \text{a.s. for all } t \geq 0. \quad (1.83)$$

Letting  $(Q^*, I^*, V^*) := (f(\xi^*, \zeta^*), g(\xi^*, \zeta^*), h(\xi^*, \zeta^*))$ , the proof is complete by (1.82)–(1.83) and Proposition 1. Finally, the distributional description of the Brownian motion process  $(\xi^*, \zeta^*)$  comes from Lemma 3.  $\square$

## 1.7 Concluding Remarks

This paper proposes a closed queueing system to model the movement of cars in a ride-hailing network. Under the assumption that the supply of cars and customer demand is perfectly balanced, our results show that the distribution of cars in the network can be approximated by a diffusion process. Crucially, this paper incorporates travel times into the ridesharing model. Modeling travel times is important because, as is often the case in large cities, drivers spend a non-trivial amount of time on the road delivering customers to their destinations. Ignoring these travel times can lead to inaccuracies when tracking cars in a city.

The results in this paper effectively assume that the cars in the ride-hailing platform are self-driving because we do not model strategic driver behavior. Therefore, a worthwhile extension to this work would be to incorporate strategic behavior into the model because, in many settings, drivers are autonomous and forward looking.

## 1.8 Appendix

### 1.8.1 Notation and Terminology

For a function  $f : X \rightarrow Y$  and a subset  $S \subseteq X$ , we denote by  $f|_S : S \rightarrow Y$  the restriction of  $f$  to  $S$ . The indicator function for a subset  $S \subseteq X$  is written as  $\mathbb{1}_S$ . For  $n \in \mathbb{N}$ , we use the notation  $[n] := \{1, 2, \dots, n\}$ . For  $k \in [n]$ , the  $k$ th unit basis vector in  $\mathbb{R}^n$  is denoted by  $e_k$ , which has one in the  $k$ th component and zeros elsewhere. Moreover, for  $k \in [n]$ , the  $k$ th projection map  $\pi_k : \mathbb{R}^n \rightarrow \mathbb{R}$  is given by  $\pi_k(x) = x_k$ , where  $x_k$  is the  $k$ th component of  $x \in \mathbb{R}^n$ . For  $a, b \in \mathbb{R}$ , we use the notation  $a \vee b := \max\{a, b\}$  and  $a \wedge b := \min\{a, b\}$ , and let  $\lfloor a \rfloor$  denote

the largest integer less than or equal to  $a$ . The abbreviation a.s. stands for “almost surely” and the notation  $\xrightarrow{p}$  means “converges in probability.” Finally, for two measures  $\mu$  and  $\nu$  on the same measurable space, we write  $\mu \ll \nu$  to mean that  $\mu$  is absolutely continuous with respect to  $\nu$ .

For each  $n \in \mathbb{N}$ , we denote by  $D^n \equiv D([0, \infty), \mathbb{R}^n)$  the set of all functions  $x : [0, \infty) \rightarrow \mathbb{R}^n$  that are right continuous on  $[0, \infty)$  and have left limits on  $(0, \infty)$ . We denote by  $\mathbf{0} : [0, \infty) \rightarrow \mathbb{R}^n$  the function that is identically zero. For each  $k \in \mathbb{N}$  and  $T > 0$ , we denote by  $D_T^k \equiv D([0, T], \mathbb{R}^k)$  the set of all functions  $x : [0, T] \rightarrow \mathbb{R}^k$  that are right continuous on  $[0, T]$  and have left limits on  $(0, T]$ . When the space  $D_T^k$  is endowed with the norm

$$\|x\|_{T,k} := \max_{1 \leq l \leq k} \sup_{0 \leq t \leq T} |x_l(t)|, \quad (1.84)$$

it is a Banach space; see, e.g., Pestman [1995]. When  $k = 1$ , we write  $D^1 = D$ ,  $D_T^1 = D_T$ , and  $\|\cdot\|_{T,1} = \|\cdot\|_T$ . The one-sided reflection mapping on  $D$  is given by the pair of functions  $(\psi, \phi) : D \rightarrow D^2$  defined as follows:

$$\psi(x)(t) := \sup_{0 \leq s \leq t} [-x(s)]^+, \quad x \in D, \quad t \geq 0, \quad (1.85)$$

$$\phi(x)(t) := x(t) + \psi(x)(t), \quad x \in D, \quad t \geq 0. \quad (1.86)$$

For  $x, y \in D$  and  $T > 0$  the following inequalities hold (see Lemma 13.5.1 in Whitt [2002]):

$$\|\psi(x) - \psi(y)\|_T \leq \|x - y\|_T, \quad (1.87)$$

$$\|\phi(x) - \phi(y)\|_T \leq 2\|x - y\|_T. \quad (1.88)$$

We regard  $D^k$  as a topological space with usual Skorokhod  $J_1$  topology (see, e.g., Billingsley [1999]). The precise definition of the underlying topology is not important for our purposes. However, for proofs in this paper, the reader must know what it means for a sequence of

functions to converge in  $D^k$ . To that end, a sequence  $\{x^n\}_{n=1}^\infty$  in  $D^k$  converges to an element  $x \in D^k$ , written  $x^n \rightarrow x$  as  $n \rightarrow \infty$ , if for all continuity points  $T > 0$  of  $x$  the following holds:

$$d_T^k(x^n|_{[0,T]}, x|_{[0,T]}) \rightarrow 0 \quad \text{as } n \rightarrow \infty, \quad (1.89)$$

where  $d_T^k : D_T^k \times D_T^k \rightarrow [0, \infty)$  is given by

$$d_T^k(x, y) := \inf_{\lambda \in \Lambda_T} \{ \|x \circ \lambda - y\|_{T,k} \vee \|\lambda - e\|_T \}, \quad (1.90)$$

where  $e : [0, T] \rightarrow \mathbb{R}$  is the identity map, i.e.,  $e(t) = t$  for all  $t \in [0, T]$ , and

$$\Lambda_T := \{ \lambda : [0, T] \rightarrow [0, T] \mid \lambda \text{ is an increasing homeomorphism} \}. \quad (1.91)$$

Equivalently, (1.89) can be written with  $\tilde{d}_T^k$  in place of  $d_T^k$  (see, e.g., page 226 of Pang et al. [2007] and Billingsley [1999]), where  $\tilde{d}_T^k : D_T^k \times D_T^k \rightarrow [0, \infty)$  is given by

$$\tilde{d}_T^k(x, y) := \inf_{\lambda \in \tilde{\Lambda}_T} \{ \|x \circ \lambda - y\|_{T,k} \vee \|\dot{\lambda} - \mathbf{1}\|_T \}, \quad (1.92)$$

where  $\mathbf{1} : [0, T] \rightarrow \mathbb{R}^k$  is the function that takes one everywhere,  $\dot{\lambda}$  is the derivative of  $\lambda$ , and

$$\tilde{\Lambda}_T := \{ \lambda : [0, T] \rightarrow [0, T] \mid \lambda \in \Lambda_T \text{ and } \lambda \ll \text{Lebesgue measure} \}. \quad (1.93)$$

As before, when  $k = 1$ , we write  $d_T^1 = d_T$  and  $\tilde{d}_T^1 = \tilde{d}_T$ .

All random variables in this paper are assumed to live on a common probability space  $(\Omega, \mathcal{F}, P)$ . We denote by  $\mathcal{M}^k$  the Borel  $\sigma$ -algebra on  $D^k$  induced by the Skorokhod  $J_1$  topology, so that  $(D^k, \mathcal{M}^k)$  forms a measurable space. Each stochastic process in this paper is assumed to be a measurable function from  $(\Omega, \mathcal{F}, P)$  to  $(D^k, \mathcal{M}^k)$ , with appropriate dimension  $k$ . For a sequence of stochastic processes  $\{\xi^n\}_{n=1}^\infty$  in  $D^k$ , where  $\xi^n = \{\xi^n(t) : t \geq 0\}$ ,

we write  $\xi^n \Rightarrow \xi$  as  $n \rightarrow \infty$  to mean that the sequence of probability measures on  $(D^k, \mathcal{M}^k)$  induced by the  $\xi^n$  converge weakly to the probability measure on  $(D^k, \mathcal{M}^k)$  induced by the stochastic process  $\xi$  (see, e.g., Billingsley [1999] and Whitt [2002] for further details).

### 1.8.2 Proofs for Section 1.5

The following lemma is useful in the proof of Proposition 1 and is proven in Appendix 1.8.3.

To state it, given  $(\xi, \zeta) \in D^{J+K}$ , for  $k = 1, \dots, K$  consider the following equation:

$$y_k(t) = \zeta_k(t) - \eta_k \int_0^t y_k(s) ds - \sum_{j=1}^J p_{jk} \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) ds \right)(t), \quad t \geq 0. \quad (1.94)$$

**Lemma 4.** *For each  $(\xi, \zeta) \in D^{J+K}$ , there exists a unique  $y \in D^K$  satisfying (1.94).*

Below is a proof of Proposition 1.

*Proof of Proposition 1.* Fix  $(\xi, \zeta) \in D^{J+K}$  satisfying (1.56)–(1.57). We first prove existence.

By Lemma 4, there exists a  $y \in D^K$  satisfying (1.94). For  $j = 1, \dots, J$ , define

$$u_j := \mu_j^{-1} \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) ds \right), \quad (1.95)$$

$$x_j := \phi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) ds \right). \quad (1.96)$$

Since  $y \in D^K$ , it follows that  $u \in D^J$  and  $x \in D^J$ , so that  $(x, u, y) \in D^{2J+K}$ . It remains to show that  $(x, u, y)$  satisfy (1.58)–(1.62). Equation (1.58) follows from the definitions of  $u$  and  $x$  in (1.95) and (1.96), respectively, as well as definition of the one-sided reflection map  $(\psi, \phi)$  in (1.85)–(1.86). Equation (1.59) holds by the fact that  $y$  satisfies (1.94) and the definition of  $u$  in (1.95). Equation (1.60) follows from (1.56), (1.58)–(1.59), and the fact that

$$\sum_{k=1}^K p_{jk} = 1 \quad \text{for all } j = 1, \dots, J,$$

$$\sum_{j=1}^J q_{kj} = 1 \quad \text{for all } k = 1, \dots, K.$$

Equation (1.61) follows from the definition of  $u$  in (1.95), the definition of  $\psi$  in (1.85), and (1.57). Finally, if we let

$$z_j := \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) ds, \quad (1.97)$$

we see that  $u_j = \mu_j^{-1} \psi(z_j)$  and  $x_j = \phi(z_j)$ . Therefore, (1.62) says that

$$\int_0^\infty \mathbb{1}_{\{\phi(z_j) > 0\}} d(\mu_j^{-1} \psi(z_j))(t) = 0, \quad (1.98)$$

which is equivalent to

$$\int_0^\infty \mathbb{1}_{\{\phi(z_j) > 0\}} d\psi(z_j)(t) = 0. \quad (1.99)$$

But (1.99) holds true by the definition of  $(\psi, \phi)$  given by (1.85)–(1.86). Thus, (1.62) holds.

We now prove uniqueness. Let  $(x, u, y) \in D^{2J+K}$  and  $(\tilde{x}, \tilde{u}, \tilde{y}) \in D^{2J+K}$  both satisfy (1.58)–(1.62). Then, by (1.58), we have that

$$x_j = z_j + \mu_j u_j \geq 0, \quad (1.100)$$

$$\tilde{x}_j = \tilde{z}_j + \mu_j \tilde{u}_j \geq 0, \quad (1.101)$$

where  $z_j$  and  $\tilde{z}_j$  are both given by (1.97), using  $y_l$  and  $\tilde{y}_l$ , respectively. Since  $(x, u)$  and  $(\tilde{x}, \tilde{u})$  satisfy (1.61)–(1.62), it follows that  $(x, \mu_j u_j)$  and  $(\tilde{x}, \mu_j \tilde{u}_j)$  satisfy (1.61)–(1.62). By this and (1.100)–(1.101), it follows that

$$\mu_j u_j = \psi(z_j), \quad (1.102)$$

$$\mu_j \tilde{u}_j = \psi(\tilde{z}_j). \quad (1.103)$$

It then follows by (1.59), (1.97), (1.102)–(1.103), and Lemma 4 that  $y_k = \tilde{y}_k$  for all  $k = 1, \dots, K$ . By uniqueness of  $y$ , it then follows by (1.97) and (1.102)–(1.103) that  $u_j = \tilde{u}_j$  for all  $j = 1, \dots, J$ . Finally, by uniqueness of  $y$  and  $u$ , it follows by (1.97) and (1.100)–(1.101) that  $x_j = \tilde{x}_j$  for all  $j = 1, \dots, J$ . This completes the proof.  $\square$

Before a proof of Proposition 2, we provide a description of the mappings  $f$ ,  $g$ , and  $h$  in the statement of Corollary 1. First, by Lemma 4, we can indirectly write  $h : D^{J+K} \rightarrow D^K$  as the unique mapping sending  $(\xi, \zeta) \in D^{J+K}$  to  $h(\xi, \zeta) \in D^K$  satisfying (1.94). For our purposes, this indirect description of  $h$  will be enough. On the other hand, the proof of Proposition 1 shows that the mappings  $f : D^{J+K} \rightarrow D^J$  and  $g : D^{J+K} \rightarrow D^J$  are uniquely given by the following:

$$f := (\xi, \zeta) \mapsto \left( \phi \left( \pi_j \circ \xi + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot (\pi_l \circ h(\xi, \zeta))(s) \, ds \right) \right)_{j=1, \dots, J}, \quad (1.104)$$

$$g := (\xi, \zeta) \mapsto \left( \mu_j^{-1} \psi \left( \pi_j \circ \xi + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot (\pi_l \circ h(\xi, \zeta))(s) \, ds \right) \right)_{j=1, \dots, J}. \quad (1.105)$$

*Proof of Proposition 2.* We first prove continuity of  $f$ ,  $g$ , and  $h$  separately. Then we use these results to argue that the joint mapping  $(f, g, h)$  is continuous.

**Continuity of  $h$ .** Recall that for each  $(\xi, \zeta) \in D^{J+K}$  the element  $y = h(\xi, \zeta) \in D^K$  satisfies

$$y_k(t) = \zeta_k(t) - \eta_k \int_0^t y_k(s) \, ds - \sum_{j=1}^J p_{jk} \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) \, ds \right)(t), \quad t \geq 0, \quad (1.106)$$

for all  $k = 1, \dots, K$ . Suppose that  $(\xi^n, \zeta^n) \rightarrow (\xi, \zeta)$  in  $D^{J+K}$  as  $n \rightarrow \infty$  and let  $\tilde{T} > 0$  be a

continuity point of  $f(\xi, \zeta)$ . To complete the proof, we must show that

$$d_{\tilde{T}}^K \left( h(\xi^n, \zeta^n)|_{[0, \tilde{T}]}, h(\xi, \zeta)|_{[0, \tilde{T}]} \right) \rightarrow 0 \quad \text{as } n \rightarrow \infty, \quad (1.107)$$

where  $d_{\tilde{T}}^K$  is given by (1.90). However, by Lemma 5.1 in Ethier and Kurtz [2005], since  $(\xi, \zeta) \in D^{J+K}$ , it has at most countably many points of discontinuity. As a result, there exists a point  $T > \tilde{T}$  that is a continuity point of  $(\xi, \zeta)$ . By Lemma 1 on page 167 in Billingsley [1999], it suffices to show that

$$d_T^K \left( h(\xi^n, \zeta^n)|_{[0, T]}, h(\xi, \zeta)|_{[0, T]} \right) \rightarrow 0 \quad \text{as } n \rightarrow \infty, \quad (1.108)$$

for then (1.108) would imply (1.107). To avoid cumbersome notation, we write

$$d_T^K \left( h(\xi^n, \zeta^n), h(\xi, \zeta) \right) \rightarrow 0 \quad \text{as } n \rightarrow \infty \quad (1.109)$$

to mean (1.108). The remainder of the proof aims at proving (1.109). But, since the metrics  $d_T^K$  and  $\tilde{d}_T^K$  are topologically equivalent, it suffices to prove (1.109) with  $\tilde{d}_T^K$  in place of  $d_T^K$ .

Since  $T > 0$  is a continuity point of  $(\xi, \zeta)$  and  $(\xi^n, \zeta^n) \rightarrow (\xi, \zeta)$  in  $D^{J+K}$ , there exists a sequence of homeomorphisms  $\lambda^n \in \tilde{\Lambda}_T$  such that

$$\|(\xi, \zeta) \circ \lambda^n - (\xi^n, \zeta^n)\|_{T, J+K} \vee \|\dot{\lambda}^n - \mathbf{1}\|_T \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (1.110)$$

Letting  $y = h(\xi, \zeta)$  and  $y^n = h(\xi^n, \zeta^n)$ , for any  $0 \leq t \leq T$  we have

$$\begin{aligned} & \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_t \\ &= \max_{1 \leq k \leq K} \left\| \left( \zeta_k \circ \lambda^n - \eta_k \int_0^{\lambda^n(\cdot)} y_k(s) ds - \sum_{j=1}^J p_{jk} \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_j \int_0^{\cdot} y_l(s) ds \right) \circ \lambda^n \right) \right\| \end{aligned}$$



$$\begin{aligned}
& - \left( \zeta_k^n - \eta_k \int_0^\cdot y_k^n(s) ds - \sum_{j=1}^J p_{jk} \psi \left( \xi_j^n + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l^n(s) ds \right) \right) \Big\|_t \Big] \\
& \leq \max_{1 \leq k \leq K} \left[ \left\| \zeta_k \circ \lambda^n - \zeta_k^n \right\|_t + \eta \left\| \int_0^{\lambda^n(\cdot)} y_k(s) ds - \int_0^\cdot y_k^n(s) ds \right\|_t \right. \\
& \quad \left. + \sum_{j=1}^J \left\| \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) ds \right) \circ \lambda^n - \psi \left( \xi_j^n + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l^n(s) ds \right) \right\|_t \right], \quad (1.111)
\end{aligned}$$

where  $\eta := \max_{1 \leq k \leq K} \eta_k$ . We now bound each of the terms in (1.111). To that end, let

$$M_T := \max_{1 \leq k \leq K} \|y_k\|_T < \infty. \quad (1.112)$$

Then by (1.112) and the chain rule we have that

$$\begin{aligned}
& \left\| \int_0^{\lambda^n(\cdot)} y_k(s) ds - \int_0^\cdot y_k^n(s) ds \right\|_t \\
& = \left\| \int_0^\cdot y_k(\lambda^n(s)) \dot{\lambda}^n(s) ds - \int_0^\cdot y_k^n(s) ds \right\|_t \\
& \leq \left\| \int_0^\cdot y_k(\lambda^n(s)) (\dot{\lambda}^n(s) - 1) ds \right\|_t + \left\| \int_0^\cdot (y_k(\lambda^n(s)) - y_k^n(s)) ds \right\|_t \\
& \leq TM_T \|\dot{\lambda} - 1\|_T + \int_0^t \|y_k \circ \lambda^n - y_k^n\|_s ds. \quad (1.113)
\end{aligned}$$

Similar to (1.113), we have

$$\begin{aligned}
& \left\| \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) ds \right) \circ \lambda^n - \psi \left( \xi_j^n + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l^n(s) ds \right) \right\|_t \\
& \leq \left\| \xi_j \circ \lambda^n - \xi_j^n \right\|_T + \eta \sum_{l=1}^K \left\| \int_0^{\lambda^n(\cdot)} y_l(s) ds - \int_0^\cdot y_l^n(s) ds \right\|_t \\
& \leq \left\| \xi_j \circ \lambda^n - \xi_j^n \right\|_T + \eta \sum_{l=1}^K \left\| \int_0^\cdot y_l(\lambda^n(s)) \dot{\lambda}^n(s) ds - \int_0^\cdot y_l^n(s) ds \right\|_t \\
& \leq \left\| \xi_j \circ \lambda^n - \xi_j^n \right\|_T + \eta K T M_T \|\dot{\lambda}^n - 1\|_T + \eta \sum_{l=1}^K \int_0^t \|y_l \circ \lambda^n - y_l^n\|_s ds, \quad (1.114)
\end{aligned}$$

where the first inequality holds by (1.87) and the fact that (see Lemma 13.5.2 in Whitt [2002])

$$\psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_j \int_0^\cdot y_l(s) ds \right) \circ \lambda^n = \psi \left( \xi_j \circ \lambda^n + \sum_{l=1}^K q_{lj} \eta_j \int_0^{\lambda^n(\cdot)} y_l(s) ds \right). \quad (1.115)$$

By (1.111) and (1.113)–(1.114), it follows that for all  $0 \leq t \leq T$ ,

$$\begin{aligned} & \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_t \\ & \leq \max_{1 \leq k \leq K} \left[ \|\zeta_k \circ \lambda^n - \zeta_k^n\|_T + \eta T M_T \|\dot{\lambda} - 1\|_T + \eta \int_0^t \|y_k \circ \lambda^n - y_k^n\|_s ds \right. \\ & \quad \left. + \sum_{j=1}^J \left( \|\xi_j \circ \lambda^n - \xi_j^n\|_T + \eta K T M_T \|\dot{\lambda}^n - 1\|_T + \eta \sum_{l=1}^K \int_0^t \|y_l \circ \lambda^n - y_l^n\|_s ds \right) \right] \\ & \leq \max_{1 \leq k \leq K} \|\zeta_k \circ \lambda^n - \zeta_k^n\|_T + \eta T M_T \|\dot{\lambda} - 1\|_T + \eta \int_0^t \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_s ds \\ & \quad + J \max_{1 \leq j \leq J} \|\xi_j \circ \lambda^n - \xi_j^n\|_T + \eta J K T M_T \|\dot{\lambda}^n - 1\|_T \\ & \quad + \eta J K \int_0^t \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_s ds \\ & \leq 2J \|(\xi, \zeta) \circ \lambda^n - (\xi^n, \zeta^n)\|_{D_T^{J+K}} + (JK + 1) \eta T M_T \|\dot{\lambda}^n - 1\|_T \\ & \quad + 2\eta J K \int_0^t \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_s ds. \end{aligned} \quad (1.116)$$

For  $\epsilon > 0$  fixed, let  $n_0$  be large enough such that for all  $n \geq n_0$  we have

$$\max \left\{ 2J \|(\xi, \zeta) \circ \lambda^n - (\xi^n, \zeta^n)\|_{T, J+K}, (JK + 1) \eta T M_T \|\dot{\lambda}^n - 1\|_T \right\} < \frac{\epsilon}{2e^{2\eta J K T}}. \quad (1.117)$$

Then by (1.116)–(1.117), it follows for all  $n \geq n_0$  and  $0 \leq t \leq T$  that

$$\max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_t \leq \frac{\epsilon}{e^{2\eta J K T}} + 2\eta J K \int_0^t \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_s ds. \quad (1.118)$$

By Gronwall's inequality (see, e.g., Lemma 4.1 in Pang et al. [2007]) and (1.118),

$$\max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_t \leq \frac{\epsilon}{e^{2\eta JKT}} e^{2\eta JKT} \quad \text{for all } 0 \leq t \leq T. \quad (1.119)$$

In particular, using (1.119) with  $t = T$ , it follows that

$$\|y \circ \lambda^n - y^n\|_{T,K} = \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_T \leq \epsilon. \quad (1.120)$$

Finally, by (1.110) let  $n_1$  be large enough such that for all  $n \geq n_1$  we have

$$\|\dot{\lambda}^n - 1\|_T \leq \epsilon. \quad (1.121)$$

Therefore, by (1.120)–(1.121), for all  $n \geq \max\{n_0, n_1\}$  we have

$$\|y \circ \lambda^n - y^n\|_{T,K} \vee \|\dot{\lambda}^n - 1\|_T \leq \epsilon, \quad (1.122)$$

completing the proof.

**Continuity of  $g$ .** The continuity proof for  $g$  (see (1.105) for its expression) proceeds in the same way as in the continuity proof for  $f$ . Suppose  $(\xi^n, \zeta^n) \rightarrow (\xi, \zeta)$  in  $D^{J+K}$  as  $n \rightarrow \infty$  and let  $T > 0$  be a continuity point of  $(\xi, \zeta)$ . Then there exists a sequence of homeomorphisms  $\lambda^n \in \tilde{\Lambda}_T$  such that (1.110) holds. Letting  $u = g(\xi, \zeta)$ ,  $u^n = g(\xi^n, \zeta^n)$ ,  $y = h(\xi, \zeta)$  and  $y^n = h(\xi^n, \zeta^n)$ , we have:

$$\begin{aligned} & \max_{1 \leq j \leq J} \|u_j \circ \lambda^n - u_j^n\|_T \\ &= \max_{1 \leq j \leq J} \left\| \mu_j^{-1} \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) ds \right) \circ \lambda^n - \mu_j^{-1} \psi \left( \xi_j^n + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l^n(s) ds \right) \right\|_T \\ &\leq \mu^{-1} \max_{1 \leq j \leq J} \left\| \psi \left( \xi_j \circ \lambda^n + \sum_{l=1}^K q_{lj} \eta_l \int_0^{\lambda^n(\cdot)} y_l(s) ds \right) - \psi \left( \xi_j^n + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l^n(s) ds \right) \right\|_T \end{aligned}$$

$$\begin{aligned}
&\leq \mu^{-1} \max_{1 \leq j \leq J} \|\xi_j \circ \lambda^n - \xi_j^n\|_T + \mu^{-1} \eta K T M_T \|\dot{\lambda}^n - 1\|_T \\
&\quad + \mu^{-1} \eta K T \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_T \\
&\leq \mu^{-1} \|(\xi, \zeta) \circ \lambda^n - (\xi^n, \zeta^n)\|_{T, J+K} + \mu^{-1} \eta K T M_T \|\dot{\lambda}^n - 1\|_T \\
&\quad + \mu^{-1} \eta K T \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_T, \tag{1.123}
\end{aligned}$$

where  $\mu := \min_{1 \leq j \leq J} \mu_j$  and where the second inequality follows from (1.87). Since  $(\xi^n, \zeta^n) \rightarrow (\xi, \zeta)$  in  $D^{J+K}$ , for  $\epsilon > 0$  fixed there exists an  $n_0$  such that for all  $n \geq n_0$ ,

$$\mu^{-1} \|(\xi, \zeta) \circ \lambda^n - (\xi^n, \zeta^n)\|_{T, J+K} \vee \mu^{-1} \eta K T M_T \|\dot{\lambda}^n - 1\|_T \leq \frac{\epsilon}{3}. \tag{1.124}$$

Furthermore, by (1.120), there exists an  $n_1$  such that for all  $n \geq n_1$ ,

$$\mu^{-1} \eta K T \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_T \leq \frac{\epsilon}{3}. \tag{1.125}$$

Finally, let  $n_2$  be large enough such that (1.121) holds for all  $n \geq n_2$ . Then by (1.123)–(1.125), it follows that for all  $n \geq \max\{n_0, n_1, n_2\}$ ,

$$\|u \circ \lambda^n - u^n\|_{T, J} \vee \|\dot{\lambda}^n - 1\|_T \leq \epsilon, \tag{1.126}$$

which is the desired result.

**Continuity of  $f$ .** The continuity proof for  $f$  is nearly identical to the continuity proof for  $g$  (see (1.104) for its expression). Suppose  $(\xi^n, \zeta^n) \rightarrow (\xi, \zeta)$  in  $D^{J+K}$  as  $n \rightarrow \infty$  and let  $T > 0$  be a continuity point of  $(\xi, \zeta)$ . Then there exists a sequence of homeomorphisms  $\lambda^n \in \tilde{\Lambda}_T$  such that (1.110) holds. Letting  $x = f(\xi, \zeta)$ ,  $x^n = f(\xi^n, \zeta^n)$ ,  $y = h(\xi, \zeta)$  and  $y^n = h(\xi^n, \zeta^n)$ ,

we have:

$$\begin{aligned}
& \max_{1 \leq j \leq J} \|x_j \circ \lambda^n - x_j^n\|_T \\
&= \max_{1 \leq j \leq J} \left\| \phi \left( \xi_j \circ \lambda^n + \sum_{l=1}^K q_{lj} \eta_l \int_0^{\lambda^n(\cdot)} y_l(s) ds \right) - \phi \left( \xi_j^n + \sum_{l=1}^K q_{lj} \eta_l \int_0^{\cdot} y_l^n(s) ds \right) \right\|_T \\
&\leq 2 \max_{1 \leq j \leq J} \|\xi_j \circ \lambda^n - \xi_j^n\|_T + 2\eta \sum_{l=1}^K \left\| \int_0^{\cdot} y_l(\lambda^n(s)) \dot{\lambda}^n(s) ds - \int_0^{\cdot} y_l^n(s) ds \right\|_T \\
&\leq 2 \|(\xi, \zeta) \circ \lambda^n - (\xi^n, \zeta^n)\|_{T, J+K} + 2\eta K T M_T \|\dot{\lambda}^n - 1\|_T \\
&\quad + 2\eta K T \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_T, \tag{1.127}
\end{aligned}$$

where the first inequality follows from (1.88). We see that the right-hand side (1.127) is the same as the right-hand side (1.123), except with  $\mu^{-1}$  replaced by 2. Thus, the same arguments used in showing the continuity of  $g$  can be used to show the continuity of  $h$ .

**Continuity of  $(f, g, h)$ .** We regard  $(f, g, h)$  as a function from  $D^{J+K} \rightarrow D^{2J+K}$  defined by  $(\xi, \zeta) \mapsto (f(\xi, \zeta), g(\xi, \zeta), h(\xi, \zeta))$ . Suppose that  $(\xi^n, \zeta^n) \rightarrow (\xi, \zeta)$  in  $D^{J+K}$  and let  $T > 0$  be a continuity point of  $(\xi, \zeta)$ . Therefore, there exists a sequence of homeomorphisms  $\lambda^n \in \tilde{\Lambda}_T$  such that (1.110) holds. For notational purposes, let  $F^n = (f(\xi^n, \zeta^n), g(\xi^n, \zeta^n), h(\xi^n, \zeta^n))$ ,  $F = (f(\xi, \zeta), g(\xi, \zeta), h(\xi, \zeta))$ ,  $x^n = f(\xi^n, \zeta^n)$ ,  $x = f(\xi, \zeta)$ ,  $u^n = g(\xi^n, \zeta^n)$ ,  $u = g(\xi, \zeta)$ ,  $y^n = h(\xi^n, \zeta^n)$ , and  $y = h(\xi, \zeta)$ . Using this notation, to prove continuity of  $(f, g, h)$  we it suffices to show that

$$\|F \circ \lambda^n - F^n\|_{T, 2J+K} \vee \|\dot{\lambda}^n - 1\| \rightarrow 0 \quad \text{as } n \rightarrow \infty. \tag{1.128}$$

But note that

$$\begin{aligned}
& \|F \circ \lambda^n - F^n\|_{T, 2J+K} \\
&\equiv \max_{1 \leq j \leq J} \|x_j \circ \lambda^n - x_j^n\|_T \vee \max_{1 \leq j \leq J} \|u_j \circ \lambda^n - u_j^n\|_T \vee \max_{1 \leq k \leq K} \|y_k \circ \lambda^n - y_k^n\|_T. \tag{1.129}
\end{aligned}$$

Since each of the three terms on the right-hand side of (1.129) converge to zero (by continuity of  $f$ ,  $g$ , and  $h$  separately), (1.128) will follow. The proof is complete.  $\square$

### 1.8.3 Proof of Lemma 4

The general proof technique we use parallels that of Lemma 1 in Reed and Ward [2004]. We prove that for each  $T > 0$ , there exists a unique  $y \in D_T^K$  satisfying (1.94), then extend this to  $D^K$  in the obvious way. To improve the readability of this proof, we break it up into a few separate steps.

#### Existence of an Element in $D_K^T$ Satisfying (1.94)

We prove existence via the method of successive approximations. In particular, we show that the sequence formed by this method is Cauchy in  $D_K^T$ , and then argue that the limit of the sequence (by completeness of  $D_K^T$ ) satisfies (1.94). To that end, let  $y_k^0 \equiv 0 \in D$  and let  $y_k^n \in D$ ,  $n \geq 1$ , be defined by

$$y_k^n := \xi_k - \eta_k \int_0^\cdot y_k^{n-1}(s) \, ds - \sum_{j=1}^J p_{jk} \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l^{n-1} \, ds \right), \quad (1.130)$$

for each  $k = 1, \dots, K$ . We claim that the sequence  $\left\{ (y_1^n|_{[0,T]}, \dots, y_K^n|_{[0,T]}) : n = 0, 1, \dots \right\}$  defined by (1.130) is Cauchy in  $D_T^K$ ; see Claim 1 at the end of Appendix 1.8.3 for a proof of this claim. By completeness of  $(D_T^K, \|\cdot\|_{T,K})$ , it follows that

$$\left( y_1^n|_{[0,T]}, \dots, y_K^n|_{[0,T]} \right) \rightarrow \left( y_{1,T}^\infty, \dots, y_{K,T}^\infty \right) \in D_T^K \quad \text{as } n \rightarrow \infty. \quad (1.131)$$

To show that  $(y_{1,T}^\infty, \dots, y_{K,T}^\infty)$  satisfies (1.94), define the mapping  $L : D_T^K \rightarrow D_T^K$  by

$$(y_1, \dots, y_K) \mapsto \left( \zeta_k - \eta_k \int_0^\cdot y_k(s) \, ds - \sum_{j=1}^J p_{jk} \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) \, ds \right) \right)_{k=1, \dots, K}. \quad (1.132)$$

Then for  $y, \tilde{y} \in D_T^K$  we have:

$$\begin{aligned}
& \|L(y) - L(\tilde{y})\|_{T,K} \\
&= \max_{1 \leq k \leq K} \left\| \eta_k \int_0^\cdot (y_k(s) - \tilde{y}_k(s)) \, ds \right. \\
&\quad \left. + \sum_{j=1}^J p_{jk} \left[ \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) \, ds \right) - \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot \tilde{y}_l(s) \, ds \right) \right] \right\|_T \\
&\leq \max_{1 \leq k \leq K} \left\{ \eta T \|y_k - \tilde{y}_k\|_T + \sum_{j=1}^J \sum_{l=1}^K \eta T \|y_l - \tilde{y}_l\|_T \right\} \\
&\leq \eta T \|y - \tilde{y}\|_{T,K} + \eta T JK \|y - \tilde{y}\|_{T,K} \\
&\leq \eta T JK \|y - \tilde{y}\|_{T,K}. \tag{1.133}
\end{aligned}$$

Equation (1.133) shows that  $L$  is Lipschitz continuous, so by (1.130)–(1.131) we have

$$\begin{aligned}
& \left( y_{1,T}^\infty, \dots, y_{K,T}^\infty \right) \leftarrow \left( y_1^{n+1}|_{[0,T]}, \dots, y_K^{n+1}|_{[0,T]} \right) \\
&= L \left( y_1^n|_{[0,T]}, \dots, y_K^n|_{[0,T]} \right) \rightarrow L \left( y_{1,T}^\infty, \dots, y_{K,T}^\infty \right) \quad \text{as } n \rightarrow \infty. \tag{1.134}
\end{aligned}$$

By uniqueness of limits in metric spaces, it follows that

$$L \left( y_{1,T}^\infty, \dots, y_{K,T}^\infty \right) = \left( y_{1,T}^\infty, \dots, y_{K,T}^\infty \right), \tag{1.135}$$

implying that  $(y_{1,T}^\infty, \dots, y_{K,T}^\infty)$  satisfies (1.94).

Uniqueness of the Element in  $D_T^K$  Satisfying (1.94)

We show that  $(y_{1,T}^\infty, \dots, y_{K,T}^\infty)$  is the unique element in  $D_T^K$  satisfying (1.94). To that end, suppose that  $(y_1, \dots, y_K), (\tilde{y}_1, \dots, \tilde{y}_K) \in D_T^K$  both satisfy (1.94). Define

$$m := \inf \left\{ n \geq 1 : \frac{n}{2} (2JK\eta)^{-1} > T \right\}. \tag{1.136}$$

Then for  $t_1 = \frac{1}{2} (2JK\eta)^{-1}$  we have

$$\begin{aligned}
& \max_{1 \leq k \leq K} \|y_k - \tilde{y}_k\|_t \\
&= \max_{1 \leq k \leq K} \left\| \eta_k \int_0^\cdot (y_k(s) - \tilde{y}(s)) ds \right. \\
&\quad \left. + \sum_{j=1}^J p_{jk} \left[ \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l(s) ds \right) - \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot \tilde{y}_l(s) ds \right) \right] \right\|_{t_1} \\
&\leq \max_{1 \leq k \leq K} \eta \int_0^{t_1} \|y_k - \tilde{y}_k\|_{t_1} ds + \sum_{j=1}^J \sum_{l=1}^K \eta \int_0^{t_1} \|y_l - \tilde{y}_l\|_{t_1} ds \\
&\leq 2JK\eta t_1 \max_{1 \leq k \leq K} \|y_k - \tilde{y}_k\|_{t_1}. \tag{1.137}
\end{aligned}$$

Since  $2JK\eta t_1 = \frac{1}{2} < 1$ , (1.137) implies that  $\max_{1 \leq k \leq K} \|y_k - \tilde{y}_k\|_{t_1} = 0$ , so that

$$y = \tilde{y} \quad \text{on} \quad [0, t_1]. \tag{1.138}$$

Now for  $t_2 = (2JK\eta)^{-1}$  we have

$$\begin{aligned}
\max_{1 \leq k \leq K} \|y_k - \tilde{y}_k\|_{t_2} &\leq \eta \left\| \int_0^\cdot |y_k(s) - \tilde{y}_k(s)| ds \right\|_{t_2} + \sum_{j=1}^J \sum_{l=1}^K \eta \left\| \int_0^\cdot |y_l(s) - \tilde{y}_l(s)| ds \right\|_{t_2} \\
&= \max_{1 \leq k \leq K} \eta \left\| \int_0^{t_1} |y_k(s) - \tilde{y}_k(s)| ds + \int_{t_1}^\cdot |y_k(s) - \tilde{y}_k(s)| ds \right\|_{t_2} \\
&\quad + \eta J \sum_{l=1}^K \left\| \int_0^{t_1} |y_l(s) - \tilde{y}_l(s)| ds + \int_{t_1}^\cdot |y_l(s) - \tilde{y}_l(s)| ds \right\|_{t_2} \\
&\leq \eta t_1 \max_{1 \leq k \leq K} \|y_k - \tilde{y}_k\|_{t_1} + (t_2 - t_1) \max_{1 \leq k \leq K} \|y_k - \tilde{y}_k\|_{t_2} \\
&\quad + \eta J \sum_{l=1}^K \left[ t_1 \|y_l - \tilde{y}_l\|_{t_1} + (t_2 - t_1) \|y_l - \tilde{y}_l\|_{t_2} \right] \\
&\leq 2\eta JK(t_2 - t_1) \max_{1 \leq k \leq K} \|y_k - \tilde{y}_k\|_{t_2}, \tag{1.139}
\end{aligned}$$



where in the last inequality we use the fact that (1.138) holds. Since  $2\eta JK(t_2 - t_1) = \frac{1}{2} < 1$ , (1.139) implies that  $\|y_k - \tilde{y}_k\|_{t_2} = 0$ , so that

$$y = \tilde{y} \quad \text{on} \quad [0, t_2].$$

We can continue in this iterative manner to show that for each  $n = 1, \dots, m - 1$ ,

$$y = \tilde{y} \quad \text{on} \quad [0, t_n], \tag{1.140}$$

where  $t_n = \frac{n}{2}(2JK\eta)^{-1}$  for all  $n = 1, \dots, m - 1$ . If  $t_{m-1} = T$ , then we are done by (1.140). On the other hand, if  $t_{m-1} < T$  then we can take  $t_m = T$  and show using the same argument that  $y = \tilde{y}$  on  $[0, t_m]$ , which would then complete the proof.

### Extension to a Unique Element in $D^K$ Satisfying (1.94)

We use the constructed unique solution in  $D_T^K$  to define a  $(y_1^\infty, \dots, y_K^\infty)$  that is the unique element in  $D^K$  satisfying (1.94). To that end, define  $(y_1^\infty, \dots, y_K^\infty) \in D^K$  by

$$(y_1^\infty, \dots, y_K^\infty)(t) := (y_{1,T}^\infty, \dots, y_{K,T}^\infty)(t) \quad \text{for} \quad t \in [0, T]. \tag{1.141}$$

To complete the proof, we must show that  $(y_1^\infty, \dots, y_K^\infty)$  is well-defined and is the unique element in  $D^K$  satisfying (1.94). To prove that  $(y_1^\infty, \dots, y_K^\infty)$  is well-defined, we must show that whenever  $t \in [0, T_1] \cap [0, T_2]$  that

$$(y_{1,T_1}^\infty, \dots, y_{K,T_1}^\infty)(t) = (y_{1,T_2}^\infty, \dots, y_{K,T_2}^\infty)(t). \tag{1.142}$$

Without loss of generality, suppose that  $T_1 \leq T_2$ . Then  $(y_{1,T_2}^\infty, \dots, y_{K,T_2}^\infty)|_{[0,T_1]} \in D_{T_1}^K$  satisfies (1.94) for all  $t \leq T_1$ . By uniqueness,

$$(y_{1,T_1}^\infty, \dots, y_{K,T_1}^\infty) = (y_{1,T_2}^\infty, \dots, y_{K,T_2}^\infty)|_{[0,T_1]},$$

so we have that

$$(y_{1,T_1}^\infty, \dots, y_{K,T_1}^\infty)(t) = (y_{1,T_2}^\infty, \dots, y_{K,T_2}^\infty)|_{[0,T_1]}(t) = (y_{1,T_2}^\infty, \dots, y_{K,T_2}^\infty)(t),$$

proving (1.142) holds. Next we show that  $(y_1^\infty, \dots, y_K^\infty)$  satisfies (1.94). Indeed, for any  $t \in [0, \infty)$ , there exists a  $T > 0$  such that (1.141) holds. But since the right-hand side of (1.141) satisfies (1.94) at  $t$ , so does  $(y_1^\infty, \dots, y_K^\infty)$ , which gives the desired result. Finally, we show that  $(y_1^\infty, \dots, y_K^\infty)$  is unique. To that end, let  $(y_1, \dots, y_K) \in D^K$  also satisfy (1.94) and fix  $t \in [0, \infty)$ . Let  $T > 0$  be such that  $t \in [0, T]$ . Then  $(y_1, \dots, y_K)|_{[0,T]} \in D_T^K$  satisfies (1.94) for all  $t \in [0, T]$ . Since  $(y_1^\infty, \dots, y_K^\infty)|_{[0,T]}$  also satisfies (1.94) for all  $t \in [0, T]$ , uniqueness implies that

$$(y_1^\infty, \dots, y_K^\infty)(t) = (y_1^\infty, \dots, y_K^\infty)|_{[0,T]}(t) = (y_1, \dots, y_K)|_{[0,T]}(t) = (y_1, \dots, y_K)(t).$$

Therefore,  $(y_1^\infty, \dots, y_K^\infty) = (y_1, \dots, y_K)$ , as desired. This completes the proof.  $\square$

We conclude Appendix 1.8.3 with a proof showing that the sequence defined by (1.130) is a Cauchy sequence.

**Claim 1.** *For each  $T > 0$ , the sequence  $\left\{ \left( y_1^n|_{[0,T]}, \dots, y_K^n|_{[0,T]} \right) : n = 0, 1, \dots \right\}$  defined by (1.130) is Cauchy in  $D_K^T$  (with respect to the norm  $\|\cdot\|_{T,K}$ ).*

*Proof.* For any  $t \in [0, T]$  note that

$$\|y_k^1 - y_k^0\|_t = \left\| \zeta_k - \sum_{j=1}^J p_{jk} \psi(\xi_j) \right\|_t \leq \|\zeta_k\|_t + \sum_{j=1}^J \|\psi(\xi_j)\|_t$$

$$\leq \max_{1 \leq k \leq K} \|\zeta_k\|_T + \sum_{j=1}^J \|\psi(\xi_j)\|_T =: C_T. \quad (1.143)$$

Let  $\eta := \max_{1 \leq k \leq K} \eta_k$  and fix  $\delta \in (0, T)$  such that

$$2\delta\eta JK < 1. \quad (1.144)$$

Then for  $n \geq 2$  we have

$$\begin{aligned} & \|y_k^n - y_k^{n-1}\|_\delta \\ &= \left\| \eta_k \int_0^\cdot (y_k^{n-2}(s) - y_k^{n-1}(s)) \, ds \right. \\ & \quad \left. + \sum_{j=1}^J p_{jk} \left[ \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l^{n-2} \, ds \right) - \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l^{n-1} \, ds \right) \right] \right\|_\delta \\ &\leq \eta_k \delta \|y_k^{n-1} - y_k^{n-2}\|_\delta + \sum_{j=1}^J p_{jk} \left\| \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot (y_l^{n-1}(s) - y_l^{n-2}(s)) \, ds \right\|_\delta \\ &\leq \eta_k \delta \|y_k^{n-1} - y_k^{n-2}\|_\delta + \sum_{j=1}^J p_{jk} \sum_{l=1}^K q_{lj} \eta_l \delta \|y_l^{n-1} - y_l^{n-2}\|_\delta \\ &\leq \delta \eta \left( \|y_k^{n-1} - y_k^{n-2}\|_\delta + J \sum_{l=1}^K \|y_l^{n-1} - y_l^{n-2}\|_\delta \right) \\ &\leq 2\delta\eta J \sum_{l=1}^K \|y_l^{n-1} - y_l^{n-2}\|_\delta, \end{aligned} \quad (1.145)$$

where the first inequality follows from (1.87). Doing another iteration of (1.145) gives

$$\|y_k^n - y_k^{n-1}\|_\delta \leq 2\delta\eta J \sum_{l=1}^K \|y_l^{n-1} - y_l^{n-2}\|_\delta \leq (2\delta\eta J)^2 K \sum_{l=1}^K \|y_l^{n-2} - y_l^{n-3}\|_\delta. \quad (1.146)$$

Continuing in the way and using (1.143), we find that

$$\|y_k^n - y_k^{n-1}\|_\delta \leq (2\delta\eta JK)^{n-1} C_T \quad \text{for all } n \geq 1. \quad (1.147)$$

For each  $k = 1, \dots, K$ , we now prove that for all  $m \geq 1$  and  $n \geq 1$  that

$$\|y_k^n - y_k^{n-1}\|_{m\delta} \leq mn^m (2\delta\eta JK)^{n-1} C_T. \quad (1.148)$$

We proceed by (strong) induction on  $m$ . The base case follows by (1.147) since for any  $n \geq 1$ ,

$$\|y_k^n - y_k^{n-1}\|_{1\cdot\delta} \leq (2\delta\eta JK)^{n-1} C_T \leq 1 \cdot n^1 (2\delta\eta JK)^{n-1} C_T, \quad (1.149)$$

which is the  $m = 1$  case. For the inductive step, assume that for all  $n \geq 2$  we have

$$\|y_k^n - y_k^{n-1}\|_{r\delta} \leq rn^r (2\delta\eta JK)^{n-1} C_T \quad \text{for all } r = 1, \dots, m. \quad (1.150)$$

By (1.150), for  $n \geq 2$  we have

$$\begin{aligned} \|y_k^n - y_k^{n-1}\|_{(m+1)\delta} &= \left\| \eta_k \int_0^\cdot (y_k^{n-2}(s) - y_k^{n-1}(s)) \, ds \right. \\ &\quad \left. + \sum_{j=1}^J p_{jk} \left[ \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l^{n-2} \, ds \right) \right. \right. \\ &\quad \left. \left. - \psi \left( \xi_j + \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot y_l^{n-1} \, ds \right) \right] \right\|_{(m+1)\delta} \\ &\leq \eta_k \sum_{r=1}^{m+1} \int_{(r-1)\delta}^{r\delta} \|y_k^{n-1} - y_k^{n-2}\|_{r\delta} \, ds \\ &\quad + \sum_{j=1}^J p_{jk} \left\| \sum_{l=1}^K q_{lj} \eta_l \int_0^\cdot |y_l^{n-1}(s) - y_l^{n-2}(s)| \, ds \right\|_{(m+1)\delta} \\ &\leq \delta\eta \sum_{r=1}^{m+1} \|y_k^{n-1} - y_k^{n-2}\| + \eta \sum_{j=1}^J \sum_{l=1}^K \sum_{r=1}^{m+1} \int_{(r-1)\delta}^{r\delta} \|y_l^{n-1} - y_l^{n-2}\|_{r\delta} \, ds \\ &= \delta\eta \left[ \sum_{r=1}^{m+1} \|y_k^{n-1} - y_k^{n-2}\| + J \sum_{l=1}^K \sum_{r=1}^{m+1} \|y_l^{n-1} - y_l^{n-2}\|_{r\delta} \right] \\ &\leq 2\delta\eta J \sum_{l=1}^K \sum_{r=1}^{m+1} \|y_l^{n-1} - y_l^{n-2}\|_{r\delta} \end{aligned}$$

$$\begin{aligned}
&= 2\delta\eta J \sum_{l=1}^K \sum_{r=1}^m \|y_l^{n-1} - y_l^{n-2}\|_{r\delta} + 2\delta\eta J \sum_{l=1}^K \|y_l^{n-1} - y_l^{n-2}\|_{(m+1)\delta} \\
&\leq 2\delta\eta J \sum_{l=1}^K \sum_{r=1}^m r(n-1)^r (2\delta\eta JK)^{n-2} C_T + 2\delta\eta J \sum_{l=1}^K \|y_l^{n-1} - y_l^{n-2}\|_{(m+1)\delta} \\
&= (2\delta\eta JK)^{n-1} C_T \sum_{r=1}^m r(n-1)^r + 2\delta\eta J \sum_{l=1}^K \|y_l^{n-1} - y_l^{n-2}\|_{(m+1)\delta} \\
&\leq (2\delta\eta JK)^{n-1} C_T m n^m + 2\delta\eta J \sum_{l=1}^K \|y_l^{n-1} - y_l^{n-2}\|_{(m+1)\delta}. \tag{1.151}
\end{aligned}$$

By (1.143) and (1.151), for  $n = 2$  we have

$$\begin{aligned}
\|y_k^2 - y_k^1\|_{(m+1)\delta} &\leq (2\delta\eta JK) C_T m 2^m + 2\delta\eta J \sum_{l=1}^K \|y_l^1 - y_l^0\|_{(m+1)\delta} \\
&\leq (2\delta\eta JK) C_T m 2^m + 2\delta\eta JK C_T \\
&= (2\delta\eta JK) C_T (m 2^m + 1). \tag{1.152}
\end{aligned}$$

Similarly, by (1.151)–(1.152), for  $n = 3$  we have

$$\begin{aligned}
\|y_k^3 - y_k^2\|_{(m+1)\delta} &\leq (2\delta\eta JK)^2 C_T m 3^m + 2\delta\eta J \sum_{l=1}^K \|y_l^2 - y_l^1\|_{(m+1)\delta} \\
&\leq (2\delta\eta JK)^2 C_T m 3^m + 2\delta\eta J \sum_{l=1}^K (2\delta\eta JK) C_T (m 2^m + 1) \\
&= (2\delta\eta JK)^2 C_T (m 3^m + m 2^m + 1). \tag{1.153}
\end{aligned}$$

Continuing in this iterative fashion, we find that for  $n \geq 2$ ,

$$\begin{aligned}
\|y_k^n - y_k^{n-1}\|_{(m+1)\delta} &\leq (2\delta\eta JK)^{n-1} C_T \left( m \sum_{i=2}^n i^m + 1 \right) \\
&\leq (2\delta\eta JK)^{n-1} C_T (m n^{m+1} + 1) \\
&\leq (2\delta\eta JK)^{n-1} C_T (m+1) n^{m+1}. \tag{1.154}
\end{aligned}$$

Because (1.148) also holds for all  $m \geq 1$  when  $n = 1$  [by (1.143)], this completes the inductive step. Thus, (1.148) holds for all  $m \geq 1$  and  $n \geq 1$ .

By (1.148), for all  $n \geq 1$  and  $k = 1, \dots, K$  we have

$$\|y_k^n - y_k^{n-1}\|_T \leq \|y_k^n - y_k^{n-1}\|_{[\delta^{-1}T]\delta} \leq [\delta^{-1}T] n^{[\delta^{-1}T]} (2\delta JK)^{n-1} C_T,$$

implying that

$$\max_{1 \leq k \leq K} \|y_k^n - y_k^{n-1}\|_T \leq [\delta^{-1}T] n^{[\delta^{-1}T]} (2\delta JK)^{n-1} C_T.$$

Since by (1.144) we have

$$\limsup_{n \rightarrow \infty} \left| \frac{[\delta^{-1}T](n+1)^{[\delta^{-1}T]} (2\delta JK)^n C_T}{[\delta^{-1}T] n^{[\delta^{-1}T]} (2\delta JK)^{n-1} C_T} \right| = \limsup_{n \rightarrow \infty} \left[ \frac{(n+1)^{[\delta^{-1}T]}}{n^{[\delta^{-1}T]}} \cdot 2\delta JK \right] = 2\delta JK < 1,$$

the ratio test implies that

$$\sum_{n=1}^{\infty} [\delta^{-1}T] n^{[\delta^{-1}T]} (2\delta JK)^{n-1} C_T < \infty,$$

implying that the sequence  $\left\{ \left( y_1^n|_{[0,T]}, \dots, y_K^n|_{[0,T]} \right) : n = 0, 1, \dots \right\}$  is Cauchy in  $D_T^K$ .  $\square$

#### 1.8.4 Miscellaneous Proofs

Below is the proof of a result used in the proof of Lemma 1 in Section 1.6.

**Lemma 5.** *Let  $\{X^n\}_{n=1}^{\infty}$  be a random sequence in  $D$  such that  $\|X^n\|_T \Rightarrow 0$  for all  $T > 0$ .*

*Then  $X^n \Rightarrow \mathbf{0}$  in  $D$ .*

*Proof.* Note that  $\|X^n\|_T \Rightarrow 0$  for all  $T > 0$  is equivalent to  $\|X^n\|_T \xrightarrow{p} 0$  for all  $T > 0$ . We claim

that  $X^n \xrightarrow{p} \mathbf{0}$  in  $D$ . This amounts to showing that for all  $0 < \epsilon < 1$ ,

$$P\left(\int_0^\infty e^{-t} [d_t(X^n, 0) \wedge 1] dt > \epsilon\right) \rightarrow 0 \quad \text{as } n \rightarrow \infty; \quad (1.155)$$

see, e.g., Chapter 3, Section 3 of Whitt [2002]. More explicitly, we must show that

$$P\left(\int_0^\infty e^{-t} \left[\inf_{\lambda \in \Lambda_t} \{\|X^n \circ \lambda\|_t \vee \|\lambda - e\|_t\} \wedge 1\right] dt > \epsilon\right) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (1.156)$$

However, for all  $T > 0$ , we have the following inequalities:

$$\begin{aligned} & P\left(\int_0^\infty e^{-t} \left[\inf_{\lambda \in \Lambda_t} \{\|X^n \circ \lambda\|_t \vee \|\lambda - e\|_t\} \wedge 1\right] dt > \epsilon\right) \\ & \leq P\left(\int_0^\infty e^{-t} [\|X^n\|_t \wedge 1] dt > \epsilon\right) \\ & \leq P\left(\int_0^T e^{-t} \|X^n\|_t dt + \int_T^\infty e^{-t} dt > \epsilon\right). \end{aligned} \quad (1.157)$$

Let  $T > 0$  be such that  $\int_T^\infty e^{-t} dt = \epsilon/2$ . Then continuing from (1.157) yields

$$\begin{aligned} P\left(\int_0^T e^{-t} \|X^n\|_t dt + \int_T^\infty e^{-t} dt > \epsilon\right) &= P\left(\int_0^T e^{-t} \|X^n\|_t dt > \frac{\epsilon}{2}\right) \\ &\leq P\left(\|X^n\|_T \left(1 - \frac{\epsilon}{2}\right) > \frac{\epsilon}{2}\right) \\ &= P\left(\|X^n\|_T > \frac{\epsilon}{2} \left(1 - \frac{\epsilon}{2}\right)\right). \end{aligned} \quad (1.158)$$

Since  $\|X^n\|_T \xrightarrow{p} 0$ , it follows from (1.158) that  $X^n \xrightarrow{p} \mathbf{0}$  in  $D$ . Since convergence in probability implies convergence in distribution, we conclude that  $X^n \Rightarrow \mathbf{0}$  in  $D$ .  $\square$

### Omitted Details in the Proof of Lemma 1

We show that the right-hand sides of (1.64) and (1.65) converge weakly to a nondegenerate limit and conclude that  $\|\bar{\xi}_j^n\|_T \Rightarrow 0$  and  $\|\bar{\zeta}_k^n\|_T \Rightarrow 0$  for each  $T > 0$ . We show this only for

(1.64); the proof for (1.65) is identical. For the remainder of this argument, we fix  $T > 0$ . By (1.46) and the continuous mapping theorem we have  $\|\widehat{Q}_j^n(0)\|_T \Rightarrow \|Q_j(0)\|_T$ , so that

$$n^{-1/2} \|\widehat{Q}_j^n(0)\|_T \Rightarrow 0 \cdot \|Q_j(0)\|_T = 0. \quad (1.159)$$

Define  $\widetilde{M}_k(t) := M_k(\eta_k t)$  for  $t \geq 0$ , so that  $\widetilde{M}_k$  is a Poisson process with rate  $\eta_k$ . Therefore, by the functional central limit theorem for renewal processes (see, e.g., Theorem 14.6 in Billingsley [1999]), it follows that

$$\widehat{M}_k^n \Rightarrow \eta_k^{1/2} B_k, \quad (1.160)$$

where  $B_k$  is a standard Brownian motion and

$$\widehat{M}_k^n(t) := \frac{\widetilde{M}_k(nt) - n\eta_k t}{\sqrt{n}} = \frac{M_k(n\eta_k t) - n\eta_k t}{\sqrt{n}} = \widehat{M}_k^n(\eta_k t), \quad t \geq 0. \quad (1.161)$$

By (1.160)–(1.161) and the continuous mapping theorem, note that

$$n^{-1/2} \|\widehat{M}_k^n(\eta_k \cdot)\|_T \Rightarrow 0 \cdot \|\eta^{1/2} B_k\|_T = 0. \quad (1.162)$$

Furthermore, observe by (1.35) and (1.161) that

$$\bar{M}_k^n(\eta_k t) = \frac{\widehat{M}_k^n(\eta_k t)}{\sqrt{n}} + \eta_k t, \quad t \geq 0. \quad (1.163)$$

Therefore, by (1.160) and (1.163) we have that

$$\bar{M}_k^n(\eta_k \cdot) = \frac{\widehat{M}_k(\eta_k \cdot)}{\sqrt{n}} + \eta_k e \Rightarrow 0 \cdot \eta_k^{1/2} B_k + \eta_k e = \eta_k e, \quad (1.164)$$



where  $e : [0, \infty) \rightarrow [0, \infty)$  is the identity map, i.e.,  $e(t) = t$  for  $t \geq 0$ . By (1.28) and Donsker's theorem (see, e.g., Theorem 4 in Glynn [1990]), it follows that

$$\widehat{\Phi}_{kj}^n \Rightarrow \sqrt{q_{kj}(1-q_{kj})} B_{kj}, \quad (1.165)$$

where  $B_{kj}$  is a standard Brownian motion. Since  $(t \mapsto \bar{M}_k^n(\eta t)) : [0, \infty) \rightarrow [0, \infty)$  is non-decreasing function, the random time change theorem (see, e.g., Proposition 5 in Glynn [1990]) and (1.164)–(1.165) yields

$$\widehat{\Psi}_{kj}^n(\bar{M}_k^n(\eta_k \cdot)) \Rightarrow \sqrt{q_{kj}(1-q_{kj})} B_{kj}(\eta_k \cdot). \quad (1.166)$$

By (1.166) and the continuous mapping theorem, we have that

$$n^{-1/2} \|\widehat{\Psi}_{kj}^n(\bar{M}_k^n(\eta_k \cdot))\|_T \rightarrow 0 \cdot \|\sqrt{q_{kj}(1-q_{kj})} B_{kj}(\eta_k \cdot)\|_T = 0. \quad (1.167)$$

Similar to (1.160), we have that

$$\widehat{N}_j^n(\mu_j \cdot) \Rightarrow \mu_j^{1/2} B_j, \quad (1.168)$$

where  $B_j$  is standard Brownian motion. Another application of the continuous mapping theorem along with (1.168) gives

$$n^{-1/2} \|\widehat{N}_j^n(\mu_j \cdot)\|_T \Rightarrow 0 \cdot \|\mu_j^{1/2} B_j\|_T = 0. \quad (1.169)$$

Finally, by (1.64), (1.159), (1.162), (1.167), and (1.169), it follows that

$$\|\bar{\xi}_j^n\|_T \leq n^{-1/2} \|\hat{Q}_j^n(0)\|_T + \sum_{k=1}^K n^{-1/2} \|\widehat{\Psi}_{kj}^n(\bar{M}_k^n(\eta_k \cdot))\|_T + n^{-1/2} \|\widehat{N}_j^n(\mu_j \cdot)\|_T$$

$$+ \sum_{k=1}^K n^{-1/2} \|\widehat{M}_k^n(\eta_k \cdot)\|_T \Rightarrow 0.$$

### 1.8.5 Miscellaneous Derivations

Derivation of (1.40)–(1.41)

We begin by deriving (1.40), the diffusion-scaled system equation for the single-server stations.

By (1.8), (1.10), and (1.12)–(1.13), we have

$$\begin{aligned} Q_j^n(t) &= Q_j^n(0) + A_j^n(t) - D_j^n(t) \\ &= Q_j^n(0) + \sum_{k=1}^K \Psi_{kj} \left( M_k \left( \eta_k \int_0^t V_k^n(s) ds \right) \right) - N_j \left( n\mu_j T_j^n(t) \right). \end{aligned} \quad (1.170)$$

Elementary algebraic manipulations and the appropriate scaling in (1.23)–(1.35) applied to (1.170) yields

$$\begin{aligned} \widehat{Q}_j^n(t) &= n^{-1/2} Q_j^n(0) + \sum_{k=1}^K n^{-1/2} \Psi_{kj} \left( M_k \left( \eta_k \int_0^t V_k^n(s) ds \right) \right) - n^{-1/2} N_j \left( n\mu_j T_j^n(t) \right) \\ &= \widehat{Q}_j^n(0) + \sum_{k=1}^K n^{-1/2} \left[ \Psi_{kj} \left( M_k \left( \eta_k \int_0^t V_k^n(s) ds \right) \right) - q_{kj} M_k \left( \eta_k \int_0^t V_k^n(s) ds \right) \right] \\ &\quad + \sum_{k=1}^K q_{kj} n^{-1/2} M_k \left( \eta_k \int_0^t V_k^n(s) ds \right) - n^{-1/2} \left[ N_j \left( n\mu_j T_j^n(t) \right) - n\mu_j T_j^n(t) \right] \\ &\quad - n^{-1/2} n\mu_j T_j^n(t) \\ &= \widehat{Q}_j^n(0) + \sum_{k=1}^K n^{-1/2} \left[ \Psi_{kj} \left( nn^{-1} M_k \left( n\eta_k \int_0^t n^{-1} V_k^n(s) ds \right) \right) \right. \\ &\quad \left. - q_{kj} nn^{-1} M_k \left( n\eta_k \int_0^t n^{-1} V_k^n(s) ds \right) \right] \\ &\quad + \sum_{k=1}^K q_{kj} n^{-1/2} \left[ M_k \left( n\eta_k \int_0^t n^{-1} V_k^n(s) ds \right) - n\eta_k \int_0^t n^{-1} V_k^n(s) ds \right] \\ &\quad + \sum_{k=1}^K q_{kj} n^{-1/2} \eta_k \int_0^t V_k^n(s) ds \end{aligned}$$

$$\begin{aligned}
& - n^{-1/2} \left[ N_j \left( n\mu_j T_j^n(t) \right) - n\mu_j T_j^n(t) \right] - \sqrt{n}\mu_j \left( t - I_j^n(t) \right) \\
= & \widehat{Q}_j^n(0) + \sum_{k=1}^K n^{-1/2} \left[ \Psi_{kj} \left( n\bar{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) \right) - nq_{kj} \bar{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) \right] \\
& + \sum_{k=1}^K q_{kj} n^{-1/2} \left[ M_k \left( n\eta_k \int_0^t \bar{V}_k^n(s) ds \right) - n\eta_k \int_0^t \bar{V}_k^n(s) ds \right] \\
& + \sum_{k=1}^K q_{kj} n^{-1/2} \eta_k \int_0^t \left[ \sqrt{n} \widehat{V}_k^n(s) + nm_k \right] ds \\
& - n^{-1/2} \left[ N_j \left( n\mu_j T_j^n(t) \right) - n\mu_j T_j^n(t) \right] - \sqrt{n}\mu_j t - \sqrt{n}\mu_j I_j^n(t) \\
= & \widehat{Q}_j^n(0) + \sum_{k=1}^K \widehat{\Psi}_{kj}^n \left( \bar{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) \right) + \sum_{k=1}^K q_{kj} \widehat{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) \\
& - \widehat{N}_j^n \left( \mu_j T_j^n(t) \right) + \sum_{k=1}^K q_{kj} \eta_k \int_0^t \widehat{V}_k^n(s) ds - \mu_j \widehat{I}_j^n(t) + t\sqrt{n} \left[ \sum_{k=1}^K q_{kj} \eta_k m_k - \mu_j \right] \\
= & \widehat{\xi}_j^n(t) + \sum_{k=1}^K q_{kj} \eta_k \int_0^t \widehat{V}_k^n(s) ds - \mu_j \widehat{I}_j^n(t),
\end{aligned}$$

which is (1.40). Note that the final equality in the above calculation follows from (2.51), the definition of  $m_k$  (given in Section 1.4), and the heavy traffic condition in (1.21).

Next we derive (1.41), the diffusion-scaled system equation for the infinite-server stations.

By (1.9) and (1.11)–(1.13), we have

$$\begin{aligned}
V_k^n(t) &= V_k^n(0) + E_k^n(t) - F_k^n(t) \\
&= V_k^n(0) + \sum_{j=1}^J \Phi_{jk} \left( N_j \left( n\mu_j T_j^n(t) \right) \right) - M_k \left( \eta_k \int_0^t V_k^n(s) ds \right). \tag{1.171}
\end{aligned}$$

As before, elementary algebraic manipulations and the appropriate scaling in (1.23)–(1.35) applied to (1.171) yields

$$\begin{aligned}
\widehat{V}_k^n(t) &= n^{-1/2} V_k^n(0) + \sum_{j=1}^J n^{-1/2} \Phi_{jk} \left( N_j \left( n\mu_j T_j^n(t) \right) \right) - n^{-1/2} M_k \left( \eta_k \int_0^t V_k^n(s) ds \right) \\
&= \widehat{V}_k^n(0) + \sum_{j=1}^J n^{-1/2} \left[ \Phi_{jk} \left( N_j \left( n\mu_j T_j^n(t) \right) \right) - p_{jk} N_j \left( n\mu_j T_j^n(t) \right) \right]
\end{aligned}$$

$$\begin{aligned}
& + \sum_{j=1}^J n^{-1/2} p_{jk} N_j(n\mu_j T_j^n(t)) - n^{-1/2} \left[ M_k \left( \eta_k \int_0^t V_k^n(s) ds \right) - \eta_k \int_0^t V_k^n(s) ds \right] \\
& - n^{-1/2} \eta_k \int_0^t V_k^n(s) ds \\
= & \widehat{V}_k^n(0) + \sum_{j=1}^J n^{-1/2} \left[ \Phi_{jk} \left( n n^{-1} N_j(n\mu_j T_j^n(t)) \right) - p_{jk} n n^{-1} N_j(n\mu_j T_j^n(t)) \right] \\
& + \sum_{j=1}^J n^{-1/2} p_{jk} \left[ N_j(n\mu_j T_j^n(t)) - n\mu_j T_j^n(t) \right] + \sum_{j=1}^J n^{-1/2} p_{jk} n\mu_j T_j^n(t) \\
& - n^{-1/2} \left[ M_k \left( n\eta_k \int_0^t n^{-1} V_k^n(s) ds \right) - n\eta_k \int_0^t n^{-1} V_k^n(s) ds \right] \\
& - n^{-1/2} \eta_k \int_0^t \left[ \sqrt{n} \widehat{V}_k^n(s) + n m_k \right] ds \\
= & \widehat{V}_k^n(0) + \sum_{j=1}^J n^{-1/2} \left[ \Phi_{jk} \left( n \bar{N}_j^n(\mu_j T_j^n(t)) \right) - p_{jk} n \bar{N}_j(\mu_j T_j^n(t)) \right] \\
& + \sum_{j=1}^J p_{jk} \widehat{N}_j^n(\mu_j T_j^n(t)) + \sum_{j=1}^J \sqrt{n} p_{jk} \mu_j (t - I_j^n(t)) \\
& - n^{-1/2} \left[ M_k \left( n\eta_k \int_0^t \bar{V}_k^n(s) ds \right) - n\eta_k \int_0^t \bar{V}_k^n(s) ds \right] \\
& - \eta_k \int_0^t \widehat{V}_k^n(s) ds - t\sqrt{n} \eta_k m_k \\
= & \widehat{V}_k^n(0) + \sum_{j=1}^J \widehat{\Phi}_{jk}^n \left( \bar{N}_j^n(\mu_j T_j^n(t)) \right) + \sum_{j=1}^J p_{jk} \widehat{N}_j^n(\mu_j T_j^n(t)) - \widehat{M}_k^n \left( \eta_k \int_0^t \bar{V}_k^n(s) ds \right) \\
& - \eta_k \int_0^t \widehat{V}_k^n(s) ds - \sum_{j=1}^J p_{jk} \mu_j \widehat{I}_j^n(t) + t\sqrt{n} \left[ \sum_{j=1}^J p_{jk} \mu_j - \eta_k m_k \right] \\
= & \widehat{\zeta}_k^n(t) - \eta_k \int_0^t \widehat{V}_k^n(s) ds - \sum_{j=1}^J p_{jk} \mu_j \widehat{I}_j^n(t),
\end{aligned}$$

which is (1.41). Note that the final equality in the above calculation follows from (1.39) and the definition of  $m_k$ .

# CHAPTER 2

## A QUEUEING MODEL OF DYNAMIC PRICING AND DISPATCH CONTROL FOR RIDE-HAILING SYSTEMS INCORPORATING TRAVEL TIMES

### 2.1 Introduction

This paper studies a dynamic control problem for a queueing model motivated by taxi and ride hailing systems. In those systems, customers and drivers can be matched centrally by a platform using web or mobile applications. In addition, the platform can adjust the prices dynamically over time. We consider a city partitioned into a set of geographical regions. Each such region should be thought of as a pick-up or drop-off location. Simultaneously, cars reside in these regions waiting to pick up customers. We use a queueing model to study this problem, following a growing number of papers in the operations research literature. However, much of the relevant literature assumes away the travel times between the pick-up and drop-off locations; see, for example, Ata et al. [2020a] and the references therein. A key novelty of our model is that it incorporates travel times, but this leads to a significantly more challenging analysis.

We assume that the platform, also referred to as the system manager hereafter, has two levers: pricing and dispatch controls. She seeks an effective policy that makes both dynamic pricing and dynamic dispatch control decisions in order to maximize the long-run average profit. We allow the prices to depend on time and the customer location. Dynamically adjusting prices elicits two competing effects. On the one hand, increasing prices increase the per-ride revenue for the platform. On the other hand, customers are price sensitive, so higher prices result in lower customer demand. Dispatching refers to the process of matching a car with a customer requesting a ride and constitutes an important operational decision for the platform.

We model a ride-hailing or taxi system as a closed queueing network with a fixed number of jobs, denoted by  $n$ . There are  $I$  buffers,  $I$  single-server nodes, and an infinite-server node in the SPN. The terms “server” and “resource” will be used interchangeably to refer to a single-server node. Similarly, the terms “buffer” and “class” will be used interchangeably. As such, jobs in buffer  $i$  will be referred to as class  $i$  jobs, for  $i = 1, \dots, I$ . In addition to choosing prices dynamically, the system manager can engage in  $J$  possible (dispatch) activities, where each activity corresponds to a server serving jobs in a buffer. Following service at a single-server node, jobs are routed to the infinite-server node. Jobs then continue their service at the infinite-server node, after which they are probabilistically routed back to the buffers. The infinite-server node models the travel times. This process continues indefinitely.

In the context of our motivating application, jobs correspond to cars that circulate in the system perpetually. The  $I$  buffers correspond to  $I$  city regions where cars wait to get matched with a customer. In addition, the service rates at a single-server node can be thought of as the customer arrival rate to the corresponding region, which depends on the price. As a result, customer demand dynamically changes over time as the platform varies the prices of rides. An activity corresponds to dispatching a car from one region to serving an arriving customer possibly in another region. Thus, a service completion at a single-server node corresponds to a car getting matched with a customer. After getting matched with a customer, the car must travel to pick up the customer and bring him to his destination. We assume that all customer requests that are not met immediately are lost. In the queueing model, this corresponds to jobs getting routed to and served at the infinite-server node. That is, the infinite-server node models the travel time of a car from its initial dispatch time to the drop off time of the customer. Upon completing service at the infinite server node, the job is routed to the buffer that is associated with the customer’s destination. This is modeled through a probabilistic routing structure as is usually done in the queueing literature. Although the

SPN we study is motivated by the ride-hailing and taxi systems, in what follows we use the queueing terminology that is standard in the literature. However, we will occasionally make reference to our motivating applications when intuition or interpretation are needed.

As mentioned above, incorporating travel times makes the problem significantly more challenging. To ease the analysis, we assume there is a single travel node. This assumption has two implications: First, the travel times between any two regions have the same distribution. Second, upon completing service at the infinite-server node all job classes share the same probabilistic routing structure. Admittedly, this is a restrictive assumption, but it simplifies the analysis and allows us to incorporate the travel times into the model. We view our model as an important first step in the analysis of ride-sharing network models that incorporate travel times.

However, even under the single travel node assumption, the problem is not amenable to exact analysis. As such, we consider a diffusion approximation to it in the heavy traffic asymptotic regime. In that regime, under the so called complete resource pooling condition, see Harrison and López [1999], we solve the problem analytically and derive a closed-form solution for the optimal dynamic prices.

Notwithstanding these restrictive assumptions, the paper makes two contributions. First, it incorporates the travel times in the model and solves the resulting dynamic pricing and dispatch control problem analytically in the heavy traffic regime. Second, it makes a methodological contribution by solving a drift-rate control problem on an unbounded domain, which could be of interest in its own right.

The rest of the paper is structured as follows. Section 2.2 reviews the literature. Section 2.3 presents the control problem for the ride-hailing platform, and the associated Brownian control problem is derived formally in Section 2.4. The equivalent workload formulation is formulated in Section 2.5 and it is solved in Sections 2.6 and 2.7 by studying a related Bellman equation. Section 2.8 interprets the solution of the equivalent workload formulation in the

context of the original control problem and proposes a pricing and dispatch policy. Section 2.9 conducts a simulation study to illustrate the effectiveness of the proposed policy. Section 2.10 concludes the paper. There are two appendices: Appendix 2.11.1 provides a formal derivation of the Brownian control problem, and Appendix 2.11.2 contains all miscellaneous proofs.

## 2.2 Literature Review

Our paper is related to two streams of literature: the modeling and analysis of ride-hailing and taxi systems and the dynamic control of queueing networks.

In recent years several authors have modeled ride-hailing and taxi systems using queueing networks. A majority of this literature has focused on how pricing, dispatch (matching), and relocation decisions can improve system performance. From a modeling perspective, Ata et al. [2020a] and Braverman et al. [2019] are most closely related to ours. Ata et al. [2020a] model a ride-hailing system closed stochastic processing network with dispatch and relocation control. Under heavy traffic conditions, they approximate the original control problem by a Brownian control problem (BCP). After reducing the BCP to an equivalent workload formulation, they propose an algorithm to solve it numerically. However, their model does not include travel times, whereas ours does. Incorporating travel times leads to a significantly more challenging problem in the heavy traffic limit under the diffusion scaling. On the other hand, Braverman et al. [2019] model a ride-hailing system as a closed queueing networks with travel times and relocation control. By solving a suitable linear program, they propose a static routing policy and prove that it is asymptotically optimal in a large market asymptotic regime under fluid scaling. Hosseini et al. [2021] extends the analysis of Braverman et al. [2019] by designing a dynamic relocation that outperforms the asymptotically optimal static policy in realistic problem instances. In a related study, Zhang and Pavone [2016] uses a combination of single-server and infinite-server queueing model



to study the control of autonomous vehicles. The authors derive an open loop policy by solving a linear program. Building on this solution, they also propose an effective dynamic rebalancing policy.

Several other papers are at the intersection of ride-hailing and queueing, but differ more in their modeling choices and analysis. Banerjee et al. [2016] study pricing on a single-region model with a single travel time node and show that an optimal static pricing policy performs well. Banerjee et al. [2022] develop an approximation framework to study vehicle sharing systems under pricing, matching, and repositioning policies for several objective functions and under various system constraints. In particular, they develop algorithms and show that the approximation ratio of the resulting policy improves as the number of cars in each region grows. Banerjee et al. [2020] study matching for a general closed queueing network that can be used to model ride-hailing systems. They propose a family of state-dependent matching policies that do not use any demand arrival rate information. Under a complete resource pooling assumption, they show that the proportion of dropped demand under any such policy decays exponentially as the number of supply units in the network grows. Afèche et al. [2023] develop a game-theoretic fluid model to study admission control and repositioning in a ride-hailing system with strategic drivers. Their analysis provides insights into spatial demand imbalances and how demand admission control can impact the strategic behavior of drivers in the network. Afèche et al. [2023] studies the optimal dynamic pricing and dispatch control under demand shocks. Özkan and Ward [2020] model a ride-hailing system as an open queueing network model with impatient customers. They propose a matching policy and prove asymptotically optimality in the fluid scale in a large market regime. Özkan [2020] studies a fluid model with strategic drivers that incorporates both pricing and matching decisions, highlighting the importance of looking at multiple controls simultaneously. Besbes et al. [2022] study the effect of pick up and travel times on capacity planning for a ride-hailing system by modeling it as a spatial multi-server queue. Chen et al. [2020] proposes static

and dynamic policies that are asymptotically optimal. Varma et al. [2022] studies an open network model and proposes an asymptotically optimal policy. Examples of other papers that use spatial models for pricing include Yang et al. [2018], Jacob and Roet-Green [2021], and Hu et al. [2022].

Several other researchers focused on different aspects of the ride-hailing and taxi systems without using queueing theoretic models. These include Wang et al. [2017], Ata et al. [2020b], Bertsimas et al. [2019], Besbes et al. [2021], Bimpikis et al. [2019], Cachon et al. [2017], Castillo et al. [2018], Chen and Sheldon [2016], Garg and Nazerzadeh [2021], Gokpinar and Selcuk [2019], Guda and Subramanian [2019b], He et al. [2020], Hu et al. [2022], Hu and Zhou [2021], Yan et al. [2020], and Lu et al. [2018].

This paper also contributes to the broader literature on dynamic control of queueing systems. Two prominent approaches in that literature are: (i) Markov decision process (MDP) formulations, and (ii) heavy traffic approximations. Intuitively, the workload problem studied in Sections 2.5–2.7 relates to the service rate and admission control problems studied using MDP formulations, see for example Jr. and Weber [1989] and references therein. The most closely related papers are George and Harrison [2001] and Ata [2005]. These papers study the service rate control problems for an  $M/M/1$  queue and provide closed-form solutions; also see Ata and Shneorson [2006], Ata and Zachariadis [2007], Adusumilli and Hasenbein [2010], and Kumar et al. [2013].

The second approach is pioneered by Harrison [1988], also see Harrison [2000, 2003]. In particular, a number of papers studied drift rate control problems for one-dimensional diffusions arising under heavy traffic approximations, see Ata et al. [2005], Ata [2006], Ghosh and Weerasinghe [2007, 2010], Rubino and Ata [2009], Kim and Ward [2013], and Ata and Tongarlak [2013]. More recently, Ata and Barjesteh [2023] and Ata et al. [2024] studied drift-rate control problems arising in different contexts such as volunteer capacity management and make-to-stock manufacturing. The analysis of the drift-rate control problem solved in

this paper differs significantly from the analysis in those papers because it involves a quadratic cost of drift rate, unbounded set of feasible drift rates, and an unbounded state space. The combination of these features lead to a more challenging analysis. Our paper also makes a modeling contribution by formulating the dynamic dispatch and pricing control problem that incorporates travel times. Furthermore, it proposes an analytically tractable approximation in the heavy traffic limit and solves that in closed form.

Lastly, our paper draws on the literature of the asymptotic analysis of closed queueing networks with infinite-server queues, see for example Kogan et al. [1986], Smorodinskii [1986], Kogan and Lipster [1993], and Krichagina and Puhalskii [1997].

## 2.3 Ridesharing Model

Motivated by the taxi and ride-hailing application described in the introduction, we consider a closed queueing network with  $n$  jobs,  $I$  buffers,  $I$  single-server nodes, and one infinite-server node. Figure 2.1 displays an illustrative network with  $I = 4$  and  $J = 10$ , also see Section 2.9 for the motivation behind this example.

As mentioned earlier, in addition to dynamic pricing decisions, the system manager also makes dispatch decisions dynamically. There are  $J$  dispatch activities she can choose from. Each dispatch activity involves a unique buffer and a unique server—we use the terms single-server node and server interchangeably. Let  $s(j)$  and  $b(j)$  denote the server and the buffer, respectively, associated with activity  $j$  for  $j = 1, \dots, J$ . In other words, activity  $j$  is undertaken by server  $s(j)$  and it servers jobs in buffer  $b(j)$ . We describe the association between activities and resources by the capacity consumption matrix  $A$  and the association between activities and buffers by the constituency matrix  $C$ . That is,  $A$  is the  $I \times J$  matrix




Figure 2.1: A network with four regions and ten dispatch activities. The open rectangles are the buffers, the circles are the single servers, and the oval is an infinite-server node. The ten activities are represented by the arrows between the buffers and servers. The numbers on the arrows indicate their index. Activities 1, 2, 3, and 4 are local dispatch activities while activities 5 through 10 are non-local dispatch activities. The arrows from the infinite-server to the buffers represent probabilistic rerouting of jobs in the network.

given by

$$A_{ij} := \begin{cases} 1, & \text{if } s(j) = i, \\ 0, & \text{otherwise,} \end{cases} \quad (2.1)$$

and  $C$  is the  $I \times J$  matrix given by

$$C_{ij} := \begin{cases} 1, & \text{if } b(j) = i, \\ 0, & \text{otherwise.} \end{cases} \quad (2.2)$$

Let  $\mathcal{A}_i$  denote the set of activities server  $i$  undertakes. Similarly, let  $\mathcal{C}_i$  denote the set of

activities that serve buffer  $i$ . In other words, we have that

$$\mathcal{A}_i := \{j : A_{ij} = 1\}, \quad (2.3)$$

$$\mathcal{C}_i := \{j : C_{ij} = 1\}. \quad (2.4)$$

For each activity  $j = 1, \dots, J$ , we associate a unit rate Poisson process  $N_j$ . We also associate a unit rate Poisson process  $N_0$  with the infinite-server node. The processes  $N_0, N_1, \dots, N_J$  are mutually independent. The service rate at the infinite-server node is denoted by  $\eta > 0$ . We denote the service rate for activity  $j$  at time  $t$  by  $\mu_j(t)$  for  $t \geq 0$  and  $j = 1, \dots, J$ . The system manager chooses prices  $p(t) = (p_i(t))$  dynamically over time, where  $p_i(t)$  denotes the price charged to customers who seek rides from region  $i$  at time  $t$ . As the reader will see below, these prices ultimately determine activity service rates  $\mu_j(t)$  for  $j = 1, \dots, J$  and  $t \geq 0$ . We assume  $p_i(t) \in [\underline{p}_i, \bar{p}_i]$  for  $t \geq 0$ , where  $0 \leq \underline{p}_i < \bar{p}_i < \infty$ . The price sensitivity of demand is captured by a nonnegative demand function  $\Lambda : \mathcal{P} \rightarrow \mathbb{R}_+^I$ , where  $\mathcal{P} = \prod_{i=1}^I [\underline{p}_i, \bar{p}_i]$ . Namely, the demand rate vector at time  $t$ , denoted by  $\lambda(t)$ , is given by<sup>1</sup>

$$\lambda(t) = \Lambda(p(t)) = (\Lambda_1(p_1(t)), \dots, \Lambda_I(p_I(t)))', \quad t \geq 0. \quad (2.5)$$

We make the following monotonicity assumption to simplify the analysis:

**Assumption 3.** The demand rate function is strictly decreasing in price, i.e.,  $\Lambda_i(p_i)$  is strictly decreasing in  $p_i$  for  $i = 1, \dots, I$ .

From this monotonicity assumption, it follows that  $\Lambda_i(\cdot)$  has an inverse function, denoted by  $\Lambda_i^{-1}(\cdot)$ . Moreover, the pricing decisions can be replaced with choosing the demand rate vector  $\lambda(t)$  dynamically over time. This is convenient for our analysis. In order to proceed

---

1. The customer demand rate in region  $i$ ,  $\lambda_i(t)$ , depends only on the price  $p_i(t)$ .

with that approach, we define the set of admissible demand rate vectors  $\mathcal{L} \subseteq \mathbb{R}_+^I$  as follows:

$$\mathcal{L} := \prod_{i=1}^I \mathcal{L}_i, \quad (2.6)$$

where  $\mathcal{L}_i := [\Lambda_i(\bar{p}_i), \Lambda_i(\underline{p}_i)]$  for  $i = 1, \dots, I$ . Denoting  $\Lambda^{-1}(x) = (\Lambda_1^{-1}(x_1), \dots, \Lambda_I^{-1}(x_I))'$  for  $x \in \mathcal{L}$ , it is easy to see that  $\Lambda^{-1}$  is the inverse function of  $\Lambda$ . Viewing the demand rates as the platform's pricing control, we define the revenue rate function  $\pi : \mathcal{L} \rightarrow \mathbb{R}$  as follows:

$$\pi(x) := \sum_{i=1}^I x_i \Lambda_i^{-1}(x_i), \quad x \in \mathcal{L}. \quad (2.7)$$

We also make the following regularity assumptions for the revenue rate function:

**Assumption 4.** The revenue rate function  $\pi$  is: (a) three-times continuously differentiable and strictly concave on  $\mathcal{L}$ , and (b) has a maximizer in the interior of  $\mathcal{L}$ .

Upon completing service at a single-server node, each job goes next to the infinite-server node. Once its service there is complete, the job next joins buffer  $i$  with probability  $q_i > 0$  for  $i = 1, \dots, I$  where  $\sum_{i=1}^I q_i = 1$ . The routing probability vector  $q = (q_i)$  does not depend on the single-server node the job departed from prior to joining the infinite-server node. In other words, customers' destination distribution is identical across different origins. This is a restrictive assumption, but it simplifies the analysis significantly and enables us to incorporate travel times into the model. As discussed in the Introduction, we view this as an important first step in the analysis of ride-sharing network models that incorporate travel times. In order to model this probabilistic routing structure mathematically, we let  $\psi = \{\psi(l), l \geq 1\}$  denote a sequence of  $I$ -dimensional i.i.d. random vectors with  $P(\psi(1) = e_i) = q_i$  for  $i = 1, \dots, I$ , where  $e_i$  is an  $I$ -dimensional vector with one in the  $i$ th component and zeros elsewhere. Then letting

$$\Psi(m) := \sum_{l=1}^m \psi(l), \quad m \geq 1, \quad (2.8)$$

we note that the  $i$ th component of  $\Psi(m)$ , denoted by  $\Psi_i(m)$ , represents the total number of jobs routed to buffer  $i$  among the first  $m$  jobs that have finished service at the infinite-server node.

As discussed earlier, there are two types of control decisions that the system manager must make. First, she must choose an  $I$ -dimensional demand rate process  $\lambda = \{\lambda(t), t \geq 0\}$ . This is equivalent to making dynamic pricing decisions. Recall that the customer arrival process at single-server node  $i$  corresponds to its service process. Because these customers can be transported by cars in regions corresponding to activities  $j \in \mathcal{A}_i$ , we let

$$\mu_j(t) := \lambda_i(t) \quad \text{for } j \in \mathcal{A}_i, \quad i = 1, \dots, I, \quad \text{and } t \geq 0. \quad (2.9)$$

This defines the  $J$ -dimensional service rate process  $\mu = \{\mu(t), t \geq 0\}$ , where  $\mu(t) = (\mu_j(t))$ . Second, she must decide on how servers allocate their time to various (dispatch) activities. This decision takes the form of cumulative allocation processes  $T_j = \{T_j(t), t \geq 0\}$  for  $j = 1, \dots, J$ . In particular,  $T_j(t)$  represents the cumulative amount of time server  $s(j)$  devotes to activity  $j$  (serving class  $i(j)$  jobs) during  $[0, t]$ .

Next, we introduce the system dynamics equations that govern the movement of jobs in the network. To that end, we let  $Q_0(t)$  and  $Q_i(t)$  denote the number of jobs in the infinite-server node and in buffer  $i$  at time  $t$ , respectively, for  $i = 1, \dots, I$ . We also let  $A_0(t)$  and  $A_i(t)$  be the total number of jobs that have arrived to the infinite-server node and to buffer  $i$  by time  $t$ , respectively, for  $i = 1, \dots, I$ . Then we have that

$$A_0(t) := \sum_{j=1}^J N_j \left( \int_0^t \mu_j(s) dT_j(s) \right), \quad t \geq 0, \quad (2.10)$$

$$A_i(t) := \Psi_i \left( N_0 \left( \eta \int_0^t Q_0(s) ds \right) \right), \quad t \geq 0. \quad (2.11)$$

Moreover, letting  $D_0(t)$  and  $D_i(t)$  denote the total number of jobs that have left the infinite-

server node and buffer  $i$  by time  $t$ , respectively, for  $i = 1, \dots, I$ , we have that

$$D_0(t) := N_0 \left( \eta \int_0^t Q_0(s) ds \right), \quad t \geq 0, \quad (2.12)$$

$$D_i(t) := \sum_{j \in \mathcal{C}_i} N_j \left( \int_0^t \mu_j(s) dT_j(s) \right), \quad t \geq 0. \quad (2.13)$$

We refer to the  $(I+1)$ -dimensional process  $Q = (Q_0, Q_1, \dots, Q_I)'$  as the queue length process, whose dynamics is given next:

$$Q_i(t) = Q_i(0) + A_i(t) - D_i(t), \quad i = 0, 1, \dots, I \quad \text{and} \quad t \geq 0, \quad (2.14)$$

where  $Q(0)$  is the vector of initial queue lengths such that  $\sum_{i=0}^I Q_i(0) = n$ . Letting  $I_i(t)$  denote the cumulative amount of time that server  $i$  is idle during the interval  $[0, t]$  for  $i = 1, \dots, I$ , we have that

$$I_i(t) := t - \sum_{j \in \mathcal{A}_i} T_j(t), \quad t \geq 0, \quad (2.15)$$

or in matrix notation,  $I(t) = et - AT(t)$  for  $t \geq 0$ . Note that (2.10)–(2.14) imply that

$$\sum_{i=0}^I Q_i(t) = \sum_{i=0}^I Q_i(0) = n, \quad t \geq 0,$$

expressing the fact that the total number of jobs in the system remains fixed in a closed network.

In order to state the platform's objective and its control problem formally, we introduce two vectors of cost parameters  $h = (h_0, h_1, \dots, h_I)' \in \mathbb{R}_+^{I+1}$  and  $c = (c_1, \dots, c_I)' \in \mathbb{R}_+^I$ . In the context of the ride-hailing system, the platform incurs a fuel cost at a rate of  $h_0$  per traveling car. Moreover, for  $i = 1, \dots, I$ , there is a holding cost at a rate of  $h_i$  for each car waiting for a ride in region  $i$ , reflecting the fact that no driver likes sitting idle. We assume that  $h_i > h_0$



for all  $i = 1, \dots, I$ . Finally, for  $i = 1, \dots, I$ , there is an idleness cost at the rate of  $c_i$  per unit of time server  $i$  is idle. This represents the lost revenue from picking up customers arriving to region  $i$  and goodwill loss.<sup>2</sup> A control policy is denoted by  $(T, \lambda)$  and must satisfy the following conditions:

$$T, \lambda \text{ are nonanticipating with respect to } Q, \quad (2.16)$$

$$T, I \text{ are nondecreasing and continuous with } T(0) = I(0) = 0, \quad (2.17)$$

$$\lambda(t) \in \mathcal{L} \text{ for all } t \geq 0, \quad (2.18)$$

$$Q_i(t) \geq 0 \text{ for all } t \geq 0 \text{ and } i = 0, 1, \dots, I. \quad (2.19)$$

Equation (2.16) expresses the fact that the policy can only depend on observable quantities. Equation (2.17) is natural given the interpretations of the processes  $T$  and  $I$ . Equation (2.18) requires that  $\lambda$  come from the set of achievable demand rates. Equation (2.19) expresses the fact that queue lengths are nonnegative. The arriving customer demand is allocated to cars waiting in various buffers through the dispatch activities  $j = 1, \dots, J$  (see, for example, (2.10) and (2.13)). Given a control policy  $(T, \lambda)$ , we define the cumulative profit collected up to time  $t$  as

$$V(t) := \int_0^t [\pi(\lambda(s)) - h'Q(s)] ds - c'I(t), \quad t \geq 0. \quad (2.20)$$

The platform's control problem is to choose a policy  $(T, \lambda)$  so as to

$$\text{maximize} \quad \liminf_{t \rightarrow \infty} \frac{1}{t} E[V(t)] \quad (2.21)$$

$$\text{subject to} \quad (2.10)\text{--}(2.20). \quad (2.22)$$

Because control problem (2.21)–(2.22) in its original form is not amenable to exact analysis,

---

2. One can assume  $c_i \geq p_i^* = \Lambda_i^{-1}(\lambda_i^*)$  naturally, where  $\lambda^*$  is defined in (2.23) below.

the next section considers a related control problem in an asymptotic regime where the number of cars gets large and derives the approximating Brownian control problem. The Brownian control problem is an approximation to the original problem, yet it is far more tractable.

## 2.4 Brownian Control Problem

Following an approach that is similar to the one taken in Harrison [1988], this section develops a Brownian approximation to the control problem presented in Section 2.3. Many authors have proved heavy traffic limit theorems to rigorously justify such Brownian approximations—see for example Harrison [1998], Williams [1998], Kumar [2000], Bramson and Dai [2001], Stolyar [2004], Bell and Williams [2001, 2005], Ata and Kumar [2005], Ata and Olsen [2009, 2013] and references therein. We do not attempt to prove a rigorous convergence theorem in this paper, but refer the reader to Harrison [1988, 2000, 2003] for elaborate and intuitive justifications of the approximation procedure we follow.

The approximation procedure starts by solving the following static pricing problem (existence of the optimal solution is guaranteed by Assumption 4), which helps us articulate the heavy traffic assumption that underlies the mathematical development to follow. We set

$$\lambda^* := \operatorname{argmax}_{\lambda \in \mathcal{L}} \pi(\lambda). \quad (2.23)$$

Recall from Assumption 4 that we assume  $\lambda^*$  is in the interior of  $\mathcal{L}$ , i.e.,  $\lambda^* \in \operatorname{int}(\mathcal{L})$ . The vector  $\lambda^*$  represents the average demand rates that would result in the largest revenue rate ignoring variability in the system. Note that the corresponding nominal service rates for the

various activities are given by<sup>3</sup>

$$\mu_j^* := \lambda_i^* \quad \text{for } j \in \mathcal{A}_i. \quad (2.24)$$

Using these nominal service rates, we define an  $I \times J$  input-output matrix  $R$  as follows:

$$R_{ij} := \mu_j^* C_{ij}, \quad i = 1, \dots, I, \quad j = 1, \dots, J. \quad (2.25)$$

Following Harrison [1988, 2000], we interpret  $R_{ij}$  as the long-run average rate of class  $i$  material consumed per unit of activity  $j$  under the nominal service rates  $\mu_j^*$  for  $j = 1, \dots, J$ .

We also define the  $I$ -dimensional input vector  $\nu$  as

$$\nu_i := q_i \eta, \quad i = 1, \dots, I. \quad (2.26)$$

We interpret  $\nu_i$  as the long-run average rate of input into buffer  $i$  from the infinite-server node. As a preliminary to stating the heavy traffic assumption, we introduce the notion of local activities. In the context of the motivating application, it corresponds to a customer in a region being picked up by a car in the same region. Using the terminology that is standard in queueing theory, it corresponds to a server processing its own buffer. Without loss of generality, we assume that the first  $I$  activities are local. That is,

$$s(j) = b(j) = j \quad \text{for } j = 1, \dots, I.$$

This is equivalent to assuming that the first  $I$  columns of matrices  $A$  and  $C$  constitute the  $I \times I$  dimensional identity matrix. The following is the heavy traffic assumption:

---

3. In particular, for all  $j$ ,  $\mu_j^* = \sum_{i=1}^I \lambda_i^* A_{ij}$ . This is true because there exists only one  $i$  such that  $A_{ij} \neq 0$  for each  $j = 1, \dots, J$ . That is, an activity only uses one server. In matrix notation,  $\mu^* = A' \lambda^*$ , where  $A'$  is the transpose of  $A$ .

**Assumption 5.** There exists a unique  $x^* \in \mathbb{R}^J$  such that

$$x_j^* = \min \left\{ 1, \frac{\nu_j}{\lambda_j^*} \right\}, \quad j = 1, \dots, I, \quad (2.27)$$

$$Rx^* = \nu, \quad (2.28)$$

$$Ax^* = e, \quad (2.29)$$

$$x^* \geq 0. \quad (2.30)$$

The vector  $x^*$  is referred to as the nominal processing plan and the component  $x_j^*$  can be interpreted as the long-run average rate at which activity  $j$  is undertaken. Equation (2.30) says that all nominal activity levels must be non-negative. Equation (2.29) means that under the nominal processing plan, servers are fully utilized. Equation (2.28) is a flow balance condition which says that the rate of jobs leaving the buffers equals the rate of jobs entering the buffers under the nominal processing plan. Note that by (2.25)–(2.28) we have  $\sum_{j \in \mathcal{C}_i} \mu_j^* x_j = q_i \eta$  for each  $i$ , which then implies that  $(\mu^*)'x^* = \eta$  by summing over  $i$ . We interpret  $(\mu^*)'x^*$  as the rate of jobs entering the infinite-server node under the nominal processing plan, and Assumption 5 ensures that this equals the service rate at the infinite-server node.<sup>4</sup> Equation (2.27) ensures that local activities are used at maximal rates. In the context of the motivating application, this means customer demand is met by cars in the same region as much as possible.

Following Harrison [2000], we call activity  $j$  basic if  $x_j^* > 0$ , whereas it is called nonbasic if  $x_j^* = 0$ . We let  $b$  denote the number basic activities. After possibly relabeling, we assume without loss of generality that activities  $1, \dots, b$  are basic and that activities  $b+1, \dots, J$  are nonbasic. Recall that the first  $I$  of them are the local activities. As done in Harrison [2000],

---

4. Based on intuition from the classical  $M/M/\infty$  queue, this condition implies that the steady-state fraction of jobs in the infinite-server node under the nominal processing plan is equal to one as the number of jobs in the system grows, i.e. as  $n \rightarrow \infty$ .

we partition the matrices  $R$  and  $A$  as follows:

$$R = [H \ K] \quad \text{and} \quad A = [B \ N], \quad (2.31)$$

where  $H, B \in \mathbb{R}^{I \times b}$  and  $K, N \in \mathbb{R}^{I \times (J-b)}$ . The submatrices  $H$  and  $B$  correspond to the basic activities of  $R$  and  $A$ , respectively, while the submatrices  $K$  and  $N$  correspond to the nonbasic activities.

In order to derive the approximating Brownian control problem, we consider a sequence of closely related systems indexed by the total number of jobs  $n$ . The formal limit of this sequence as  $n \rightarrow \infty$  is the approximating Brownian control problem. We attach a superscript of  $n$  to quantities associated with the  $n$ th system in the sequence. To be specific, we define the scaled demand rate function  $\Lambda^n : \mathcal{P} \rightarrow \mathbb{R}_+^I$  by

$$\Lambda^n(x) := n\Lambda(x), \quad x \in \mathcal{P}. \quad (2.32)$$

Then we define the set of admissible scaled demand rate vectors  $\mathcal{L}^n$  as the following:

$$\mathcal{L}^n := \left\{ \lambda^n \in \mathbb{R}_+^I : \lambda^n = \Lambda^n(p) \text{ for some } p \in \mathcal{P} \right\}. \quad (2.33)$$

We note from (2.5)–(2.6) and (2.32)–(2.33) that  $\mathcal{L}^n = n\mathcal{L}$ , and that  $\Lambda^n$  has the inverse function  $(\Lambda^n)^{-1}(x) = ((\Lambda_1^n)^{-1}(x_1), \dots, (\Lambda_I^n)^{-1}(x_I))'$  for  $x \in \mathcal{L}^n$ . We define the scaled revenue rate function  $\pi^n$  as follows:

$$\pi^n(x) := \sum_{i=1}^I x_i (\Lambda_i^n)^{-1}(x_i), \quad x \in \mathcal{L}^n. \quad (2.34)$$

Observing that  $nx \in \mathcal{L}^n$  if and only if  $x \in \mathcal{L}$ , it can equivalently be shown that<sup>5</sup>

$$\pi^n(nx) = n\pi(x) = n \sum_{i=1}^I x_i \Lambda_i^{-1}(x_i), \quad x \in \mathcal{L}. \quad (2.35)$$

Therefore, in the  $n$ th system, the revenue rate process is simply scaled by  $n$ . We also scale the holding cost rates  $h^n$  and the idleness cost rates  $c^n$  as follows:

$$h_i^n := \frac{h_i}{\sqrt{n}}, \quad i = 0, 1, \dots, I, \quad (2.36)$$

$$c_i^n := \sqrt{n}c_i, \quad i = 1, \dots, I. \quad (2.37)$$

Lastly, we allow the mean travel time to vary with  $n$  as follows:

$$\eta^n := \eta + \frac{\hat{\eta}}{\sqrt{n}}, \quad (2.38)$$

where  $\hat{\eta} \in \mathbb{R}$ . As observed in Kogan and Lipster [1993] and Ata et al. [2024], under our heavy traffic assumption we expect that the queue lengths at the buffers to be of order  $\sqrt{n}$  and that the number of jobs in the infinite-server node be of order  $n$ . Therefore, we define the centered and scaled queue length processes as follows:

$$Z_0^n(t) := \frac{1}{\sqrt{n}}(Q_0^n(t) - n) \quad \text{and} \quad Z_i^n(t) := \frac{1}{\sqrt{n}}Q_i^n(t), \quad i = 1, \dots, I, \quad t \geq 0. \quad (2.39)$$

Observe that since  $\sum_{i=0}^I Q_i^n(t) = n$  for all  $t \geq 0$ , it follows that  $\sum_{i=0}^I Z_i^n(t) = 0$  for all  $t \geq 0$ .

As argued in Harrison [1988] (see also Harrison [2000, 2003]), any policy worthy of consideration satisfies  $T^n(t) \approx x^*t$ , for all  $t \geq 0$  and large  $n$ . That is, the nominal allocation rate  $x^*$  given in Assumption 5 should give a first-order approximation to the allocation rates of the various activities under policy  $T^n$ . However, the system manager can choose the

---

5. The first equality in (2.35) is proved by applying (2.34) and noting that  $(\Lambda^n)^{-1}(nx) = \Lambda^{-1}(x)$  for  $x \in \mathcal{L}$ . The second equality in (2.35) then follows by (2.7).

second-order, i.e., order  $1/\sqrt{n}$ , deviations from that. In order to capture such deviations from the nominal rates, we define the centered and scaled processes as follows:

$$Y_j^n(t) := \sqrt{n} \left( x_j^* t - T_j^n(t) \right), \quad j = 1, \dots, J, \quad t \geq 0, \quad (2.40)$$

Similarly, in the heavy traffic regime, we expect the servers to be always busy to a first-order approximation, but they may incur idleness on the second order, i.e., order  $1/\sqrt{n}$ . As such, we define the scaled idleness processes as follows:

$$U_i^n(t) := \sqrt{n} I_i^n(t), \quad i = 1, \dots, I, \quad t \geq 0. \quad (2.41)$$

Then, it follows from (2.15) and (2.29) that

$$U_i^n(t) = \sum_{j \in \mathcal{A}_i} Y_j^n(t), \quad i = 1, \dots, I, \quad t \geq 0. \quad (2.42)$$

In addition, we define the centered and scaled demand and service rate processes, respectively, as follows:

$$\zeta_i^n(t) := \frac{1}{\sqrt{n}} \left( \lambda_i^n(t) - n\lambda_i^* \right), \quad i = 1, \dots, I, \quad t \geq 0, \quad (2.43)$$

$$\kappa_j^n(t) := \frac{1}{\sqrt{n}} \left( \mu_j^n(t) - n\mu_j^* \right), \quad j = 1, \dots, J, \quad t \geq 0. \quad (2.44)$$

Note that by (2.9) we have  $\kappa_j^n(\cdot) = \zeta_i^n(\cdot)$  for each  $j \in \mathcal{A}_i$ . Finally, we define the centered cumulative profit function. To do so, we first introduce the auxiliary function  $\tilde{V}^n$  that will serve as the centering function. To be specific, we define

$$\tilde{V}^n(t) := n \left[ \pi(\lambda^*) - h_0^n \right] t = n\pi(\lambda^*)t - \sqrt{n}h_0t, \quad t \geq 0, \quad (2.45)$$

where the second equality follows from the definition of  $h_0^n$  from (2.36). Note that  $\tilde{V}^n(t)$  does

not depend on the system manager's control. Therefore, instead of maximizing the average profit, she can focus on minimizing the average cost, where the cumulative cost up to time  $t$ , denoted by  $\widehat{V}^n(t)$ , is defined as follows:

$$\widehat{V}^n(t) := \widetilde{V}^n(t) - V^n(t), \quad t \geq 0. \quad (2.46)$$

We then proceed with replacing the processes  $Z^n, Y^n, U^n, \zeta^n, \kappa^n$ , and  $\widehat{V}^n$  with their formal limits  $Z, Y, U, \zeta, \kappa$ , and  $\xi$ , respectively, as  $n \rightarrow \infty$ . In particular, the cost process  $\xi$  in the Brownian approximation is given by

$$\xi(t) = \int_0^t \left( \sum_{i=1}^I \alpha_i \zeta_i^2(s) + \sum_{i=0}^I h_i Z_i(s) \right) ds + c'U(t), \quad t \geq 0, \quad (2.47)$$

where  $\alpha_i := -(\Lambda_i^{-1})'(\lambda_i^*) - (\lambda_i^*/2) \times (\Lambda_i^{-1})''(\lambda_i^*) > 0$  for  $i = 1, \dots, I$ . The steps outlining the formal derivation of the Brownian control problem and (2.47) are given in Appendix 2.11.1.

The Brownian control problem (BCP) is given as follows: Choose processes  $Y = (Y_j)$  and  $\zeta = (\zeta_i)$  that are nonanticipating with respect to  $B$  so as to

$$\text{minimize } \limsup_{t \rightarrow \infty} \frac{1}{t} E[\xi(t)] \quad (2.48)$$

subject to

$$Z_i(t) = B_i(t) - q_i \eta \int_0^t \sum_{i=1}^I Z_i(s) ds - \sum_{j \in \mathcal{C}_i} \int_0^t x_j^* \kappa_j(s) ds \quad (2.49)$$

$$+ \sum_{j \in \mathcal{C}_i} \mu_j^* Y_j(t), \quad i = 1, \dots, I, \quad t \geq 0, \quad (2.50)$$

$$Z_0(t) = - \sum_{i=1}^I Z_i(t), \quad t \geq 0, \quad (2.51)$$

$$U(t) = AY(t), \quad t \geq 0, \quad (2.52)$$

$$\kappa_j(t) = \zeta_i(t) \text{ for all } j \in \mathcal{A}_i, i = 1, \dots, I, \text{ and } t \geq 0, \quad (2.53)$$



$$Z_i(t) \geq 0 \text{ for all } i = 1, \dots, I \text{ and } t \geq 0, \quad (2.54)$$

$$U \text{ is nondecreasing with } U(0) = 0, \quad (2.55)$$

where  $B = \{B(t), t \geq 0\}$  is an  $I$ -dimensional Brownian motion with starting state  $B(0) \geq 0$  that has drift rate vector  $\gamma = (\gamma_i)$  where  $\gamma_i := \hat{\eta}q_i$  and covariance matrix  $\Sigma$  given by

$$\Sigma_{ii} := q_i\eta + \sum_{j \in \mathcal{C}_i} \mu_j^* x_j^* \quad \text{and} \quad \Sigma_{ii'} := q_i q_{i'} \eta \quad \text{for } i, i' = 1, \dots, I, \quad i \neq i'. \quad (2.56)$$

Although the BCP (2.50)–(2.53) is simpler than the original control problem that it approximates, it is not easy to solve because it is a multidimensional stochastic control problem. Thus, we further simplify it in Section 2.5 and derive an equivalent workload formulation that is one-dimensional under the complete resource pooling condition which we solve analytically in Section 2.6.

## 2.5 Equivalent Workload Formulation

As a preliminary to the derivation of the workload problem, letting  $Z = (Z_1, \dots, Z_I)'$  and using (2.25), we first rewrite (2.50) in vector form as follows:

$$Z(t) = B(t) - \eta q \int_0^t e' Z(s) ds - C \text{diag}(x^*) \int_0^t \kappa(s) ds + RY(t), \quad t \geq 0, \quad (2.57)$$

where  $e$  is an  $I$ -dimensional vector of ones and  $\text{diag}(x^*)$  is the  $J \times J$  diagonal matrix whose  $(j, j)$ th element is  $x_j^*$ .

Motivated by the development in Harrison and Muehlebach [1997] and Harrison [2000], we define the space of reversible displacements as follows:

$$\mathcal{N} := \left\{ Hy_B : By_B = 0, y_B \in \mathbb{R}^b \right\}, \quad (2.58)$$

where  $y_B \in \mathbb{R}^b$  is the vector consisting of the components of  $y$  indexed by the basic activities  $j = 1, \dots, b$ . We let  $\mathcal{M} = \mathcal{N}^\perp$  be the orthogonal complement of the space  $\mathcal{N}$  and call  $d = \dim(\mathcal{M})$  the workload dimension. Any  $d \times I$  matrix  $M$  whose rows form a basis for  $\mathcal{M}$  is called a workload matrix. Lemma 6 provides a canonical choice of the workload matrix  $M$  based on the notion of communicating buffers, which is defined next, see Ata et al. [2020a]. Also see Harrison and López [1999] for a related definition of communicating servers.

**Definition 2.** Buffers  $i$  and  $i'$  are said to *communicate directly* if there exist basic activities  $j$  and  $j'$  such that  $i = b(j)$ ,  $i' = b(j')$ , and  $s(j) = s(j')$ . That is, buffers  $i$  and  $i'$  are served by a common server using basic activities. Buffers  $i$  and  $i'$  are said to *communicate* 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  $s = 1, \dots, l-1$ .

Buffer communication is an equivalence relation. Thus, the set of buffers can be partitioned into  $L$  disjoint subsets where all buffers in the same subset communicate with each other. We call each subset a buffer pool and denote the  $l$ th buffer pool by  $\mathcal{P}_l$  for  $l = 1, \dots, L$ . Associated with each buffer pool is a server pool. The  $l$ th server pool  $\mathcal{S}_l$  is defined as follows:

$$\mathcal{S}_l := \left\{ k : \exists j \in \{1, \dots, b\} \text{ s.t. } s(j) = k \text{ and } b(j) \in \mathcal{P}_l \right\}, \quad l = 1, \dots, L. \quad (2.59)$$

In words, server pool  $l$  consists of all servers that can serve a buffer in buffer pool  $l$  using a basic activity. Note that since the buffer pools partition the buffers, it follows from (2.59) that the server pools partition the servers. Thus, the buffer pools and the server pools are in a one-to-one correspondence. As a result, there is an equivalent notion of server communication, but we stick with the definition of buffer communication for mathematical convenience. The following lemma characterizes the workload dimension and the workload matrix.

**Lemma 6.** *The workload dimension equals the number of buffer pools, i.e.,  $d = L$ . Further-*

more, the  $L \times I$  matrix  $M$  given by

$$M_{li} := \begin{cases} 1, & \text{if } i \in \mathcal{P}_l, \\ 0, & \text{otherwise,} \end{cases} \quad (2.60)$$

for  $l = 1, \dots, L$  and  $i = 1, \dots, I$  constitutes a canonical workload matrix.

To facilitate the derivation of the workload state dynamics, we define the  $L \times I$  matrix  $G$  as follows:

$$G_{lk} := \lambda_k^* \mathbf{1}_{\{k \in \mathcal{S}_l\}}, \quad l = 1, \dots, L, \quad k = 1, \dots, I. \quad (2.61)$$

That is, the  $l$ th row of  $G$ ,  $(G_{l1}, \dots, G_{lI})$  contains the nominal service rates for those servers in server pool  $l$  and zeros for the rest of the servers. The next lemma provides a useful result that helps us derive the workload problem. It is proved in Appendix 2.11.2.

**Lemma 7.** *We have that  $MR = GA$ .*

We define the  $L$ -dimensional workload process  $W = \{W(t), t \geq 0\}$  as

$$W(t) := MZ(t), \quad t \geq 0, \quad (2.62)$$

whose  $l$ th component represents the total number of jobs for the  $l$ th server pool at time  $t$  for  $l = 1, \dots, L$ . By (2.62) and Lemma 7, we arrive at the following equation which describes the evolution of the workload process:

$$W(t) = \chi(t) - M\eta q \int_0^t e' Z(s) ds - MC \text{diag}(x^*) \int_0^t \kappa(s) ds + GU(t), \quad t \geq 0, \quad (2.63)$$

where  $\chi(t) := MB(t)$ , so that  $\chi = \{\chi(t), t \geq 0\}$  is a  $L$ -dimensional Brownian motion with drift vector  $M\gamma$ , covariance matrix  $M\Sigma M'$ , and starting state  $\chi(0) = MB(0) \geq 0$ .

Next, we introduce a closely related control problem referred to as the reduced Brownian control problem (RBCP). Its state descriptor is the workload process  $W$ . To be more specific, the RBCP is the following: Choose a policy  $(Z, U, \zeta)$  that is nonanticipating with respect to  $\chi$  so as to

$$\text{minimize } \limsup_{t \rightarrow \infty} \frac{1}{t} E \left[ \int_0^t \left( \sum_{i=1}^I \alpha_i \zeta_i^2(s) + \sum_{i=1}^I (h_i - h_0) Z_i(s) \right) ds + c'U(t) \right] \quad (2.64)$$

subject to

$$W(t) = MZ(t), \quad t \geq 0, \quad (2.65)$$

$$W(t) = \chi(t) - M\eta q \int_0^t e'Z(s) ds - MC \text{diag}(x^*) \int_0^t \kappa(s) ds + GU(t), \quad t \geq 0, \quad (2.66)$$

$$Z(t) \geq 0 \text{ for } t \geq 0, \quad (2.67)$$

$$U \text{ is nondecreasing with } U(0) = 0, \quad (2.68)$$

$$\kappa(t) = A'\zeta(t) \text{ for } t \geq 0. \quad (2.69)$$

The BCP (2.48)–(2.53) and the RBCP (2.64)–(2.69) are equivalent for the purposes of optimal control, as shown by the next proposition.

**Proposition 3.** *Every admissible policy  $(Y, \zeta)$  for the BCP (2.48)–(2.53) yields an admissible policy  $(Z, U, \zeta)$  for the RBCP (2.64)–(2.69) and these two policies have the same cost. On the other hand, for every admissible policy  $(Z, U, \zeta)$  of the RBCP, there exists an admissible policy  $(Y, \zeta)$  for the BCP whose cost is equal to that of the policy  $(Z, U, \zeta)$  for the RBCP.*

Hereafter, we make the complete resource pooling assumption that corresponds to having a single resource pool in our context, see Assumption 6 below. Harrison and López [1999] observes that the complete resource pooling assumption leads to a one-dimensional workload formulation, also see Ata and Kumar [2005]. Similarly, Assumption 6 allows us to formulate a one-dimensional workload formulation that is equivalent to the RBCP (2.64)–(2.69).

**Assumption 6.** All buffers communicate under the nominal processing plan, i.e.,  $L = 1$ .

This assumption says that servers have sufficiently overlapping capabilities under the nominal processing plan; see Harrison and López [1999] for further details. The following lemma allows us to simplify the RBCP under Assumption 6.

**Lemma 8.** *Under Assumption 6, we have  $M = e'$  and  $G = (\lambda^*)'$ . Moreover, we have that*

$$M\eta q = \eta \quad \text{and} \quad MC \text{diag}(x^*)A' = e'. \quad (2.70)$$

Using Lemma 8, the RBCP can be equivalently written as follows: Choose a policy  $(Z, U, \zeta)$  that is nonanticipating with respect to  $\chi$  so as to

$$\text{minimize} \quad \limsup_{t \rightarrow \infty} \frac{1}{t} E \left[ \int_0^t \left( \sum_{i=1}^I \alpha_i \zeta_i^2(s) + \sum_{i=1}^I (h_i - h_0) Z_i(s) \right) ds + c'U(t) \right] \quad (2.71)$$

subject to

$$W(t) = \sum_{i=1}^I Z_i(t), \quad t \geq 0, \quad (2.72)$$

$$W(t) = \chi(t) - \eta \int_0^t W(s) ds - \int_0^t \sum_{i=1}^I \zeta_i(s) ds + \sum_{i=1}^I \lambda_i^* U_i(t), \quad t \geq 0, \quad (2.73)$$

$$Z(t) \geq 0 \text{ for } t \geq 0, \quad (2.74)$$

$$U \text{ is nondecreasing with } U(0) = 0, \quad (2.75)$$

where  $\chi$  is a one-dimensional Brownian motion with drift rate parameter  $a := e'\gamma$  and variance parameter  $\sigma^2 := e'\Sigma e$  and starting state  $\chi(0) = \sum_{i=1}^I B_i(0) \geq 0$ .

To further simplify the RBCP, we define the cost function  $c: \mathbb{R} \rightarrow \mathbb{R}$  by

$$c(x) := \min \left\{ \sum_{i=1}^I \alpha_i \zeta_i^2 : e'\zeta = x, \zeta \in \mathbb{R}^I \right\}, \quad x \in \mathbb{R}, \quad (2.76)$$

and the optimal (state-dependent) drift rate function  $\zeta^* : \mathbb{R} \rightarrow \mathbb{R}$  by

$$\zeta^*(x) := \operatorname{argmin} \left\{ \sum_{i=1}^I \alpha_i \zeta_i^2 : e'\zeta = x, \zeta \in \mathbb{R}^I \right\}, \quad x \in \mathbb{R}. \quad (2.77)$$

Defining  $\hat{\alpha} := \sum_{i=1}^I 1/\alpha_i$ , the following lemma characterizes these functions—similar results are found in Çelik and Maglaras [2008] and Ata and Barjesteh [2023].

**Lemma 9.** *We have that  $c(x) = \frac{1}{\hat{\alpha}}x^2$  and  $\zeta_i^*(x) = \frac{1}{\alpha_i \hat{\alpha}}x$  for  $i = 1, \dots, I$  and  $x \in \mathbb{R}$ .*

In the workload formulation, it is optimal to keep all workload in the buffer with the lowest holding cost, i.e., buffer  $i^*$  given by

$$i^* := \operatorname{argmin}_{i=1, \dots, I} h_i, \quad (2.78)$$

with corresponding effective holding cost  $h := h_{i^*} - h_0 > 0$ . Moreover, the system manager will only idle the server that is cheapest to idle, i.e., server  $k^*$  given by

$$k^* := \operatorname{argmin}_{i=1, \dots, I} \frac{c_i}{\lambda_i^*}, \quad (2.79)$$

with corresponding effective idling cost  $r := c_{k^*}/\lambda_{k^*}^*$ .

The workload formulation can now be stated as follows: Choose a policy  $\theta : [0, \infty) \rightarrow \mathbb{R}$  that is nonanticipating with respect to  $\chi$  so as to

$$\text{minimize } \limsup_{t \rightarrow \infty} \frac{1}{t} E \left[ \int_0^t [c(\theta(s)) + h W(s)] ds + r L(t) \right] \quad (2.80)$$

subject to

$$W(t) = \chi(t) - \eta \int_0^t W(s) ds - \int_0^t \theta(s) ds + L(t), \quad t \geq 0, \quad (2.81)$$

$$W(t) \geq 0 \text{ for } t \geq 0, \quad (2.82)$$

$$L \text{ is nondecreasing with } L(0) = 0, \quad (2.83)$$

The RBCP (2.64)–(2.69) and the EWF (2.80)–(2.83) are equivalent for the purposes of optimal control, as shown by the next proposition.

**Proposition 4.** *Every admissible policy  $\theta$  for the EWF (2.80)–(2.83) yields an admissible policy  $(Z, U, \zeta)$  for the RBCP (2.71)–(2.75) and these two policies have the same cost. On the other hand, for every admissible policy  $(Z, U, \zeta)$  of the RBCP, there exists an admissible policy  $\theta$  for the EWF whose cost is less than or equal to that of the policy  $(Z, U, \zeta)$  for the RBCP.*

In what follows, we add two additional constraints to the equivalent workload formulation. First, we require that

$$\int_0^\infty \mathbf{1}_{\{W(t) > 0\}} dL(t) = 0, \quad (2.84)$$

which requires that the process  $L$  can increase only when  $W = 0$ . That is, the control policy must be work conserving. We include this restriction because its optimality is intuitive from the cost structure, i.e., there are both holding and idleness costs, and that the workload process is one dimensional. Second, we impose the following regularity condition:

$$\lim_{t \rightarrow \infty} \frac{E[W(t)]}{t} = 0.$$

To repeat, we further require a policy  $\theta$  to satisfy these conditions to be admissible.

## 2.6 Solution to the Equivalent Workload Formulation

This section solves the EWF (2.80)–(2.83). In order to minimize technical complexity, we restrict attention to stationary Markov policies. That is, the drift chosen at time  $t$  will be a function of the current workload only, and so we write it as  $\theta(W(t))$ . To facilitate the analysis, we next consider the Bellman equation for the workload formulation which is the

following second-order nonlinear differential equation: Find a function  $f \in \mathcal{C}^2[0, \infty)$  and a constant  $\beta \in \mathbb{R}$  satisfying

$$\begin{aligned} \beta &= \min_{x \in \mathbb{R}} \left\{ \frac{1}{2} \sigma^2 f''(w) - \eta w f'(w) - x f'(w) + a f'(w) + c(x) + h w \right\} \\ &= \min_{x \in \mathbb{R}} \left\{ \frac{1}{\hat{\alpha}} x^2 - x f'(w) \right\} + \frac{1}{2} \sigma^2 f''(w) - \eta w f'(w) + a f'(w) + h w, \quad w \geq 0, \end{aligned} \quad (2.85)$$

subject to the boundary conditions

$$f'(0) = -r \quad \text{and} \quad f' \text{ is increasing with } \lim_{w \rightarrow \infty} f'(w) = \frac{h}{\eta}. \quad (2.86)$$

The optimization problem on the right-hand side of (2.85) is convex. Therefore, its solution is easily seen to be

$$x^*(w) := \frac{\hat{\alpha}}{2} f'(w), \quad w \geq 0. \quad (2.87)$$

The Bellman equation can then be simplified as follows: Find a function  $f \in \mathcal{C}^2[0, \infty)$  and a constant  $\beta \in \mathbb{R}$  satisfying

$$\beta = -\frac{\hat{\alpha}}{4} [f'(y)]^2 + \frac{1}{2} \sigma^2 f''(y) - \eta y f'(y) + a f'(y) + h y, \quad y \geq 0, \quad (2.88)$$

subject to the boundary conditions

$$f'(0) = -r \quad \text{and} \quad f' \text{ is increasing with } \lim_{w \rightarrow \infty} f'(w) = \frac{h}{\eta}. \quad (2.89)$$

Setting  $v = f'$ , the Bellman equation can be written as follows: find a function  $v \in \mathcal{C}^1[0, \infty)$  and a constant  $\beta \in \mathbb{R}$  satisfying

$$\beta = -\frac{\hat{\alpha}}{4} v^2(y) + \frac{1}{2} \sigma^2 v'(y) - \eta y v(y) + a v(y) + h y, \quad y \geq 0, \quad (2.90)$$



subject to the boundary conditions

$$v(0) = -r \quad \text{and} \quad v \text{ is increasing with } \lim_{y \rightarrow \infty} v(y) = \frac{h}{\eta}. \quad (2.91)$$

This expresses the Bellman equation as a first-order differential equation. The following theorem guarantees the existence of a solution to the Bellman equation.

**Theorem 3.** *The Bellman equation (2.90)–(2.91) has a solution  $(\beta^*, v)$  with  $\beta^* > 0$ .*

With  $\beta^* > 0$  and  $v$  given by Theorem 3, we define

$$f(y) := \int_0^y v(x) dx, \quad y \geq 0.$$

The next result is immediate from Theorem 3 and provides a solution to the original Bellman equation:

**Corollary 3.** *The pair  $(\beta^*, f)$  solves the Bellman equation (2.85)–(2.86).*

Define the following candidate policy  $\theta^* : [0, \infty) \rightarrow \mathbb{R}$  by

$$\theta^*(w) := \frac{\hat{\alpha}}{2} v(w), \quad w \geq 0. \quad (2.92)$$

The following proposition facilitates the proof of our main result.

**Proposition 5.** *The candidate policy  $\theta^*$  is admissible for the equivalent workload formulation.*

*That is, letting  $W^* = \{W^*(t), t \geq 0\}$  denote the workload process under the candidate policy  $\theta^*$ , we have*

$$\lim_{t \rightarrow \infty} \frac{E[W^*(t)]}{t} = 0.$$

The following result establishes that the candidate policy is optimal:

**Theorem 4.** *The candidate policy  $\theta^*$  is optimal for the equivalent workload formulation (2.80)–(2.83), and its long-run average cost is  $\beta^*$ .*

Next, we state an auxiliary lemma used in the proof of Theorem 4.

**Lemma 10.** *Let  $W$  be the workload process defined by (2.81)–(2.83) under an arbitrary admissible policy. Then the following hold:*

$$(i) \quad E \int_0^t f'(W(s)) d(\chi(s) - as) = 0, \quad t \geq 0,$$

$$(ii) \quad \limsup_{t \rightarrow \infty} \frac{E[f(W(t))]}{t} = 0.$$

*Proof.* Since the stochastic process  $t \mapsto \chi(t) - at$  is standard Brownian motion, by Proposition 4.7 in Harrison [2013], to prove part (i) it suffices to show that

$$E \int_0^t [f'(W(s))]^2 ds < \infty \quad \text{for each } t \geq 0.$$

Since  $f'(w) \in [-r, h/\eta]$  for all  $w \geq 0$  by (2.86) and  $W(t) \geq 0$  for all  $t \geq 0$  by (2.82), we have

$$E \int_0^t [f'(W(s))]^2 ds \leq t \left( r + \frac{h}{\eta} \right)^2 < \infty, \quad t \geq 0,$$

proving part (i). In order to prove part (ii), that it suffices to show that

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

To this end, observe that

$$\begin{aligned} |E[f(W(t))]| &\leq E|f(W(t))| = E \left| \int_0^{W(t)} f'(s) ds \right| \leq E \int_0^{W(t)} |f'(s)| ds \\ &\leq \left( r + \frac{h}{\eta} \right) E[W(t)]. \end{aligned}$$

Thus, by definition of an admissible policy, it follows that

$$\limsup_{t \rightarrow \infty} \frac{|E[f(W(t))]|}{t} \leq \left( r + \frac{h}{\eta} \right) \limsup_{t \rightarrow \infty} \frac{E[W(t)]}{t} = 0,$$

proving part (ii). □

We conclude this section with a proof of Theorem 4.

*Proof of Theorem 4.* By (2.81), note that for an admissible policy  $\theta$ ,

$$dW(s) = d\chi(s) - \eta W(s) ds - \theta(W(s)) ds + dL(s). \quad (2.93)$$

Furthermore, since  $L(s)$  is nondecreasing in  $s$ , the processes is a VF function almost surely; see Section B.2 in Harrison (2013). Therefore,

$$\begin{aligned} [dW(s)]^2 &= [d\chi(s)]^2 + 2d\chi(s)[- \eta W(s) ds - \theta(W(s)) ds + dL(s)] \\ &\quad + [- \eta W(s) ds - \theta(W(s)) ds + dL(s)]^2 \\ &= \sigma^2 ds. \end{aligned} \quad (2.94)$$

Note that the last two terms on the right-hand side of (2.94) are zero; see Chapter 4 in Harrison (2013). Then, for  $f \in C^2[0, \infty)$ , Itô's Lemma gives

$$df(W(s)) = f'(W(s))dW(s) + \frac{1}{2}f''(W(s))[dW(s)]^2. \quad (2.95)$$

Define the differential operator  $\Gamma_\theta : C^2[0, \infty) \rightarrow C[0, \infty)$  by

$$(\Gamma_\theta f)(w) = \frac{1}{2}\sigma^2 f''(w) - [\eta w + \theta(w) - a]f'(w), \quad w \geq 0. \quad (2.96)$$

Then, combining (2.93)–(2.96) gives

$$df(W(s)) = f'(W(s))d(\chi(s) - as) + \Gamma_\theta f(W(s)) ds + f'(W(s)) dL(s). \quad (2.97)$$

Integrating both sides of (2.97) over  $[0, t]$  gives

$$f(W(t)) = f(W(0)) + \int_0^t f'(W(s)) d(\chi(s) - as) + \int_0^t \Gamma_\theta f(W(s)) ds + \int_0^t f'(W(s)) dL(s). \quad (2.98)$$

Recall that by (2.84) the process  $L$  increases only when  $W = 0$ . Thus, for  $f \in C^2[0, \infty)$  satisfying  $f'(0) = -r$  we have

$$\int_0^t f'(W(s)) dL(s) = f'(0)L(t) = -rL(t). \quad (2.99)$$

Moreover, for the solution  $(\beta^*, f)$  of the Bellman equation (2.85)–(2.86) it follows that

$$\beta^* - c(\theta(w)) - hw \leq \frac{1}{2}\sigma^2 f''(w) - [\eta w + \theta(w) - a]f'(w), \quad w \geq 0, \quad (2.100)$$

with equality holding when  $\theta = \theta^*$ . Therefore by (2.96), (2.100), Lemma 10, and taking expectations, it follows that

$$\begin{aligned} E[f(W(t))] - [f(W(0))] + E[rL(t)] &= E \int_0^t \Gamma_\theta f(W(s)) ds \\ &\geq E \int_0^t [\beta^* - c(\theta(W(s))) - hW(s)] ds, \end{aligned} \quad (2.101)$$

with equality holding when  $\theta = \theta^*$ . Rearranging terms in (2.101) and dividing by  $t$  gives

$$\frac{1}{t} E \left[ \int_0^t [c(\theta(s)) + hW(s)] ds + rL(t) \right] \geq \beta^* - \frac{1}{t} Ef(W(t)) + \frac{1}{t} Ef(W(0)), \quad (2.102)$$

with equality holding when  $\theta = \theta^*$ . Finally, taking limits on both sides of (2.102) and applying Lemma 10 gives

$$\limsup_{t \rightarrow \infty} \frac{1}{t} E \left[ \int_0^t [c(\theta(s)) + hW(s)] ds + rL(t) \right] \geq \beta^*,$$

with equality holding when  $\theta = \theta^*$ . Therefore, the policy  $\theta^*$  is optimal for the equivalent workload formulation and its long-run average cost is  $\beta^*$ .  $\square$

## 2.7 Solution to the Bellman Equation

In this section we prove Theorem 3 by considering an initial value problem that is closely related to the Bellman equation. Namely, for each fixed  $\beta \geq 0$  consider the following initial value problem, denoted by IVP( $\beta$ ): Find a function  $v \in C^1[0, \infty)$  such that

$$\frac{\sigma^2}{2} v'(y) = \beta + \frac{\hat{\alpha}}{4} v^2(y) + \eta y \left( v(y) - \frac{h}{\eta} \right) - av(y), \quad y \geq 0, \quad (2.103)$$

$$v(0) = -r. \quad (2.104)$$

The following result is standard and is proven in Appendix 2.11.2.

**Lemma 11.** *For  $\beta \geq 0$ , there exists a unique solution  $v_\beta \in C^1[0, \infty)$  to (2.103)–(2.104).*

For the remainder of this section, we analyze the (unique) solution to (2.103)–(2.104), focusing on how the behavior of the solution varies with the parameter  $\beta$ . Using this approach, we ultimately find a  $\beta^* > 0$ , with corresponding solution  $v_{\beta^*}$ , such that the pair  $(\beta^*, v_{\beta^*})$  solves the original Bellman equation. Namely, we look for  $\beta^*$  such that  $v_{\beta^*}$  satisfies the second condition in (2.89) that  $v_{\beta^*}$  is increasing with  $\lim_{y \rightarrow \infty} v_{\beta^*}(y) = h/\eta$ .

For much of our analysis, we consider parameters that satisfy one of two cases, given in Assumption 7 below. To state the assumption, let

$$\underline{\beta}_1 = 0 \quad \text{and} \quad \underline{\beta}_2 = -ar - \frac{\hat{\alpha}r^2}{4}.$$

**Assumption 7.** One of the following holds:

- (a) Case 1:  $a > -\frac{\hat{\alpha}}{4}r$  and  $\beta \geq \underline{\beta}_1$ .

(b) Case 2:  $a \leq -\frac{\hat{\alpha}}{4}r$  and  $\beta > \underline{\beta}_2$ .

**Remark.** Note that under Assumption 7(b), we have that  $\underline{\beta}_2 \geq 0$ .

Lemmas 12–14 below facilitate the analysis to follow.

**Lemma 12.** If  $y > 0$  is a local maximizer of  $v_\beta(y)$ , then  $v_\beta(y) \leq h/y$ .

*Proof.* Because  $y$  is a local maximizer, we have that  $v'_\beta(y) = 0$ , and  $v''_\beta(y) \leq 0$ . Differentiating both sides of (2.103) and using  $v'_\beta(y) = 0$ , we write

$$\frac{\sigma^2}{2} v''_\beta(y) = \eta \left( v_\beta(y) - \frac{h}{\eta} \right) \leq 0,$$

from which it follows that  $v_\beta(y) \leq h/y$ . □

**Lemma 13.** Under Assumption 7,  $v_\beta$  increases to its supremum.

*Proof.* First, note that  $v'_\beta(0) = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{2\sigma^2}r^2 + \frac{2ar}{\sigma^2} > 0$  in either case of Assumption 7. Aiming for a contradiction, suppose  $v_\beta$  does not increase to its maximum. Then we must have  $0 \leq x_1 < x_2 < x_3$  such that

$$\begin{aligned} v_\beta(x_1) &= v_\beta(x_2) = v_\beta(x_3) = v, \\ v'_\beta(x_1) &> 0, \quad v'_\beta(x_2) < 0, \quad v'_\beta(x_3) > 0. \end{aligned}$$

In particular, we have the following equations:

$$v'_\beta(x_1) = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{2\sigma^2}v^2 + \frac{2\eta}{\sigma^2}x_1 \left( v - \frac{h}{\eta} \right) - \frac{2av}{\sigma^2} > 0, \quad (2.105)$$

$$v'_\beta(x_2) = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{2\sigma^2}v^2 + \frac{2\eta}{\sigma^2}x_2 \left( v - \frac{h}{\eta} \right) - \frac{2av}{\sigma^2} < 0, \quad (2.106)$$

$$v'_\beta(x_3) = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{2\sigma^2}v^2 + \frac{2\eta}{\sigma^2}x_3 \left( v - \frac{h}{\eta} \right) - \frac{2av}{\sigma^2} > 0. \quad (2.107)$$

On the one hand, subtracting (2.106) from (2.105) yields

$$\frac{2\eta}{\sigma^2}(x_1 - x_2)\left(v - \frac{h}{\eta}\right) > 0. \quad (2.108)$$

Because  $x_1 - x_2 < 0$ , we conclude from (2.108) that

$$v - \frac{h}{\eta} < 0 \quad (2.109)$$

On the other hand, subtracting (2.106) from (2.107) gives

$$\frac{2\eta}{\sigma^2}(x_3 - x_2)\left(v - \frac{h}{\eta}\right) > 0. \quad (2.110)$$

But, we deduce from (2.109) and from  $x_3 - x_2 > 0$  that the left-hand side of (2.110) is negative, which is a contradiction. This completes the proof.  $\square$

**Lemma 14.** *Let  $0 \leq x_1 < x_2$ . Under Assumption 7, the following condition is necessary for  $v_\beta(x)$  to be constant on  $(x_1, x_2)$ :*

$$\beta = a \frac{h}{\eta} - \frac{\hat{\alpha}}{4} \left(\frac{h}{\eta}\right)^2. \quad (2.111)$$

*Moreover, if  $v_\beta$  is constant on  $(x_1, x_2)$ , then  $v_\beta(x) = h/\eta$  for  $x \in (x_1, x_2)$ , and letting  $\hat{x} := \inf\{x \geq 0 : v_\beta(x) = h/\eta\}$ , it follows that  $v_\beta$  is nondecreasing on  $[0, \hat{x}]$  and stays constant at value  $h/\eta$  thereafter.*

*On the other hand, if (2.111) does not hold, then there is no interval on which  $v_\beta$  is constant, i.e., the set  $\{y \geq 0 : v'_\beta(y) = 0\}$  has Lebesgue measure zero.*

*Proof.* Suppose the condition in (2.111) is violated, which implies

$$\beta + \frac{\hat{\alpha}}{4} \left(\frac{h}{\eta}\right)^2 - a \frac{h}{\eta} \neq 0 \quad (2.112)$$

Aiming for a contraction, suppose there exist an interval  $(x_1, x_2)$  such that  $v_\beta(y) = v$  on it. This implies  $v'_\beta(y) = v''_\beta(y) = 0$  on  $(x_1, x_2)$ . Differentiating both sides of (2.103) and using  $v'_\beta(y) = 0$  on  $(x_1, x_2)$  gives

$$\frac{\sigma^2}{2} v''_\beta(y) = \eta \left( v_\beta(y) - \frac{h}{\eta} \right), \quad y \in (x_1, x_2).$$

Thus,  $v_\beta(y) = h/\eta$  on  $(x_1, x_2)$ . Substituting this into (2.103) yields

$$\frac{\sigma^2}{2} v'_\beta(y) = \beta + \frac{\hat{\alpha}}{4} \left( \frac{h}{\eta} \right)^2 - a \frac{h}{\eta} \neq 0, \quad y \in (x_1, x_2),$$

which follows from (2.112) and contradicts that  $v'_\beta(y) = 0$  on  $(x_1, x_2)$ . Therefore, if (2.111) does not hold, then there is no interval on which  $v_\beta$  is constant.

Now, we turn to the first part of the lemma. If  $v_\beta$  is constant on  $(x_1, x_2)$ , then  $v'_\beta(x) = v''_\beta(x) = 0$  on  $(x_1, x_2)$ . As argued above, these imply  $v_\beta(x) = h/\eta$  on  $(x_1, x_2)$ . In addition, it follows from (2.103) and  $v'_\beta(x) = -h/\eta$  on  $(x_1, x_2)$  that

$$\beta + \frac{\hat{\alpha}}{4} \left( \frac{h}{\eta} \right)^2 - a \frac{h}{\eta} = 0,$$

proving the necessary condition (2.112). Building on these, because at any local maximum  $v_\beta(x) \leq h/\eta$  by Lemma 12 and  $\hat{x}$  is the first time  $v_\beta$  reaches to its maximum by Lemma 13, we conclude that  $v_\beta$  is nondecreasing on  $[0, \hat{x}]$ . To conclude the proof, consider an auxiliary IVP involving (2.103) on  $[\hat{x}, \infty)$  with the initial condition  $v(\hat{x}) = h/\eta$ . Then setting  $v(x) = h/\eta$  solves it. Moreover, combining that with  $v_\beta$  on  $[0, \hat{x})$  constitutes a solution to the IVP (2.103)–(2.104). By Lemma 11, this is the unique solution.  $\square$

To facilitate the analysis below, we define the following four sets. First, consider Case 1



of Assumption 7, and let

$$\mathcal{I}_1 := \{\beta \geq 0 : v_\beta \text{ is nondecreasing on } (0, \infty)\},$$

$$\mathcal{D}_1 := \{\beta \geq 0 : \exists x_\beta \geq 0 \text{ s.t. } v_\beta \text{ is nondecreasing on } (0, x_\beta) \text{ and decreasing on } (x_\beta, \infty)\}.$$

Similarly, for Case 2 of Assumption 7, we define

$$\mathcal{I}_2 := \{\beta > \underline{\beta}_2 : v_\beta \text{ is nondecreasing on } (0, \infty)\},$$

$$\mathcal{D}_2 := \{\beta > \underline{\beta}_2 : \exists x_\beta \geq 0 \text{ s.t. } v_\beta \text{ is nondecreasing on } (0, x_\beta) \text{ and decreasing on } (x_\beta, \infty)\}.$$

**Lemma 15.** *We have the following:*

(i) *Under Assumption 7(a),  $\beta \in \mathcal{D}_1$  if and only if  $\exists x_0 \in (0, \infty)$  such that  $v'_\beta(x_0) < 0$ .*

(ii) *Under Assumption 7(b),  $\beta \in \mathcal{D}_2$  if and only if  $\exists x_0 \in (0, \infty)$  such that  $v'_\beta(x_0) < 0$ .*

*Proof.* First, note from Lemma 14 that it is necessary that  $v_\beta$  increases to  $h/\eta$  and stay constant thereafter for it to be constant on any interval. In that case, we would have  $\beta \in \mathcal{I}_i$  ( $i = 1$  under Assumption 7(a) and  $i = 2$  under Assumption 7(b)). Thus, for the remainder of the proof, we assume there is no interval on which  $v_\beta$  is constant.

We prove Cases (i) and (ii) simultaneously because their proofs are identical. For  $i = 1, 2$ , let  $\beta \in \mathcal{D}_i$ . Aiming for a contradiction, assume there does not exist  $x_0 > 0$  such that  $v'_\beta(x_0) < 0$ . Then,  $v'_\beta(x) \geq 0$  for all  $x \geq 0$ , i.e.,  $v_\beta$  is nondecreasing on  $(0, \infty)$ . Thus,  $\beta \in \mathcal{I}_i$ , a contradiction. Therefore, there exists  $x_0 > 0$  such that  $v'_\beta(x_0) < 0$ .

For the other direction, suppose there exists  $x_0 > 0$  such that  $v'_\beta(x_0) < 0$ . Because  $v_\beta$  increases to its maximum (Lemma 13), it is not constant on any interval (by the argument given in the opening paragraph of this proof) and  $v'_\beta(x_0) < 0$ , it achieves its maximum at

some  $x^* < x_0$ . Thus, by Lemma 12, we have that

$$v_\beta(x) \leq v_\beta(x^*) \leq \frac{h}{\eta}, \quad x \geq 0. \quad (2.113)$$

Aiming for a contradiction, suppose that  $\beta \notin \mathcal{D}_i$ . Then  $v_\beta$  cannot be decreasing over  $[x^*, \infty)$ .

Thus, there exist  $x_1$  and  $x_2$  such that

$$\begin{aligned} x^* &< x_1 < x_2, \\ v &= v_\beta(x_1) = v_\beta(x_2) \leq \frac{h}{\eta}, \\ v'_\beta(x_1) &< 0 < v'_\beta(x_2). \end{aligned}$$

In particular, the following holds:

$$v'_\beta(x_1) = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{2\sigma^2} v^2 + \frac{2\eta}{\sigma^2} x_1 \left( v - \frac{h}{\eta} \right) - av < 0, \quad (2.114)$$

$$v'_\beta(x_2) = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{2\sigma^2} v^2 + \frac{2\eta}{\sigma^2} x_2 \left( v - \frac{h}{\eta} \right) - av > 0. \quad (2.115)$$

Subtracting (2.114) from (2.115) gives

$$0 < v'_\beta(x_2) - v'_\beta(x_1) = \eta(x_2 - x_1) \left( v - \frac{h}{\eta} \right) \leq 0,$$

where the last inequality follows because  $x_2 - x_1 > 0$  and  $v \leq \frac{h}{\eta}$  by (2.113), leading to a contradiction. We conclude that  $\beta \in \mathcal{D}_i$ .  $\square$

**Corollary 4.** *Under Assumption 7, we have the following:*

(i) *In Case 1 of Assumption 7, the sets  $\mathcal{I}_1$  and  $\mathcal{D}_1$  partition  $[0, \infty)$ .*

(ii) *In Case 2 of Assumption 7, the sets  $\mathcal{I}_2$  and  $\mathcal{D}_2$  partition  $(\underline{\beta}_2, \infty)$ .*

*Proof.* Consider Case 1 of Assumption 7. Then for  $\beta \geq 0$ , if  $v'_\beta(x) < 0$  for some  $x > 0$ , it

follows that  $\beta \in \mathcal{D}_1$  by Lemma 15. Otherwise,  $v'_\beta(x) \geq 0$  for all  $x > 0$ , in which case  $\beta \in \mathcal{I}_1$  by definition. This proves (i). Proof of (ii) follows similarly.  $\square$

**Corollary 5.** *For  $i = 1, 2$ , under Case  $i$  of Assumption 7, we have that  $v_\beta$  achieves its maximum whenever  $\beta \in \mathcal{D}_i$ , and*

$$\sup_{x \geq 0} v_\beta(x) < \frac{h}{\eta}.$$

*Proof.* For  $i = 1, 2$ , by definition of  $\mathcal{D}_i$ ,  $\exists x^* \geq 0$  such that  $v_\beta$  is nondecreasing on  $(0, x^*)$  and it is decreasing on  $(x^*, \infty)$ . First, note that if  $x^* = 0$ , then  $v_\beta(x)$  is decreasing everywhere and  $v_\beta(0) = -r$  and the result follows. Thus, we assume  $x^* > 0$ . Note that  $v_\beta$  achieves its maximum at  $x^*$ . Also, we conclude from Lemma 12 that  $v_\beta(x^*) \leq h/\eta$ . Aiming for a contradiction, suppose  $v_\beta(x^*) = h/\eta$ . Note that  $v'_\beta(x^*) = 0$  because  $x^*$  is the maximizer. From these, by differentiating both sides of (2.103), we conclude that  $v''_\beta(x^*) = 0$ . Then we can argue as in the proof of Lemma 14 that  $v_\beta(x) = h/\eta$  for  $x \geq x^*$ , implying  $\beta \notin \mathcal{D}_i$ , a contradiction. Thus,  $v_\beta(x^*) \neq h/\eta$ , completing the proof.  $\square$

**Lemma 16.** *For  $i = 1, 2$ , under Case  $i$  of Assumption 7, we have that  $\lim_{x \rightarrow \infty} v_\beta(x) = -\infty$  whenever  $\beta \in \mathcal{D}_i$ .*

*Proof.* It follows from Corollary 5 that  $v_\beta$  has a maximizer  $x^*$  such that

$$v_\beta(x) \leq v_\beta(x^*) < \frac{h}{\eta}, \quad x \geq 0. \quad (2.116)$$

Also define the constant

$$\epsilon := \frac{h}{\eta} - v_\beta(x^*) > 0. \quad (2.117)$$

To prove  $\lim_{x \rightarrow \infty} v_\beta(x) = -\infty$ , we argue by contradiction. To that end, suppose there exists a

$K_1 > 0$  such that  $v_\beta(x) \geq -K_1$  for  $x \geq 0$ . Then we have that

$$|v_\beta(x)| \leq K_2 := \max\left\{K_1, \frac{h}{\eta}\right\}. \quad (2.118)$$

Recalling IVP( $\beta$ ), we bound  $v'_\beta(\cdot)$  using (2.116)–(2.118) as follows:

$$\begin{aligned} \frac{\sigma^2}{2} v'_\beta(y) &\leq \beta + \frac{\hat{\alpha}}{4} K_2^2 + \eta y \left( v_\beta(x^*) - \frac{h}{\eta} \right) + |a|K_2 \\ &= \left[ \beta + \frac{\hat{\alpha}}{4} K_2^2 + |a|K_2 \right] - \epsilon \eta y, \quad y \geq 0. \end{aligned} \quad (2.119)$$

Integrating both sides of (2.119) over  $[0, y]$  and using the initial condition  $v_\beta(0) = -r$  gives

$$\frac{\sigma^2}{2} v'_\beta(y) \leq -\frac{\sigma^2}{2} r + \left[ \beta + \frac{\hat{\alpha}}{4} K_2^2 + |a|K_2 \right] y - \frac{\eta \epsilon}{2} y^2, \quad y \geq 0. \quad (2.120)$$

Since  $\eta \epsilon / 2 > 0$ , the right-hand side of (2.120) tends to  $-\infty$  as  $y \rightarrow \infty$ , implying that  $v(y) \rightarrow -\infty$  as  $y \rightarrow \infty$ , a contradiction.  $\square$

**Lemma 17.** *For  $i = 1, 2$ , under Case  $i$  of Assumption 7, the following are equivalent:*

- (i)  $\beta \in \mathcal{D}_i$ ,
- (ii)  $\exists x > 0$  such that  $v'_\beta(x) < 0$ ,
- (iii)  $\exists x > 0$  such that  $v_\beta(x) < -r$ ,
- (iv)  $\lim_{x \rightarrow \infty} v_\beta(x) = -\infty$ .

*Proof.* Parts (i) and (ii) are equivalent by Lemma 15. Part (i) implies (iv) by Lemma 16. Clearly, (iv) implies (iii). Therefore, it suffices to prove that (iii) implies (ii). To that end, let  $x_0 > 0$  be such that  $v_\beta(x_0) < -r$ . Since  $v_\beta(0) = -r$ , it follows from the mean value theorem that there exists a  $\hat{x}_0 \in (0, x_0)$  such that

$$v'_\beta(\hat{x}_0) = \frac{v_\beta(x_0) - v_\beta(0)}{x_0 - 0} < 0,$$

proving part (ii). □

**Lemma 18.** *Under Assumption 7, we have that  $\lim_{x \rightarrow \infty} v_\beta(x) = \infty$  if and only if there exists an  $x_0 > 0$  such that  $v_\beta(x_0) \geq \frac{h}{\eta}$ .*

*Proof.* First, if  $\lim_{x \rightarrow \infty} v_\beta(x) = \infty$ , then clearly, there exists an  $x_0 > 0$  such that  $v_\beta(x_0) > \frac{h}{\eta}$ . To prove the other direction, suppose there exists  $x_0 > 0$  such that  $v_\beta(x_0) > \frac{h}{\eta}$ , and define

$$x_1 := \inf \left\{ x > 0 : v_\beta(x) \geq \frac{h}{\eta} \right\},$$

Because  $v_\beta(0) = -r < \frac{h}{\eta} < v_\beta(x_0)$ , by the intermediate value theorem,  $v_\beta(x_1) = \frac{h}{\eta}$ . Next, we argue that  $v_\beta(x) > \frac{h}{\eta}$  for all  $x > x_1$ . If not, then there exists an  $x_2 > x_1$  such that  $v_\beta(x_2) \leq \frac{h}{\eta}$ . Then let

$$x_3 := \inf \left\{ x > x_1 : v_\beta(x) \leq \frac{h}{\eta} \right\}.$$

Note that  $v_\beta(x_3) = \frac{h}{\eta}$  by continuity of  $v_\beta$ . Moreover, note that  $x_3 > x_1$  since  $v_\beta(x_1) = \frac{h}{\eta}$  and

$$v'_\beta(x_1) = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{2\sigma^2} \left( \frac{h}{\eta} \right)^2 - a \frac{h}{\eta} > 0, \tag{2.121}$$

where the inequality holds because (i)  $v'_\beta(x_1) \geq 0$  (by definition of  $x_1$ ), (ii)  $x_1 < x_0$ ,  $v_\beta(x_0) > \frac{h}{\eta}$ , and  $v_\beta$  increases to its maximum, and (iii)  $v'_\beta(x_1) \neq 0$  by Lemma 14 (because  $v_\beta(x_0) > \frac{h}{\eta}$ ).

Thus, (2.121) follows. Consequently, we have that

$$v_\beta(x) > \frac{h}{\eta} \quad \text{for } x \in (x_1, x_3). \tag{2.122}$$

By continuity,  $v_\beta(x)$  achieves a local maximum at some  $\hat{x} \in (x_1, x_3)$  and  $v_\beta(\hat{x}) > \frac{h}{\eta}$ , but this

contradicts Lemma 12. Therefore, we conclude that

$$v_\beta(x) > \frac{h}{\eta} \quad \text{for } x \geq x_1, \quad (2.123)$$

implying that  $\beta \notin \mathcal{D}_i$ . Therefore,  $\beta \in \mathcal{I}_i$  and  $v_\beta$  is nondecreasing by Corollary 4, and so

$$v_\beta(x) \geq v_\beta(x_0) > \frac{h}{\eta} \quad \text{for } x > x_0. \quad (2.124)$$

To conclude the proof, we consider two cases:  $a \leq 0$  and  $a > 0$ . When  $a \leq 0$ , we note from (2.103) that

$$\frac{\sigma^2}{2} v'_\beta(y) \geq \beta + \frac{\hat{\alpha}}{4} \left(\frac{h}{\eta}\right)^2, \quad y \geq x_0. \quad (2.125)$$

Integrating both sides of (2.125) over  $[x_0, y]$  gives

$$v_\beta(y) \geq \frac{h}{\eta} + \frac{2}{\sigma^2} \left[ \beta + \frac{\hat{\alpha}}{4} \left(\frac{h}{\eta}\right)^2 \right] (y - x_0), \quad y \geq x_0, \quad (2.126)$$

where the right-hand side tends to  $\infty$ , completing the proof when  $a \leq 0$ .

When  $a > 0$ , we note from (2.103) that

$$v'_\beta(y) + \frac{2a}{\sigma^2} v_\beta(y) = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{2\sigma^2} v_\beta^2(y) + \eta y \left( v_\beta(y) - \frac{h}{\eta} \right), \quad y \geq x_0. \quad (2.127)$$

We let  $\epsilon = v_\beta(x_0) - h/\eta > 0$  and write from (2.127) that

$$v'_\beta(y) + \frac{2a}{\sigma^2} v_\beta(y) = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{2\sigma^2} \left(\frac{h}{\eta}\right)^2 + \epsilon \eta y. \quad (2.128)$$

Multiplying both sides of (2.128) with the integrating factor  $\exp\left\{\frac{2a}{\sigma^2}y\right\}$  yields:

$$\left(v_\beta(y) \exp\left\{\frac{2a}{\sigma^2}y\right\}\right)' \geq C \exp\left\{\frac{2a}{\sigma^2}y\right\} + \epsilon\eta y \exp\left\{\frac{2a}{\sigma^2}y\right\},$$

where  $C = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{\sigma^2}(h/\eta)^2 > 0$ . Integrating both sides of this on  $[x_0, y]$  yields

$$v_\beta(y) \geq v_\beta(x_0) + C \left(1 - \exp\left\{-\frac{2a}{\sigma^2}(y - x_0)\right\}\right) + \epsilon\eta \frac{\sigma^4}{4a^2} \left(\frac{2a}{\sigma^2}y - 1\right),$$

where the right-hand side tends to  $\infty$  as  $y \rightarrow \infty$ , completing the proof when  $a > 0$ .  $\square$

**Lemma 19.** *For  $0 \leq \beta_1 < \beta_2$ , we have that  $v_{\beta_1}(x) < v_{\beta_2}(x)$  for all  $x > 0$ . That is,  $v_\beta(x)$  is an increasing function of  $\beta$  for each  $x > 0$ .*

*Proof.* Let  $\beta_2 > \beta_1 \geq 0$ . We argue by contradiction. Suppose  $v_{\beta_1}(x) \geq v_{\beta_2}(x)$  for some  $x > 0$ , and let

$$\hat{x} := \inf \{x > 0 : v_{\beta_1}(x) \geq v_{\beta_2}(x)\}.$$

Then there exists a sequence  $\{x_n\}$  that decreases to  $\hat{x}$ , i.e.,  $x_n \searrow \hat{x}$  as  $n \rightarrow \infty$ , such that  $v_{\beta_1}(x_n) \geq v_{\beta_2}(x_n)$  for all  $n$ . Recall that  $v_{\beta_1}(0) = v_{\beta_2}(0) = -r$  and  $v'_{\beta_2}(0) > v'_{\beta_1}(0)$ . Hence,  $v_{\beta_2} > v_{\beta_1}$  in a neighborhood around zero. This and continuity of  $v_{\beta_1}$  and  $v_{\beta_2}$  imply that

$$v_{\beta_1}(\hat{x}) = v_{\beta_2}(\hat{x}). \tag{2.129}$$

Consequently, we can write

$$\frac{v_{\beta_1}(x_n) - v_{\beta_1}(\hat{x})}{x_n - \hat{x}} \geq \frac{v_{\beta_2}(x_n) - v_{\beta_2}(\hat{x})}{x_n - \hat{x}}, \quad n \geq 1.$$

Passing to the limit as  $n \rightarrow \infty$ , we conclude that

$$v'_{\beta_1}(\hat{x}) \geq v'_{\beta_2}(\hat{x}). \quad (2.130)$$

However, note that from IVP( $\beta$ ) for  $\beta = \beta_1, \beta_2$  we have

$$\frac{\sigma^2}{2} v'_{\beta_1}(\hat{x}) = \beta_1 + \frac{\hat{\alpha}}{4} v_{\beta_1}^2(\hat{x}) + \eta \hat{x} \left( v_{\beta_1}(\hat{x}) - \frac{h}{\eta} \right) - av_{\beta_1}(\hat{x}), \quad (2.131)$$

$$\frac{\sigma^2}{2} v'_{\beta_2}(\hat{x}) = \beta_2 + \frac{\hat{\alpha}}{4} v_{\beta_2}^2(\hat{x}) + \eta \hat{x} \left( v_{\beta_2}(\hat{x}) - \frac{h}{\eta} \right) - av_{\beta_1}(\hat{x}). \quad (2.132)$$

Subtracting (2.131) from (2.132) and using (2.129) yield

$$\frac{\sigma^2}{2} [v'_{\beta_2}(\hat{x}) - v'_{\beta_1}(\hat{x})] = \beta_2 - \beta_1 > 0,$$

which contradicts (2.130). Thus, we conclude that  $v_{\beta_2}(x) > v_{\beta_1}(x)$  for  $x > 0$ .  $\square$

**Lemma 20.** *For  $x > 0$ , we have that  $v_{\beta}(x)$  is continuous in  $\beta$  on  $[0, \infty)$ . That is, for  $x > 0$ , given  $\beta \geq 0$  and  $\epsilon > 0$ , there exists a  $\delta > 0$  such that  $|v_{\beta}(x) - v_{\tilde{\beta}}(x)| < \epsilon$  for all  $\tilde{\beta} \in (\beta - \delta, \beta + \delta) \cap [0, \infty)$ .*

*Proof.* Let  $x > 0$  and  $\beta_2 > \beta_1 \geq 0$ . Integrating IVP( $\beta$ ) over  $[0, x]$  for  $\beta = \beta_1, \beta_2$ , we arrive at the following two equations:

$$\frac{\sigma^2}{2} v_{\beta_1}(x) = -\frac{\sigma^2}{2}r + \beta_1 x + \frac{\hat{\alpha}}{4} \int_0^x v_{\beta_1}^2(y) dy + \eta \int_0^x y \left( v_{\beta_1}(y) - \frac{h}{\eta} \right) dy - \int_0^x av_{\beta_1}(y) dy, \quad (2.133)$$

$$\frac{\sigma^2}{2} v_{\beta_2}(x) = -\frac{\sigma^2}{2}r + \beta_2 x + \frac{\hat{\alpha}}{4} \int_0^x v_{\beta_2}^2(y) dy + \eta \int_0^x y \left( v_{\beta_2}(y) - \frac{h}{\eta} \right) dy - \int_0^x av_{\beta_2}(y) dy. \quad (2.134)$$



Subtracting (2.133) from (2.134) gives the following:

$$\begin{aligned} \frac{\sigma^2}{2} [v_{\beta_2}(x) - v_{\beta_1}(x)] &= (\beta_2 - \beta_1)x + \frac{\hat{\alpha}}{4} \int_0^x [v_{\beta_2}^2(y) - v_{\beta_1}^2(y)] dy \\ &\quad + \eta \int_0^x y [v_{\beta_2}(x) - v_{\beta_1}(x)] dy - a \int_0^x [v_{\beta_1}(y) - v_{\beta_2}(y)] dy. \end{aligned} \quad (2.135)$$

In order to facilitate the bound, let  $\bar{\beta} > \beta_2 > \beta_1 \geq 0$  and note from Lemma 19 that

$$v_0(y) \leq v_{\beta_1}(y) \leq v_{\beta_2}(y) \leq v_{\bar{\beta}}(y), \quad y \geq 0.$$

Hence for  $y \geq 0$  we have that

$$2v_0(y) \leq v_{\beta_1}(y) + v_{\beta_2}(y) \leq 2v_{\bar{\beta}}(y),$$

from which we conclude that

$$|v_{\beta_1}(y) + v_{\beta_2}(y)| \leq 2 \max(|v_0(y)| + |v_{\bar{\beta}}(y)|).$$

Thus, letting

$$K(\bar{\beta}) := 2 \sup_{0 \leq y \leq x} \left\{ \max(|v_0(y)| + |v_{\bar{\beta}}(y)|) \right\},$$

we arrive at the following for  $y \in [0, x]$ :

$$|v_{\beta_2}^2(y) - v_{\beta_1}^2(y)| = |v_{\beta_2}(y) + v_{\beta_1}(y)| \cdot |v_{\beta_2}(y) - v_{\beta_1}(y)| \leq K(\bar{\beta}) |v_{\beta_2}(y) - v_{\beta_1}(y)|.$$

Combining this with (2.135) and letting

$$h(y) := |v_{\beta_2}(y) - v_{\beta_1}(y)|, \quad y \in [0, x],$$

yields the following inequality:

$$h(x) \leq \frac{2x}{\sigma^2} |\beta_2 - \beta_1| + \left[ \frac{\hat{\alpha}}{2\sigma^2} K(\bar{\beta}) + \eta x + |a| \right] \int_0^x h(y) dy.$$

By Gronwall's inequality (see, e.g., page 498 of Ethier and Kurtz [2005]) we conclude that

$$h(x) \leq \frac{2x}{\sigma^2} |\beta_2 - \beta_1| \exp \left\{ - \left( \eta x + \frac{\hat{\alpha}}{2\sigma^2} K(\bar{\beta}) + |a| \right) x \right\}.$$

Thus, given  $\epsilon > 0$ , we can let

$$\delta := \frac{\epsilon \sigma^2}{2x} \exp \left\{ - \left( \eta x + \frac{\hat{\alpha}}{2\sigma^2} K(\bar{\beta}) + |a| \right) x \right\},$$

so that  $|\beta_2 - \beta_1| < \delta$  implies that  $h(x) = |v_{\beta_2}(x) - v_{\beta_1}(x)| < \epsilon$ . This completes the proof.  $\square$

**Lemma 21.** *Under Assumption 7, we have the following:*

- (i) *In Case 1 of Assumption 7, for  $0 \leq \beta_1 < \beta_2$ , if  $\beta_2 \in \mathcal{D}_1$ , then  $\beta_1 \in \mathcal{D}_1$ . That is,  $[0, \beta_2] \subseteq \mathcal{D}_1$  whenever  $\beta_2 \in \mathcal{D}_1$ .*
- (ii) *In Case 2 of Assumption 7, for  $\underline{\beta}_2 < \beta_1 < \beta_2$ , if  $\beta_2 \in \mathcal{D}_2$ , then  $\beta_1 \in \mathcal{D}_2$ . That is,  $(\underline{\beta}_2, \beta_2] \subseteq \mathcal{D}_2$  whenever  $\beta_2 \in \mathcal{D}_2$ .*

*Proof.* Consider part (i), and let  $\beta_2 > \beta_1 \geq 0$ . Then by Lemma 17, there exists  $x_0 > 0$  such that  $v_{\beta_2}(x_0) < -r$ . In turn, by Lemma 19, we have that

$$v_{\beta_1}(x_0) < v_{\beta_2}(x_0) < -r,$$

Thus,  $\beta_1 \in \mathcal{D}_1$  by Lemma 17. Proof of part (ii) follows similarly.  $\square$

**Lemma 22.** *Under Assumption 7, we have the following:*

- (i) *In Case 1 of Assumption 7,  $\mathcal{D}_1 \neq \emptyset$ . In particular,  $0 \in \mathcal{D}_1$  and there exists a  $\tilde{\beta}_1 > 0$  such that  $[0, \tilde{\beta}] \subseteq \mathcal{D}_1$ .*

(ii) In Case 2 of Assumption 7,  $\mathcal{D}_2 \neq \emptyset$ . In particular, there exists  $\tilde{\beta}_2 > \underline{\beta}_2$  such that  $(\underline{\beta}_2, \tilde{\beta}_2] \subseteq \mathcal{D}_2$ .

*Proof.* Consider part (i). We first show  $0 \in \mathcal{D}_1$ . Aiming for a contradiction, suppose  $0 \notin \mathcal{D}_1$  so that  $0 \in \mathcal{I}_1$  by Corollary 4. We consider the following two cases:

- Case A:  $v_0(y) \leq 0$  for all  $y > 0$ .
- Case B:  $v_0(y) > 0$  for some  $y > 0$ .

We first consider Case A. Because  $0 \in \mathcal{I}_1$ ,  $v'_0(y) \geq 0$  for all  $y \geq 0$ . Then, we have that  $-r \leq v_0(y) \leq 0$  for all  $y \geq 0$ . Substituting this into IVP( $\beta$ ) for  $\beta = 0$ , we consider the following two subcases of Case A:  $a \geq 0$  and  $a \in (-\frac{\alpha r}{4}, 0)$ .

For  $a \geq 0$ , we conclude that

$$0 \leq \frac{\sigma^2}{2} v'_0(y) \leq \frac{\hat{\alpha}}{4} r^2 - hy + ar,$$

where the right-hand side tends to  $-\infty$ . Thus, there exists  $y > 0$  such that  $v'_0(y) < 0$ , contradicting  $0 \in \mathcal{I}_1$ .

For  $a \in (-\frac{\alpha r}{4}, 0)$ , we conclude that

$$0 \leq \frac{\sigma^2}{2} v'_0(y) \leq \frac{\hat{\alpha}}{4} r^2 - hy,$$

where the right-hand side tends to  $-\infty$ . Once again, there exists  $y > 0$  such that  $v'_0(y) < 0$ , contradicting  $0 \in \mathcal{I}_1$ .

We now consider Case B. In this case, we let  $y_0 = \inf \{y > 0 : v_0(y) > 0\}$ . By continuity of  $v_0$  and  $v_0(0) = -r < 0$ , we have that  $v_0(y_0) = 0$  and  $y_0 > 0$ . Substituting this into IVP( $\beta$ ) for  $\beta = 0$  at  $y = y_0$  gives

$$\frac{\sigma^2}{2} v'_0(y_0) = -hy_0 < 0.$$

Thus,  $0 \in \mathcal{D}_1$  by Lemma 17, a contradiction. Combining Cases A and B, we conclude that  $0 \in \mathcal{D}_1$ . Then it follows from Lemma 17 that  $v_0(y) \rightarrow -\infty$  as  $y \rightarrow \infty$ . Thus, there exists a  $x_0 > 0$  such that  $v_0(x_0) < -2r$ . Then, by continuity of  $v_\beta(x_0)$  in  $\beta$  (see Lemma 20), there exists a  $\tilde{\beta}_1 > 0$  such that  $v_{\tilde{\beta}_1}(x_0) < -r$ . By Lemma 17, we conclude  $\tilde{\beta}_1 \in \mathcal{D}_1$ . Then we conclude by Lemma 21 that  $[0, \tilde{\beta}_1] \subseteq \mathcal{D}_1$ .

Consider part (ii). Recall that in Case 2 of Assumption 7,  $a \leq -\hat{\alpha}r/4$  and  $\underline{\beta}_2 = -ar - \hat{\alpha}r^2/4 \geq 0$ . Consider  $v_{\underline{\beta}_2}$  and note that  $v_{\underline{\beta}_2}(0) = -r$ . It follows from (2.103) that  $v'_{\underline{\beta}_2}(0) = 0$ . Moreover, differentiating both sides of (2.103) and using  $v'_{\underline{\beta}_2}(0) = 0$ , we conclude that

$$v''_{\underline{\beta}_2}(0) = -\frac{2\eta}{\sigma^2} \left( r + \frac{h}{\eta} \right) < 0.$$

Thus,  $v_{\underline{\beta}_2}$  is decreasing and below  $-r$  in a neighborhood of zero. Next, we argue that  $v_{\underline{\beta}_2}(x) \leq -r$  for all  $x > 0$ .

Suppose not, and let  $x_1 = \inf\{x > 0 : v_{\underline{\beta}_2}(x) = -r\}$ . By continuity of  $v_\beta$ , we have  $v_{\underline{\beta}_2}(x_1) = -r$ . We also have by its definition that  $v'_{\underline{\beta}_2}(x_1) \geq 0$  and  $x_1 > 0$ . Then by combining these with (2.103), we write

$$\begin{aligned} 0 \leq v'_{\underline{\beta}_2}(x_1) &= -ar - \frac{\hat{\alpha}}{4}r^2 + \frac{\hat{\alpha}}{4}r^2 - \eta x_1 \left( r + \frac{h}{\eta} \right) + ar \\ &= -\eta x_1 \left( r + \frac{h}{\eta} \right) < 0, \end{aligned}$$

a contradiction. Thus,  $v_{\underline{\beta}_2}(x) \leq -r$  for all  $x \geq 0$ .

Next, we argue that  $\lim_{x \rightarrow \infty} v_{\underline{\beta}_2}(x) = -\infty$ . Suppose not (Note that we can rule out oscillatory behavior following the same technique in the proof of Lemma 13). Then, there exists  $k > r$  such that

$$v_{\underline{\beta}_2}(x) \geq -k, \quad x \geq 0.$$

But using (2.103), we conclude that

$$\frac{\sigma^2}{2} v'_{\underline{\beta}_2}(x) \leq \beta_2 + \frac{\hat{\alpha}}{4} K^2 - \eta x \left( r + \frac{h}{\eta} \right) - ar.$$

Integrating both sides from 0 to  $y$  yields

$$\frac{\sigma^2}{2} v_{\underline{\beta}_2}(y) \leq -r \frac{\sigma^2}{2} + \left[ \underline{\beta}_2 - ar + \frac{\hat{\alpha}}{4} \right] y - \frac{\eta}{2} \left( r + \frac{h}{\eta} \right) \frac{y^2}{2},$$

where the right-hand side tends to  $-\infty$  as  $y \rightarrow \infty$ . Thus,  $v_{\underline{\beta}_2}(x) \rightarrow -\infty$  as  $x \rightarrow \infty$  and there exists  $x_2$  such that  $v_{\underline{\beta}_2}(x_2) < -2r$ . Then, by Lemma 19, there exists  $\tilde{\beta}_2 > \underline{\beta}_2$  such that  $v_{\tilde{\beta}_2}(x_2) < -r$ . In particular,  $\tilde{\beta}_2 \in \mathcal{D}_2$  by Lemma 17. Then, by Lemma 21, we conclude that  $(\underline{\beta}_2, \tilde{\beta}_2] \subset \mathcal{D}_2$ .  $\square$

**Lemma 23.** *Under Assumption 7, we have  $\mathcal{I}_i \neq \emptyset$  for  $i = 1, 2$ . In particular,*

$$\left( \frac{\sigma^2 h}{2\eta} + 2\sigma \left( r + \frac{h}{\eta} \right) \sqrt{\frac{\eta}{\pi}} \exp \left\{ -\frac{\sigma^2 a^2}{4\eta} \right\}, \infty \right) \subseteq \mathcal{I}_i, \quad i = 1, 2.$$

*Proof.* We establish the result by showing that  $v_\beta(x) \rightarrow \infty$  as  $x \rightarrow \infty$  for sufficiently large  $\beta > 0$ . The result then follows from Corollary 4 and Lemmas 17 and 19. To that end, we rewrite IVP( $\beta$ ) as follows:

$$v'_\beta(y) - \frac{2\eta}{\sigma^2} y v_\beta(y) + a v_\beta(y) = \frac{2\beta}{\sigma^2} + \frac{\hat{\alpha}}{2\sigma^2} v_\beta^2(y) - \frac{2h}{\sigma^2} y, \quad y \geq 0.$$

Multiplying both sides with the integrating factor  $\exp \left\{ -\frac{\eta}{\sigma^2} y^2 + ay \right\}$  yields the following bound:

$$\left[ \exp \left\{ -\frac{\eta}{\sigma^2} y^2 + ay \right\} v_\beta(y) \right]' \geq \frac{2\beta}{\sigma^2} \exp \left\{ -\frac{\eta}{\sigma^2} y^2 + ay \right\} - \frac{2h}{\sigma^2} y \exp \left\{ -\frac{\eta}{\sigma^2} y^2 + ay \right\}.$$

Integrating both sides of the above inequality over  $[0, x]$  and using  $v_\beta(0) = -r$  gives:

$$\exp\left\{-\frac{\eta}{\sigma^2}x^2 + ax\right\}v_\beta(x) \geq -r + \frac{2\beta}{\sigma^2}I_1 - \frac{2h}{\sigma^2}I_2, \quad (2.136)$$

where

$$I_1 := \int_0^x \exp\left\{-\frac{\eta}{\sigma^2}y^2 + ay\right\} dy \quad \text{and} \quad I_2 := \int_0^x y \exp\left\{-\frac{\eta}{\sigma^2}y^2 + ay\right\} dy.$$

First, we consider  $I_1$  and write

$$I_1 = \exp\left\{\frac{\sigma^2}{4\eta}a^2\right\} \int_0^x \exp\left\{-\frac{\eta}{\sigma^2}\left(y - \frac{a\sigma^2}{2\eta}\right)^2\right\} dy.$$

Applying the change of variable  $u = \frac{\sqrt{2\eta}}{\sigma}\left(y - \frac{\sigma^2}{2\eta}\right)$  yields

$$\begin{aligned} I_1 &= \sqrt{\frac{\pi}{\eta}}\sigma \exp\left\{\frac{\sigma^2}{4\eta}a^2\right\} \int_{-\frac{\sigma}{\sqrt{2\eta}}}^{\frac{\sqrt{2\eta}}{\sigma}\left(x - \frac{\sigma^2}{2\eta}\right)} \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{u^2}{2}\right\} du \\ &= \sqrt{\frac{\pi}{\eta}}\sigma \exp\left\{\frac{\sigma^2}{4\eta}a^2\right\} \left[ \Phi\left(\frac{\sqrt{2\eta}}{\sigma}\left(x - \frac{\sigma^2}{2\eta}\right)\right) - \Phi\left(-\frac{\sigma}{\sqrt{2\eta}}\right) \right], \end{aligned} \quad (2.137)$$

where  $\Phi$  is the CDF for the standard normal distribution. Next, we turn to  $I_2$  and facilitate its derivative by first deriving

$$I_3 = \int_0^x \left(y - \frac{\sigma^2 a}{2\eta}\right) \exp\left\{-\frac{\eta}{\sigma^2}y^2 + ay\right\} dy.$$

Note that  $I_3 = I_2 - \frac{\sigma^2 a}{2\eta}I_1$ . Using the change of variable  $u = -\frac{\eta}{\sigma^2}y^2 + ay$ , we write

$$I_3 = \int_0^{-\frac{\eta}{\sigma^2}x^2+ax} -\frac{\sigma^2}{2u}e^u du = \frac{\sigma^2}{2\eta} \left[ 1 - \exp\left\{-\frac{\eta}{\sigma^2}x^2 + ax\right\} \right].$$

Then, using  $I_2 = I_3 + \frac{a\sigma^2}{2\eta}I_1$ , we arrive at

$$I_2 = \frac{\sigma^2}{2\eta} - \frac{\sigma^2}{2\eta} \exp\left\{-\frac{\eta}{\sigma^2}x^2 + ax\right\} + \frac{\sigma^2 a^2}{2\eta} \exp\left\{\frac{a^2\sigma^2}{4\eta}\right\} \sqrt{\frac{\pi}{\eta}} \sigma \left[ \Phi\left(\frac{\sqrt{2\eta}}{\sigma}x - \frac{\sigma}{\sqrt{2\eta}}\right) - \Phi\left(-\frac{\sigma}{\sqrt{2\eta}}\right) \right]. \quad (2.138)$$

Substituting (2.137)–(2.138) into (2.136) then gives

$$\begin{aligned} \exp\left\{-\frac{\eta}{\sigma^2}x^2 + ax\right\} v_\beta(x) &\geq -r - \frac{h}{\eta} + \frac{h}{\eta} \exp\left\{-\frac{\eta}{\sigma^2}x^2 + ax\right\} \\ &\quad + \frac{2\beta}{\sigma^2} \sqrt{\frac{\pi}{\eta}} \exp\left\{\frac{\sigma^2 a^2}{4\eta}\right\} \left[ \Phi\left(\frac{\sqrt{2\eta}}{\sigma}x - \frac{\sigma}{\sqrt{2\eta}}\right) - \Phi\left(-\frac{\sigma}{\sqrt{2\eta}}\right) \right] \\ &\quad - \sigma \frac{h}{\eta} \exp\left\{\frac{a^2\sigma^2}{4\eta}\right\} \sqrt{\frac{\pi}{\eta}} \left[ \Phi\left(\frac{\sqrt{2\eta}}{\sigma}x - \frac{\sigma}{\sqrt{2\eta}}\right) - \Phi\left(-\frac{\sigma}{\sqrt{2\eta}}\right) \right]. \end{aligned} \quad (2.139)$$

Note that there exists  $x_0 > 0$  large enough so that

$$\Phi\left(\frac{\sqrt{2\eta}}{\sigma}x - \frac{\sigma}{\sqrt{2\eta}}\right) - \Phi\left(-\frac{\sigma}{\sqrt{2\eta}}\right) \geq \frac{1}{4}. \quad (2.140)$$

Then for  $x \geq x_0$  and  $\beta > \frac{\sigma^2 h}{2\eta}$ , combining (2.139) and (2.140), we write,

$$\exp\left\{-\frac{\eta}{\sigma^2}x^2 + ax\right\} v_\beta(x) \geq -r + \left(\frac{2\beta}{\sigma} - \frac{\sigma h}{\eta}\right) \exp\left\{\frac{\sigma^2 a^2}{4\eta}\right\} \sqrt{\frac{\pi}{\eta}} \frac{1}{\eta} - \frac{h}{\eta} + \frac{h}{\eta} \exp\left\{-\frac{\eta}{\sigma^2}x^2 + ax\right\}.$$

Thus, we have the following lower bound on  $v_\beta(\cdot)$ :

$$v_\beta(x) \geq \left[ \frac{1}{4} \left( \frac{2\beta}{\sigma} - \frac{\sigma h}{\eta} \right) \exp\left\{\frac{\sigma^2 a^2}{4\eta}\right\} \sqrt{\frac{\pi}{\eta}} - \left( r + \frac{h}{\eta} \right) \right] \exp\left\{\frac{\eta}{\sigma^2}x^2 + ax\right\} + \frac{h}{\eta}, \quad x > x_0. \quad (2.141)$$

In particular, we note that for  $\beta > \frac{\sigma^2 h}{2\eta} + 2\sigma \left( r + \frac{h}{\eta} \right) \sqrt{\frac{\eta}{\pi}} \exp\left\{-\frac{\sigma^2 a^2}{4\eta}\right\}$ , the right-hand side of (2.141) tends to  $\infty$  as  $x \rightarrow \infty$ . Thus,  $\beta \in \mathcal{I}_i$  for  $i = 1, 2$  whenever it is above  $\frac{\sigma^2 h}{2\eta} + 2\sigma \left( r + \frac{h}{\eta} \right) \sqrt{\frac{\eta}{\pi}} \exp\left\{-\frac{\sigma^2 a^2}{4\eta}\right\}$ , completing the proof.  $\square$

To facilitate the analysis, under Case  $i$  of Assumption 7, we define  $\beta_i^* = \inf \mathcal{I}_i$  for  $i = 1, 2$ . The remaining results will prove that this  $\beta_i^*$  along with its corresponding  $v_{\beta_i^*}$ , solve the Bellman equation in Case  $i$  for  $i = 1, 2$ .

**Lemma 24.** *For  $i = 1, 2$ , under Case  $i$  of Assumption 7, we have that  $\beta_i^* > 0$ .*

*Proof.* Recall from Lemma 22 that there exists a  $\tilde{\beta}_i > 0$  such that  $\tilde{\beta}_i \in \mathcal{D}_i$  for  $i = 1, 2$ . Clearly, we must have  $\beta \geq \tilde{\beta}_i$  for  $\beta \in \mathcal{I}_i$  and  $i = 1, 2$ . Thus, we conclude that  $\beta_i^* = \inf \mathcal{I}_i \geq \tilde{\beta}_i > 0$  for  $i = 1, 2$ .  $\square$

**Lemma 25.** *For  $i = 1, 2$ , under Case  $i$  of Assumption 7, we have that  $\beta_i^* \in \mathcal{I}_i$  and  $v_{\beta_i^*}$  is bounded.*

*Proof.* Consider Case  $i$  of Assumption 7, for  $i = 1, 2$ . We argue by contradiction. Suppose  $\beta_i^* \notin \mathcal{I}_i$ . Then, by Corollary 4,  $\beta_i^* \in \mathcal{D}_i$ . In particular, by Lemma 17, there exists a  $x_0 > 0$  such that  $v_{\beta_i^*}(x) < -r$ . Because  $v_{\beta}(x_0)$  is continuous in  $\beta$  (see Lemma 20), there exists a  $\delta > 0$  such that

$$v_{\beta}(x_0) < -r \quad \text{for } \beta \in (\beta_i^* - \delta, \beta_i^* + \delta). \quad (2.142)$$

However, by definition of  $\beta_i^*$ , there exists a  $\hat{\beta}_i \in (\beta_i^*, \beta_i^* + \delta)$  such that  $\hat{\beta}_i \in \mathcal{I}_i$ . Applying Lemma 17 again, it follows that  $v_{\hat{\beta}_i}(x) \geq -r$  for all  $x \geq 0$ , contradicting (2.142). Thus,  $\beta_i^* \in \mathcal{I}_i$ .

We now prove that  $v_{\beta_i^*}$  is bounded. Aiming for a contradiction, suppose it is not bounded. Then there exists a  $x_0 > 0$  such that  $v_{\beta_i^*}(x_0) > 2h/\eta$ . Then, because  $v_{\beta}(x_0)$  is continuous in  $\beta$  (by Lemma 20) and  $\beta_i^* > 0$  (by Lemma 24), there exists an  $\epsilon > 0$  such that  $v_{\beta_i^* - \epsilon}(x_0) \geq h/\eta$ . It follows that  $v_{\beta_i^* - \epsilon}$  is unbounded by Lemma 18, which in turn implies that  $\beta_i^* - \epsilon \in \mathcal{I}_i$  by Corollary 4 and Lemma 16. That  $\beta_i^* - \epsilon \in \mathcal{I}_i$ , however, contradicts the definition of  $\beta_i^*$ .  $\square$

**Lemma 26.** *Under Assumption 7, the following hold:*

$$(i) \quad \mathcal{D}_1 = [0, \beta_1^*) \text{ and } \mathcal{I}_1 = [\beta_1^*, \infty),$$



(ii)  $\mathcal{D}_2 = (\underline{\beta}_2, \beta_2^*)$  and  $\mathcal{I}_2 = [\beta_2^*, \infty)$ .

*Proof.* Consider Case  $i$  of Assumption 7, for  $i = 1, 2$ . Suppose that there exists a  $\beta > \beta_i^*$  such that  $\beta \in \mathcal{D}_i$ . Then by Lemma 21 it follows that  $\beta_i^* \in \mathcal{D}_i$ , contradicting Lemma 25. Hence, no such  $\beta$  exists. Combining this with Lemma 25 and the definition of  $\beta_i^*$  completes the proof.  $\square$

**Lemma 27.** *For  $i = 1, 2$ , under Case  $i$  of Assumption 7, we have that  $v_{\beta_i^*}$  is nondecreasing with  $\lim_{x \rightarrow \infty} v_{\beta_i^*}(x) = h/\eta$ .*

*Proof.* Consider Case  $i$  of Assumption 7, for  $i = 1, 2$ . Because  $\beta_i^* \in \mathcal{I}_i$  by Lemma 25,  $v_{\beta_i^*}$  is nondecreasing. Also, by Lemma 25 we have that  $v_{\beta_i^*}$  is bounded. Consequently, by Lemma 18, we have that

$$v_{\beta_i^*}(x) \leq \frac{h}{\eta} \quad \text{for } x \geq 0.$$

Moreover, because  $v_{\beta_i^*}$  is nondecreasing, its limit is well-defined and satisfies

$$\lim_{x \rightarrow \infty} v_{\beta_i^*}(x) \leq \frac{h}{\eta}.$$

Now let  $v = \lim_{x \rightarrow \infty} v_{\beta_i^*}(x)$  and suppose that  $v < \frac{h}{\eta}$ . Consider IVP( $\beta_i^*$ ), i.e.,

$$\frac{\sigma^2}{2} v'_{\beta_i^*}(y) = \beta_i^* + \frac{\hat{\alpha}}{4} v_{\beta_i^*}^2(y) + \eta y \left( v_{\beta_i^*}(y) - \frac{h}{\eta} \right) - av_{\beta_i^*}(y), \quad y \geq 0.$$

Passing to the limit on both sides and noting that  $v < \frac{h}{\eta}$  gives the following:

$$\frac{\sigma^2}{2} \lim_{y \rightarrow \infty} v'_{\beta_i^*}(y) = \beta_i^* + \frac{\hat{\alpha}}{4} v^2 - av + \lim_{y \rightarrow \infty} \eta y \left( v_{\beta_i^*}(y) - \frac{h}{\eta} \right) = -\infty.$$

Thus, there exists a  $x_0 > 0$  such that  $v'_{\beta_i^*}(x_0) < 0$ . We conclude by Lemma 17 that  $\beta_i^* \in \mathcal{D}_i$ , a contradiction. Therefore,  $v = \lim_{x \rightarrow \infty} v_{\beta_i^*}(x) = h/\eta$ .  $\square$

We conclude this section with a proof of Theorem 3.

*Proof of Theorem 3.* First, consider  $a > -\frac{\hat{\alpha}}{4}r$ , for which Case 1 of Assumption 7 applies. In this case,  $(\beta_1^*, v_{\beta_1^*})$  solves (2.103)–(2.104) and this solution is unique by Lemma 11. Moreover, by Lemma 27, we have that  $\lim_{x \rightarrow \infty} v_{\beta_1^*}(x) = h/\eta$ . Finally, by Lemma 24, we have that  $\beta_1^* > 0$ . Therefore,  $(\beta_1^*, v_{\beta_1^*})$  solves the Bellman equation (2.90)–(2.91). When  $a \leq -\frac{\hat{\alpha}}{4}r$ , Case 2 of Assumption 7 applies, and the proof follows from the same steps as in the first case.  $\square$

## 2.8 Proposed Policy

In this section we propose a dynamic pricing and dispatch policy for the problem introduced in Section 2.3 by interpreting the solution of the equivalent workload formulation (2.80)–(2.84) in the context of the original control problem. To describe the policy, recall that we considered a sequence of systems indexed by the number of jobs  $n$ , whose formal limit was the Brownian control problem (2.48)–(2.53) under diffusion scaling. To articulate the proposed policy, we fix the system parameter  $n$  and use it to unscale processes of interest. We define the (unscaled) workload process  $W^n = \{W^n(t), t \geq 0\}$  as follows:

$$W^n(t) := \sum_{i=1}^I Q_i^n(t), \quad t \geq 0.$$

**Proposed Pricing Policy:** Given the workload process  $W^n$ , we choose the demand rates

$$\lambda_i^n(t) := n\lambda_i^* + \frac{\sqrt{n}}{2\alpha_i} v\left(\frac{W^n(t)}{\sqrt{n}}\right), \quad i = 1, \dots, I, \quad t \geq 0,$$

where  $v$  is the solution to the Bellman equation (2.90)–(2.91). This follows from (2.43), (2.92), Lemma 9, and Theorem 4. The corresponding proposed pricing policy is given by

$$p_i^n(t) := \Lambda_i^{-1}(\lambda_i^*) + \frac{(\Lambda_i^{-1})'(\lambda_i^*)}{2\alpha_i\sqrt{n}} v\left(\frac{W^n(t)}{\sqrt{n}}\right), \quad i = 1, \dots, I, \quad t \geq 0, \quad (2.143)$$

where  $\Lambda_i^{-1}$  is the inverse of the demand rate function for region  $i$ . Equation (2.143) is derived in Appendix 2.11.1.

**Proposed Dispatch Policy:** We propose two dispatch policies and refer to them as Dispatch Policy 1 (DP1) and Dispatch Policy 2 (DP2). Dispatch Policy 1 (DP1) is motivated by the following observation. In the Brownian control problem under the complete resource pooling assumption, we set all but one of the inventory levels to zero. (The buffer with nonzero inventory corresponds to the one with lowest holding cost.) However, as articulated in Harrison [1996], zero inventory in the Brownian control problem corresponds to small positive inventory levels in the original system. Thus, we put small safety stocks in the various buffers and only serve them when inventory levels are at or above the threshold. To that end, denote by  $s_i$  the safety stock for buffer  $i$ .

To be more specific, letting  $\bar{\mathcal{A}}_i := \mathcal{A}_i \cap \{1, \dots, b\}$  denote the set of basic activities undertaken by server  $i$  and letting  $\bar{\mathcal{C}}_i := \mathcal{C}_i \cap \{1, \dots, b\}$  denote the set of basic activities that serve buffer  $i$ , our proposed dispatch policy is as follows: If server  $i$  becomes idle at time  $t$ , it serves a job from the buffer in  $\{b(j) : j \in \bar{\mathcal{A}}_i, Q_{b(j)}^n(t) \geq s_{b(j)}\}$  with largest holding cost  $h_{b(j)}$ . In words, when server  $i$  becomes idle, it looks at all buffers it servers by means of basic activities and serves the buffer with largest holding cost that is above its safety stock. To complete the policy description, suppose that at time  $t$  the inventory in buffer  $i$  increases from  $s_i - 1$  to  $s_i$ , i.e., reaches the safety stock. The system manager serves buffer  $i$  by an idle server in  $\{s(j) : j \in \bar{\mathcal{C}}_i\}$  with largest effective idling cost  $c_{s(j)}/\lambda_{s(j)}^*$  (see (2.79)). In words, when buffer  $i$  reaches the safety stock, i.e., that buffer becomes eligible for service, the system manager selects an idle server with largest effective idling cost than can serve the buffer by means of a basic activity.

Dispatch Policy 2 (DP2) is motivated by the maximum pressure policy, see for example Stolyar [2004], Dai and Lin [2005], Dai and Lin [2008], and Ata and Lin [2008]. Under this policy, each server prioritizes his own (local) buffer. If his own buffer is empty, then he checks

the other buffers that he can serve using basic activities. If there are multiple such buffers, the server works on the buffer with the largest queue length. If the server’s own (local) buffer is empty and he cannot serve any other buffers using basic activities, then he considers all remaining buffers he can serve (using nonbasic activities) and works next on the buffer with the largest queue length.

## 2.9 Simulation Study

This section presents a simulation study to illustrate the effectiveness of the proposed policy. The simulation setting and its parameters are motivated, albeit loosely, by the taxi market in Manhattan, see Ata et al. [2020b] and the references therein. We set the number of cars, i.e., the system parameter, as  $n = 10,000$ . As done in Ata et al. [2020b], we divide Manhattan into  $I = 4$  regions, see Figure 2.2.



Figure 2.2: Manhattan area that is partitioned into four regions.

We assume cars can pick up customers in their own regions as well as from the neighboring

regions. This gives rise to the following capacity consumption matrix:

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

Using the same dataset in Ata et al. [2020b], we set the demand rate (per hour) vector as follows:<sup>6</sup>

$$\lambda^n = (\lambda_1^n, \lambda_2^n, \lambda_3^n, \lambda_4^n)' = (3678, 10723, 6792, 345)'.$$

The corresponding limiting rate vector  $\lambda^*$  is then computed as  $\lambda^* = \lambda^n/n$ , which yields

$$\lambda^* = (\lambda_1^*, \lambda_2^*, \lambda_3^*, \lambda_4^*)' = (0.367, 1.072, 0.679, 0.0345)'. \quad (2.144)$$

Using this and (2.25), we derive the input-output matrix  $R$  as follows:

$$R = \begin{bmatrix} \lambda_1^* & 0 & 0 & 0 & 0 & \lambda_2^* & 0 & 0 & 0 & 0 \\ 0 & \lambda_2^* & 0 & 0 & \lambda_1^* & 0 & 0 & \lambda_3^* & 0 & 0 \\ 0 & 0 & \lambda_3^* & 0 & 0 & 0 & \lambda_2^* & 0 & 0 & \lambda_4^* \\ 0 & 0 & 0 & \lambda_4^* & 0 & 0 & 0 & 0 & \lambda_3^* & 0 \end{bmatrix}.$$

Ata et al. [2020b] reports the mean travel time as 13.2 minutes. To account for the pick up time and for other inefficiencies that are not incorporated in our model, we inflate this by a factor of two, and set the mean trip time to 26.4 minutes. Thus  $\eta^n = 2.2727$  per hour. Moreover, because we study the system under the heavy traffic assumption (Assumption 5), we set  $\eta = e'\lambda^* = 2.1539$ . Therefore, we have that  $\hat{\eta} = \sqrt{n}(\eta^n - \eta) = 11.88$ .

---

6. For simplicity, we use the preliminary results from Ata et al. [2020b] to estimate  $\lambda^n$  and  $q$  (based on a four-year dataset from January 2010 to December 2013). In doing so, we focus on the day shift of the non-holiday weekdays.

We estimate the routing probability vector  $q$  from the data as

$$q = (q_1, q_2, q_3, q_4)' = (0.1647, 0.5408, 0.2724, 0.0221)',$$

which yields the limiting arrival rate vector  $\nu$  to various buffers as follows:

$$\nu = \eta q = (0.3529, 0.1159, 0.5837, 0.0474)'$$

Thus using the data  $A, R$ , and  $\gamma$ , one can compute the unique nominal processing plan  $x^*$ , referred to in Assumption 5. It is displayed in Figure 2.3. Having characterized  $x^*$ , we next

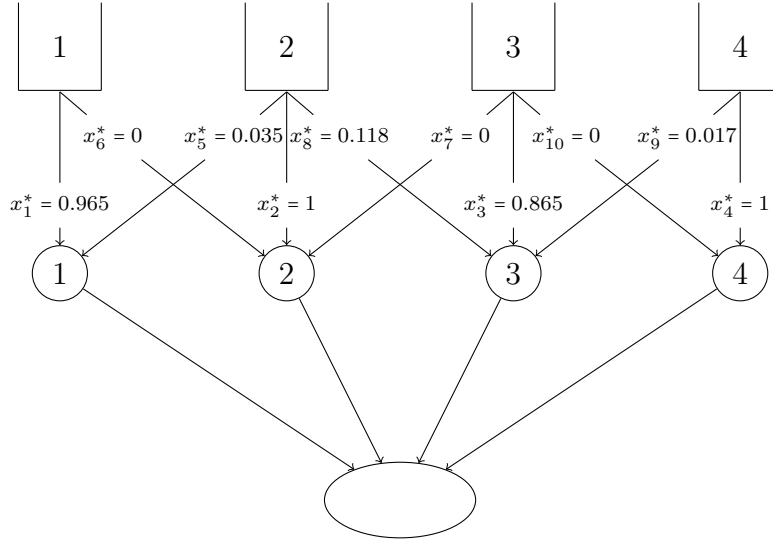


Figure 2.3: Unique solution  $x^* \in \mathbb{R}^{10}$  to the static problem from (2.28)–(2.30). We see that Activities 6, 7, and 10 are nonbasic while the rest are basic.

compute the drift parameter  $a$  and the variance parameter  $\sigma^2$  of the Brownian motion  $\chi(\cdot)$  (see (2.73)). To this end, first note that the drift vector  $\gamma$  and the covariance matrix  $\Sigma$  of the Brownian motion  $B(\cdot)$  (see (2.50), (2.56), and (2.57)) are given as follows:

$$\gamma = \hat{\eta}' q = (1.9566, 6.4247, 3.2361, 0.2625)',$$

$$\Sigma = \begin{bmatrix} 0.7097 & 0.1918 & 0.0966 & 0.0078 \\ 0.1918 & 2.3302 & 0.3173 & 0.0257 \\ 0.0966 & 0.3173 & 1.1742 & 0.0130 \\ 0.0078 & 0.0257 & 0.0130 & 0.0937 \end{bmatrix}.$$

Thus, we have that  $a = e'\gamma = 11.88$  and  $\sigma^2 = e'\Sigma e = 5.6125$ .

Next, we describe the economic primitives of our example: the demand function, and its associated profit function, the holding cost rates and the cost of idleness. We assume that the demand function is linear. That is,

$$\Lambda_i(p_i) = a_i - b_i p_i, \quad p_i \in [0, a_i/b_i] \quad \text{and} \quad i = 1, \dots, 4,$$

where  $a_i, b_i > 0$  are constants. Also, its inverse is given by

$$\Lambda_i^{-1}(\lambda_i) = \frac{a_i - \lambda_i}{b_i}, \quad \lambda_i \in [0, a_i] \quad \text{and} \quad i = 1, \dots, 4.$$

The profit function then follows from (2.7) as follows:

$$\pi(\lambda) = \sum_{i=1}^4 \frac{\lambda_i}{b_i} (a_i - \lambda_i), \quad \lambda_i \in [0, a_i] \quad \text{and} \quad i = 1, \dots, 4.$$

We set the optimal static price as  $p_i^* = 10$  for all region  $i$ , which is about the average price of a ride in the data, see Ata et al. [2020b]. Also, recall that the limiting demand rate vector  $\lambda^* = (\lambda_1^*, \dots, \lambda_4^*)$  is given by (2.144). We crucially assume that these are the optimal demand rate and the prices. This is equivalent to assuming  $a_i = 2\lambda_i^*$  and  $b^* = \lambda_i^*/p_i$  for  $i = 1, \dots, 4$ .

Namely, we set

$$\begin{aligned} a &= 2\lambda^* = (0.7356, 2.1446, 1.3584, 0.0691)', \\ b &= \lambda^*/p^* = (0.0367, 0.1072, 0.0679, 0.0035)'. \end{aligned}$$

Given these we compute the parameter  $\alpha_i$  as  $\alpha_i = -(\Lambda_i^{-1})'(\lambda_i^*) - (\lambda_i^*/2)(\Lambda_i^{-1})''(\lambda_i^*) = 1/b_i$  for  $i = 1, \dots, 4$ . Thus, we obtain  $\alpha = (27.18, 9.32, 14.72, 289.55)$  and  $\hat{\alpha} = \sum_{i=1}^4 1/\alpha_i = 0.2154$ .

Ata et al. [2020b] suggest that the holding cost when taxis are traveling is  $h_0^n = 1$  dollars per hour (which can be derived from their fuel cost estimates). To estimate the holding cost rates for other buffers, we consider the driver's opportunity cost. A driver can complete about two trips per hour, resulting in approximately  $2 \times 10 = 20$  dollars per hour. Thus, we set  $h_i^n = 20$  for  $i = 1, \dots, 4$ . Thus, we have  $h^n = \min_{i=1, \dots, 4} h_i^n - h_0^n = 19$ . Upon scaling, we derive the limiting holding cost rate  $h$  for the equivalent workload formulation as  $h = \sqrt{n}h^n = 1900$ . The idleness costs parameters are set to equal the lost revenue. That is,  $c_i^n = p_i^* = 10$  for  $i = 1, \dots, 4$ . Upon rescaling, the limiting idleness cost is  $c_i = c_i^n/\sqrt{n} = 0.1$ . Thus, the cheapest server to idle as  $k^* = \arg \min_{i=1, \dots, 4} c_i/\lambda_i^* = 2$  with the idling cost  $r = c_{k^*}/\lambda_{k^*}^* = 0.0933$ .

Having computed the parameters  $a, \sigma^2, h, r, \eta$ , and  $\hat{\alpha}$ , we solve the Bellman equation numerically for the example. Using this solution, we next describe our proposed policy.

**Pricing Policy.** It follows from (2.143) that

$$p_i^n(t) = 10 - \frac{1}{200}v\left(\frac{W^n(t)}{100}\right), \quad i = 1, \dots, 4 \quad \text{and} \quad t \geq 0.$$

This corresponds to the following demand rates:

$$\lambda_i^n = 10000\lambda_i^* + \frac{50}{\alpha_i}v\left(\frac{W^n(t)}{100}\right), \quad i = 1, \dots, 4 \quad \text{and} \quad t \geq 0.$$

**Dispatch Policy.** As discussed in Section 2.8, we propose two dispatch policies. Under the first proposed policy (Dispatch Policy 1), servers 2 and 4 work only on their own buffer throughout. Servers 1 and 3 prioritize their own buffers, but server 1 serves buffer 2 if buffer 1 is empty and buffer 2 exceeds threshold  $s$ . Similarly, server 3 serves buffers 2 or 4 only if buffer 3 is empty and buffer 2 or 4 exceeds threshold  $s$ . If both queues exceeds  $s$ , then server 3 serves the longest one. We determine the threshold  $s$  by a brute-force search. In particular,



we set  $s = 1$ .

Under Dispatch Policy 2, each server prioritizes his own (local) buffer. If his own buffer is empty, then he checks the other buffers that he can serve using basic activities. If there are multiple such buffers, the server works on the buffer with the largest queue length. If the server's own (local) buffer is empty and he cannot serve any other buffers using basic activities, then he considers all remaining buffers he can serve (using nonbasic activities) and works next on the buffer with the largest queue length.

In order to compare the performance of our policy, we calculate the total revenue by adding up the prices charged to each served customer. This also incorporates the cost of idleness. Also, we keep track of the holding costs incurred. Lastly, using (2.45), we set

$$\tilde{V}^n(t) = (n\pi(\lambda^*) - \sqrt{nh_0})t = \left( n \sum_{i=1}^4 \frac{\lambda_i^*}{b_i} (a_i - \lambda_i^*) - \sqrt{nh_0} \right) t, \quad t \geq 0$$

to compute the normalized cost  $\widehat{V}^n(t)$ , see (2.46).

We compare our policy against the following benchmark policies that combine alternative pricing and dispatch policies. For pricing, in addition to our dynamic pricing policy, we also consider the static pricing policy which sets  $p_i^n(t) = p_i^* = 10$  for all  $i = 1, \dots, 4$  and  $t \geq 0$ . For dispatch, in addition to our two proposed policies, we consider (i) a static dispatch policy, and (ii) the closest driver policy as described next.

**Static Dispatch Policy.** Servers 2 and 4 always serve their own buffers. If both buffers 1 and 2 are nonempty, then server 1 works on buffer 1 with probability  $x_1^*/(x_1^* + x_5^*) = 0.965$  and it works on buffer 2 with probability  $x_5^*/(x_1^* + x_5^*) = 0.035$ . If only one of the buffers 1 and 2 is nonempty, then server 1 works on that buffer. Server 3 splits its effort among buffers 2, 3, and 4 similarly, i.e., proportional to  $x_3^*, x_8^*$ , and  $x_9^*$ , respectively.

**Closest Driver Policy.** We let  $D$  be the distance matrix, i.e.,  $D_{ij}$  corresponds to the distance (in miles) between regions  $i$  and  $j$  when  $i \neq j$  and  $D_{ii} = 0$ . Using the data from Ata

et al. [2020b], we have

$$D = \begin{bmatrix} 0 & 2.6414 & 4.8132 & 8.2689 \\ 2.6414 & 0 & 1.9993 & 6.1969 \\ 4.8132 & 1.9993 & 0 & 3.9073 \\ 8.2689 & 6.1969 & 3.9073 & 0 \end{bmatrix}.$$

Server  $i$  engages in activity  $\operatorname{argmin}_{j \in \mathcal{A}_i} D_{i,b(j)}(t)$  at time  $t$ . In other words, under the closest driver policy each server prioritizes the buffer that is closest to him.

The result of the numerical study are given in Table 2.1. The simulated results are obtained based on a run-length of 1000 hours and the estimated average cost is computed by excluding the statistics from the first 200 hours warm-up period. The corresponding confidence intervals are calculated based on 10 macro-replications. We observe that the proposed dispatch policies (DP1, DP2) offer significant improvement (9.74%–55.01%) over the benchmark policies. More importantly, we observe that dynamic pricing can lead to significant improvement (30.96%–61.73%) for every dispatch policy considered. Among the policies considered, the dynamic pricing with Dispatch Policy 2 (DP2) has the best performance.

Table 2.1: Estimated average cost along with the 95% confidence interval based on 10 macro-replications.

Dispatch policy	Static pricing policy	Dynamic pricing policy
DP1	10075.23 ± 201.59	4302.59 ± 94.09
DP2	10607.19 ± 103.18	4059.35 ± 73.73
Static policy	13066.83 ± 457.31	9021.89 ± 204.19
Closest driver policy	12100.53 ± 193.57	4766.96 ± 122.19

Unfortunately, we do not have any data to directly estimate the holding costs and the cost of idleness. For the former, the actual holding cost may be lower because the opportunity cost we estimate is likely an upper bound. On the other hand, the latter does not account for the loss of goodwill currently. Therefore, we conduct a sensitivity analysis that considers

lower holding cost rates (Figure 2.4) and another one that considers higher cost of idleness that incorporate the loss of goodwill<sup>7</sup> (Figure 2.5). These collectively show that the insights from Table 2.1 are robust to changes in holding and idleness cost parameters.

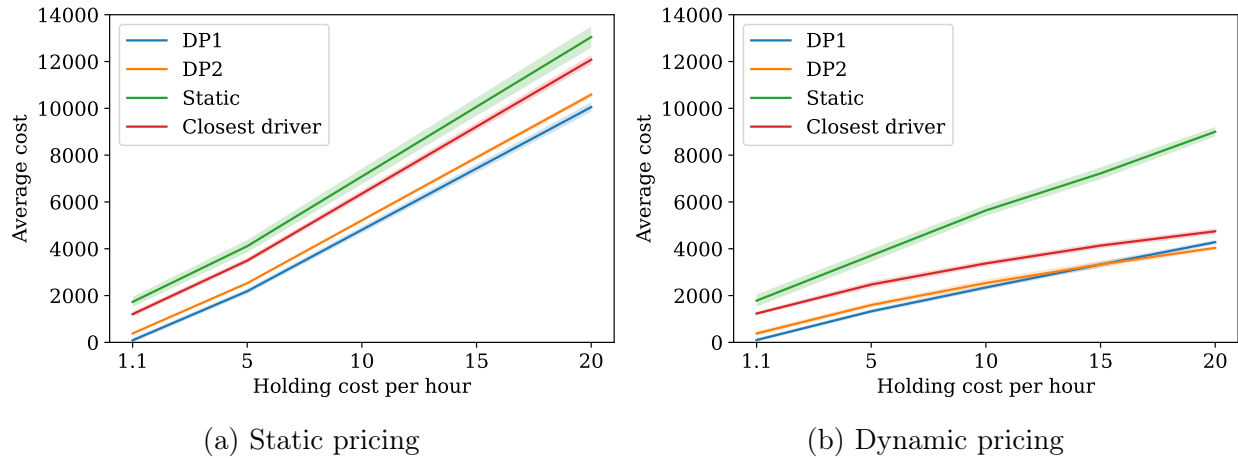


Figure 2.4: Average cost with respect to varying holding cost. The shaded area along each line shows the 95% confidence interval based on 10 macro-replications.




Figure 2.5: Average cost with respect to varying idleness cost. The shaded area along each line shows the 95% confidence interval based on 10 macro-replications.

7. The estimated performance and the corresponding confidence interval for the sensitivity analysis is based on 10 macro-replications where each replication has a run-length of 1000 hours (and the statistics of the first 200 hours are discarded as a warm-up period).

## 2.10 Concluding Remarks

We study a dynamic pricing and dispatch control problem motivated by ride-hailing systems. The novelty of our formulation is that it incorporates travel times. We solve this problem analytically in the heavy traffic regime under the complete resource pooling condition. Using this solution, we propose a closed form dynamic pricing policy as well as a dispatch policy. We compare the proposed policy against benchmarks in a simulation study and show that it is effective.

Our formulation has some limitations too. Namely, we assume there is only one travel node and that the complete resource pooling condition holds. Interesting future research directions include relaxing these assumptions.

## 2.11 Appendix

### 2.11.1 Formal Derivations

#### Formal Derivation of the Brownian Control Problem

This section provides a formal derivation of the approximating Brownian control problem introduced in Section 2.4. We do not provide a rigorous weak convergence limit theorem. However, the arguments given in support of the approximation can be viewed as a broad outline for such a proof; see Harrison [1988, 2000, 2003] for similar derivations.

We consider a sequence of systems indexed by the system parameter  $n$  under the heavy traffic assumption. Then we center the various processes by their mean, scale them appropriately by the system parameter  $n$ , and finally pass to the limit as  $n \rightarrow \infty$  formally. To that end, we first define the following (diffusion) scaled processes:

$$\widehat{\Psi}_i^n(t) := \frac{1}{\sqrt{n}} (\Psi_i(\lfloor nt \rfloor) - q_i nt), \quad t \geq 0, \quad i = 1, \dots, I, \quad (2.145)$$

$$\widehat{N}_j^n(t) := \frac{1}{\sqrt{n}}(N_j(nt) - nt), \quad t \geq 0, \quad j = 0, 1, \dots, J, \quad (2.146)$$

where  $[x]$  is the greatest integer less than or equal to  $x$ . We also define the following (fluid) scaled processes:

$$\bar{N}_0^n(t) := \frac{1}{n}N_0(nt), \quad t \geq 0, \quad (2.147)$$

$$\bar{Q}_0^n(t) := \frac{1}{n}Q_0^n(t), \quad t \geq 0, \quad (2.148)$$

$$\bar{\mu}_j^n(t) := \frac{1}{n}\mu_j^n(t), \quad j = 1, \dots, J, \quad t \geq 0. \quad (2.149)$$

By Donsker's theorem, the functional central limit theorem for renewal processes, and independence of the stochastic primitives, the processes  $\widehat{\Psi}_i^n$  and  $\widehat{N}_j^n$  converge weakly to independent standard Brownian motions, see Billingsley [1999].

As observed in Kogan and Lipster [1993], under the heavy traffic assumption, we expect that the number of jobs in the infinite-server node will be  $n$  to a first-order approximation. That is, we expect that  $\bar{Q}_0^n(t) \approx 1$  for  $t \geq 0$  as  $n$  gets large. Similarly, we expect the queue lengths at buffers  $1, \dots, I$  to be of order  $\sqrt{n}$ . As such, we expect the prices, or equivalently, the demand rates, to deviate from their nominal values only in the second order. That is, we expect  $\lambda_i^n - \lambda_i^*n = O(\sqrt{n})$ . Because the demand rates determine the service rates (in particular, recall (2.9)), we expect that  $\bar{\mu}_j^n(t) \approx \mu_j^*$  for  $t \geq 0$  as  $n$  gets large.

By combining (2.145)–(2.149) with (2.38)–(2.44), it is straightforward to derive the following scaled system dynamics equations for  $i = 1, \dots, I$ :

$$\begin{aligned} Z_i^n(t) &= B_i^n(t) + q_i\eta^n \int_0^t Z_0^n(s) ds - \sum_{j \in \mathcal{C}_i} \int_0^t \kappa_j^n(s) dT_j^n(s) + \sum_{j \in \mathcal{C}_i} \mu_j^* Y_j^n(t) \\ &\quad + t\sqrt{n} \left[ q_i\eta - \sum_{j \in \mathcal{C}_i} \mu_j^* x_j^* \right] \\ &= B_i^n(t) + q_i\eta^n \int_0^t Z_0^n(s) ds - \sum_{j \in \mathcal{C}_i} \int_0^t \kappa_j^n(s) dT_j^n(s) + \sum_{j \in \mathcal{C}_i} \mu_j^* Y_j^n(t), \end{aligned}$$

where the second equality holds by Assumption 5 and where the process  $B_i^n$  is given by

$$B_i^n(t) = Z_i^n(0) + q_i \hat{\eta} t + q_i \widehat{N}_0^n \left( \eta^n \int_0^t \bar{Q}_0^n(s) ds \right) + \widehat{\Psi}_i^n \left( \bar{N}_0^n \left( \eta^n \int_0^t \bar{Q}_0^n(s) ds \right) \right) - \sum_{j \in \mathcal{C}_i} \widehat{N}_j^n \left( \int_0^t \bar{\mu}_j^n(s) dT_j^n(s) \right).$$

Assuming that  $Z_i^n(0) \approx Z_i(0)$  for large  $n$ , it is also straightforward to argue that  $B_i^n$  can be approximated by a Brownian motion  $B_i$  with starting state  $Z_i(0)$  that has drift parameter  $\gamma_i = \hat{\eta} q_i$  and variance parameter

$$\sigma_i^2 := [q_i^2 + q_i(1 - q_i)] \eta + \sum_{j \in \mathcal{C}_i} \mu_j^* x_j^* = q_i \eta + \sum_{j \in \mathcal{C}_i} \mu_j^* x_j^*.$$

Furthermore, the covariance between the limiting Brownian motion processes is given by

$$\text{Cov}(B_i, B_{i'}) = q_i q_{i'} \eta \quad \text{for } i \neq i'.$$

Therefore, replacing  $Z^n$ ,  $Y^n$ , and  $\kappa^n$ , by their formal limits  $Z$ ,  $Y$ , and  $\kappa$ , we arrive at the following system dynamics equations in the approximating Brownian control problem for  $i = 1, \dots, I$ :

$$Z_i(t) = B_i(t) + q_i \eta \int_0^t Z_0(s) ds - \sum_{j \in \mathcal{C}_i} \int_0^t x_j^* \kappa_j(s) ds + \sum_{j \in \mathcal{C}_i} \mu_j^* Y_j(t), \quad t \geq 0.$$

Equations (2.19) and (2.39) of the system state also imply that  $Z_0^n(t) = -\sum_{i=1}^I Z_i^n(t)$  and that  $Z_i^n(t) \geq 0$  for  $i = 1, \dots, I$  and  $t \geq 0$ . Thus, in the approximating BCP, the following relationships hold for  $t \geq 0$ :

$$Z_0(t) = -\sum_{i=1}^I Z_i(t) \quad \text{and} \quad Z_i(t) \geq 0 \quad \text{for } i = 1, \dots, I.$$

Similarly, it is clear that (2.9) and (2.43)–(2.44) give rise to (2.53) in the BCP; (2.17) and (2.41) give rise to (2.55); and (2.42) gives rise to (2.52).

To complete the formal derivation of the Brownian control problem, we argue that  $\widehat{V}^n \approx \xi$  for large  $n$ , where  $\widehat{V}^n$  and  $\xi$  are given by (2.46) and (2.47), respectively. First, observe that by Taylor's theorem we have

$$\begin{aligned} \pi\left(\lambda^* + \frac{1}{\sqrt{n}}\zeta^n(s)\right) &= \pi(\lambda^*) + \nabla\pi(\lambda^*)' \frac{1}{\sqrt{n}}\zeta^n(s) + \frac{1}{2n}\zeta^n(s)' \nabla^2\pi(\lambda^*) \zeta^n(s) \\ &\quad + R_{\lambda^*,3}\left(\frac{1}{\sqrt{n}}\zeta^n(s)\right), \end{aligned}$$

where  $R_{\lambda^*,3}\left(\frac{1}{\sqrt{n}}\zeta^n(s)\right) = O(n^{-3/2})$  is a third-order remainder term.<sup>8</sup> Moreover, note that the term  $\nabla\pi(\lambda^*)' \zeta^n(s)/\sqrt{n}$  vanishes because  $\lambda^*$  is a maximizer of  $\pi(\lambda)$  and is in the interior of the feasible region  $\mathcal{L}$  (see Assumption 4), implying that  $\nabla\pi(\lambda^*) = 0$ . Therefore, we have that

$$\pi\left(\lambda^* + \frac{1}{\sqrt{n}}\zeta^n(s)\right) = \pi(\lambda^*) - \frac{1}{n}\zeta^n(s)' H \zeta^n(s) + O(n^{-3/2}),$$

where  $H := -\frac{1}{2}\nabla^2\pi(\lambda^*)$ . Using this and (2.35) and (2.43), it follows that

$$\pi^n(\lambda^n(s)) = n\pi(\lambda^*) - \zeta^n(s)' H \zeta^n(s) + O(n^{-1/2}). \quad (2.150)$$

Finally, using (2.39), (2.41)–(2.46), and (2.150), it is straightforward to derive the following:

$$\begin{aligned} \widehat{V}^n(t) &= n(\pi(\lambda^*) - h_0^n)t - \left[ \int_0^t \pi^n(\lambda^n(s)) \, ds - \int_0^t \sum_{i=0}^I h_i^n Q_i^n(s) \, ds - (c^n)' I^n(t) \right] \\ &= \int_0^t \left[ \zeta^n(s)' H \zeta^n(s) + O(n^{-1/2}) \right] \, ds + \int_0^t \sum_{i=0}^I h_i Z_i^n(s) \, ds + c' U^n(t). \end{aligned}$$

---

8. In particular, the remainder term is given by

$$R_{\lambda^*,3}\left(\frac{1}{\sqrt{n}}\zeta^n(s)\right) = \sum_{\substack{\alpha_1, \dots, \alpha_I \in \{0,1,2,3\} \\ \text{s.t. } \alpha_1 + \dots + \alpha_I = 3}} \frac{\partial^3 \pi\left(\lambda^* + \frac{C}{\sqrt{n}}\zeta^n(s)\right)}{\partial x_1^{\alpha_1} \partial x_2^{\alpha_2} \dots \partial x_I^{\alpha_I}} \prod_{i=1}^I \frac{\left(\frac{1}{\sqrt{n}}\zeta_i^n(s)\right)^{\alpha_i}}{\alpha_i!} \quad \text{for some } C \in (0,1).$$

Therefore, replacing  $\widehat{V}^n$ ,  $Z^n$ ,  $\zeta^n$ , and  $U^n$  by their formal limits  $\xi$ ,  $Z$ ,  $\zeta$ , and  $U$ , we arrive at the following cost process of the approximating Brownian control problem:

$$\xi(t) = \int_0^t \zeta(s)' H \zeta(s) ds + \int_0^t \sum_{i=0}^I h_i Z_i(s) ds + c' U(t), \quad t \geq 0.$$

Furthermore, note that  $H$  is a diagonal matrix, i.e.,  $H = \text{diag}(\alpha_1, \dots, \alpha_I)$  where

$$\alpha_i := -(\Lambda_i^{-1})(\lambda_i^*) - \frac{\lambda_i^*}{2} (\Lambda_i^{-1})''(\lambda_i^*), \quad i = 1, \dots, I.$$

Using this we further simplify the limiting cost process  $\xi(t)$  as follows:

$$\xi(t) = \int_0^t \left[ \sum_{i=1}^I \alpha_i \zeta_i^2(s) + \sum_{i=0}^I h_i Z_i(s) \right] ds + c' U(t), \quad t \geq 0.$$

### Formal Derivation of (2.143)

Recall that the proposed chosen demand rates are

$$\lambda_i^n = n\lambda_i^* + \sqrt{n}q(t), \quad i = 1, \dots, I, \quad t \geq 0,$$

where  $q(t) := \frac{1}{2\alpha_i} v\left(\frac{W^n(t)}{\sqrt{n}}\right)$  for  $t \geq 0$ . Therefore, the proposed pricing policy for region  $i$  is given as follows:

$$p_i^n(t) = (\Lambda_i^n)^{-1}(\lambda_i^n(t)) = (\Lambda_i^n)^{-1}(n\lambda_i^* + \sqrt{n}q(t)) = \Lambda_i^{-1}\left(\lambda_i^* + \frac{q(t)}{\sqrt{n}}\right),$$

where the third equality follows from the fact that  $(\Lambda_i^n)^{-1}(nx) = \Lambda_i^{-1}(x)$  for  $x \in \mathcal{L}$ . Then, by Taylor's theorem, there exists a  $c \in (0, 1)$  such that

$$p_i^n(t) = \Lambda_i^{-1}(\lambda_i^*) + (\Lambda_i^{-1})'(\lambda_i^*) \frac{q(t)}{\sqrt{n}} + \frac{1}{2} (\Lambda_i^{-1})''\left(\lambda_i^* + \frac{cq(t)}{\sqrt{n}}\right) \left(\frac{q(t)}{\sqrt{n}}\right)^2,$$



which implies that

$$\begin{aligned} p_i^n(t) &= \Lambda_i^{-1}(\lambda_i^*) + (\Lambda_i^{-1})'(\lambda_i^*) \frac{q(t)}{\sqrt{n}} + O\left(\frac{1}{n}\right) \\ &= \Lambda_i^{-1}(\lambda_i^*) + \frac{(\Lambda_i^{-1})'(\lambda_i^*)}{2\alpha_i\sqrt{n}} v\left(\frac{W^n(t)}{\sqrt{n}}\right) + O\left(\frac{1}{n}\right). \end{aligned}$$

As an aside, observe that by rearranging terms we have

$$\sqrt{n}(p_i^n(t) - \Lambda_i^{-1}(\lambda_i^*)) = \frac{(\Lambda_i^{-1})'(\lambda_i^*)}{2\alpha_i} v\left(\frac{W^n(t)}{\sqrt{n}}\right) + O\left(\frac{1}{\sqrt{n}}\right).$$

This implies that our proposed dynamic pricing policy coincides with the static prices to a first-order approximation, but deviates from the static prices on the second order, i.e., order  $1/\sqrt{n}$ .

### 2.11.2 Miscellaneous Proofs

*Proof of Lemma 6.* This proof follows in an almost identical fashion to Lemma 2 in Ata et al. [2020a], but we include it for completeness. It consists of four steps. We let  $e^n$  denote the  $n$ th unit basis vector in a Euclidean space of appropriate dimension. That is, the  $n$ th component of  $e^n$  is one, whereas its other components are zero. Moreover, recall from the discuss following Assumption 5 that for a vector  $y \in \mathbb{R}^J$ , we write  $y = (y_B, y_N)$  where  $y_B \in \mathbb{R}^b$  and  $y_N \in \mathbb{R}^{J-b}$ .

**Step 1:** Consider the set of basic activity rates that do not cause any server idleness, i.e.,  $\{y \in \mathbb{R}^J : By_B = 0, y_N = 0\}$ . First, we show that this set is the span of  $\bar{C}$ , defined next:

$$\bar{C} := \left\{ e^j - e^{j'} : (j, j') \in \bar{C} \text{ and } e^j, e^{j'} \text{ are unit basis vectors in } \mathbb{R}^J \right\}, \quad (2.151)$$

where  $\bar{C} := \{(j, j') : j, j' \in \{1, \dots, b\} \text{ such that } s(j) = s(j')\}$ . That is,  $\bar{C}$  is the set of all pairs of basic activities undertaken by the same server. Note that the difference  $e^j - e^{j'}$  in

(2.151) captures the trade-off server  $s(j)$  makes by increasing the rate at which activity  $j$  is undertaken (from its nominal value  $x_j^*$ ) at the expense of decreasing the rate of activity  $j'$ . By making such adjustments to the nominal basic activity rates  $x_B^*$ , the system manager can redistribute the workload between buffers  $b(j)$  and  $b(j')$  without incurring any idleness. As such, we intuitively expect that taking linear combinations of such activity rates in  $\bar{C}$  should yield the set of activity rates that do not result in any idleness, i.e., the set  $\{y \in \mathbb{R}^J : By_B = 0, y_N = 0\}$ . In summary, in Step 1 we prove that

$$\text{span}(\bar{C}) = \{y \in \mathbb{R}^J : By_B = 0, y_N = 0\}.$$

To prove this, we show inclusions of both sets. First, let  $y \in \{y \in \mathbb{R}^J : By_B = 0, y_N = 0\}$ . To prove that  $y \in \text{span}(\bar{C})$ , we show that there exist constants  $a_{jj'}$ ,  $(j, j') \in \bar{C}$ , such that  $y = \sum_{(j, j') \in \bar{C}} a_{jj'} (e^j - e^{j'})$ . To find these constants, it will be convenient to define the sets

$$\bar{\mathcal{A}}_i := \mathcal{A}_i \cap \{1, \dots, b\},$$

where  $\mathcal{A}_i$  is the set of activities undertaken by server  $i$  (see (2.3)). To be more specific,  $\bar{\mathcal{A}}_i$  consists of all basic activities undertaken by server  $i$ . After possibly relabeling, suppose that the basic activities are ordered so that

$$\bar{\mathcal{A}}_i = \{b_{i-1} + 1, \dots, b_i\} \quad \text{for } i = 1, \dots, I,$$

where  $0 = b_0 < b_1 < b_2 < \dots < b_I = b$ . We define constants  $a_{jj'}$  for  $(j, j') \in \bar{C}$  as follows:

$$a_{jj'} := \begin{cases} \sum_{l=b_{i-1}+1}^k y_l, & \text{if } (j, j') = (k, k+1) \text{ for } k = b_{i-1} + 1, \dots, b_i - 1 \text{ and } i = 1, \dots, I, \\ 0, & \text{otherwise.} \end{cases}$$

Therefore, we have that

$$\begin{aligned}
\sum_{(j,j') \in \bar{C}} a_{jj'} (e^j - e^{j'}) &= \sum_{i=1}^I \sum_{k=b_{i-1}+1}^{b_i-1} a_{k,k+1} (e^k - e^{k+1}) \\
&= \sum_{i=1}^I \sum_{k=b_{i-1}+1}^{b_i-1} \left[ \binom{k}{l=b_{i-1}+1} y_l \right] (e^k - e^{k+1}) \\
&= \sum_{i=1}^I \left[ y_{b_{i-1}+1} (e^{b_{i-1}+1} - e^{b_{i-1}+2}) + (y_{b_{i-1}+1} + y_{b_{i-1}+2}) (e^{b_{i-1}+2} - e^{b_{i-1}+3}) \right. \\
&\quad \left. + \dots + \binom{b_i-2}{l=b_{i-1}+1} y_l (e^{b_i-2} - e^{b_i-1}) + \binom{b_i-1}{l=b_{i-1}+1} y_l (e^{b_i-1} - e^{b_i}) \right] \\
&= \sum_{i=1}^I \left[ y_{b_{i-1}+1} e^{b_{i-1}+1} + (y_{b_{i-1}+2} + y_{b_{i-1}+1} - y_{b_{i-1}+1}) e^{b_{i-1}+2} \right. \\
&\quad \left. + \dots + \left( \sum_{l=b_{i-1}+1}^{b_i-1} y_l - \sum_{l=b_{i-1}+1}^{b_i-2} y_l \right) e^{b_i-1} - \left( \sum_{l=b_{i-1}+1}^{b_i-1} y_l \right) e^{b_i} \right] \\
&= \sum_{i=1}^I \left[ y_{b_{i-1}+1} e^{b_{i-1}+1} + y_{b_{i-1}+2} e^{b_{i-1}+2} + \dots + y_{b_i-1} e^{b_i-1} - \left( \sum_{l=b_{i-1}+1}^{b_i-1} y_l \right) e^{b_i} \right] \\
&= \sum_{i=1}^I \left[ y_{b_{i-1}+1} e^{b_{i-1}+1} + y_{b_{i-1}+2} e^{b_{i-1}+2} + \dots + y_{b_i-1} e^{b_i-1} - (-y_{b_i}) e^{b_i} \right] \\
&= \sum_{i=1}^I \sum_{k=b_{i-1}+1}^{b_i} y_k e^k \\
&= \sum_{j=1}^J y_j e^j,
\end{aligned}$$

the first two equalities following from the definition of the  $a_{jj'}$ , the fourth equality from algebraic rearrangements, and the fifth equality from canceling terms. To derive the sixth equality note that  $y$  satisfies  $By_B = 0$ , which implies

$$\sum_{l=b_{i-1}+1}^{b_i} y_l = 0 \quad \text{for } i = 1, \dots, I.$$

Equivalently, we have that

$$\sum_{l=b_{i-1}+1}^{b_i-1} y_l = -y_{b_i} \quad \text{for } i = 1, \dots, I.$$

Substituting this for the last term of the fifth equality yields the sixth equality. Finally, the eighth equality from the facts that  $y_N = 0$  and that the sets  $\bar{\mathcal{A}}_i$ ,  $i = 1, \dots, I$ , partition the basic activities. Since  $y = \sum_{j=1}^J y_j e^j$ , we conclude that  $y \in \text{span}(\bar{C})$ .

Conversely, let  $y \in \text{span}(\bar{C})$ . Then there are constants  $a_{jj'}$ ,  $(j, j') \in \bar{C}$ , such that

$$y = \sum_{(j, j') \in \bar{C}} a_{jj'} (e^j - e^{j'}).$$

Since  $\bar{C}$  consists only of pairs of basic activities, it follows that  $y_N = 0$ . Furthermore, for  $(j, j') \in \bar{C}$  and  $i \in \{1, \dots, I\}$ , we have

$$\left[ A(e^j - e^{j'}) \right]_i = \sum_{l=1}^b A_{il} (e_l^j - e_l^{j'}) = \sum_{l=1}^b \mathbf{1}_{\{s(l)=i\}} (e_l^j - e_l^{j'}) = \mathbf{1}_{\{s(j)=i\}} - \mathbf{1}_{\{s(j')=i\}} = 0,$$

the second equality holding by (2.1) and the fourth equality holding since  $s(j) = s(j')$ . Therefore,  $A(e^j - e^{j'}) = 0$  for all  $(j, j') \in \bar{C}$ , implying that  $Ay = 0$  by linearity. So,  $y \in \{y \in \mathbb{R}^J : Ay = 0, y_N = 0\}$ .

**Step 2:** In this step, we show that  $\mathcal{N} = \text{span}(\tilde{C})$ , where  $\tilde{C} = \{Ry : y \in \bar{C}\}$ . To see this, recall that  $\mathcal{N} = \{Hy_B : By_B = 0, y_B \in \mathbb{R}^b\} = \{Ry : Ay = 0, y_N = 0\}$ . Thus, it follows from Step 1 and the definition of  $\tilde{C}$  that  $\mathcal{N} = \text{span}(\tilde{C})$ .

**Step 3:** In this step, we show that  $\tilde{C} = \{\mu_j^* (e^{b(j)} - e^{b(j')}) : (j, j') \in \bar{C}, e^{b(i)}, e^{b(j')} \in \mathbb{R}^I\}$ . To see this, note that for  $(j, j') \in \bar{C}$  and  $i \in \{1, \dots, I\}$ , we have that

$$\left[ R(e^j - e^{j'}) \right]_i = \sum_{l=1}^J R_{il} (e_l^j - e_l^{j'}) = \sum_{l=1}^J \mu_l^* \mathbf{1}_{\{b(l)=i\}} (e_l^j - e_l^{j'}) = \mu_j^* \mathbf{1}_{\{b(j)=i\}} - \mu_{j'}^* \mathbf{1}_{\{b(j')=i\}}$$

$$= \mu_j^* \left( \mathbf{1}_{\{b(j)=i\}} - \mathbf{1}_{\{b(j')=i\}} \right) = \mu_j^* \left( e_i^{b(j)} - e_i^{b(j')} \right),$$

the second equality following from (2.2) and (2.25) and the fourth equality following from the fact that  $s(j) = s(j')$  (since  $(j, j') \in \bar{C}$ ) and (2.24). That is,

$$R(e^j - e^{j'}) = \mu_j^* \left( e^{b(j)} - e^{b(j')} \right) \quad \text{for } (j, j') \in \bar{C}. \quad (2.152)$$

Then using the definition of  $\bar{C}$ , we write

$$\begin{aligned} \tilde{C} &= \{Ry : y \in \bar{C}\} \\ &= \{Ry : y = e^j - e^{j'} \text{ such that } (j, j') \in \bar{C}, e^j, e^{j'} \text{ are unit basis vectors}\} \\ &= \{R(e^j - e^{j'}) : (j, j') \in \bar{C}, e^j, e^{j'} \text{ are unit basis vectors}\} \\ &= \{\mu_j^* (e^{b(j)} - e^{b(j')}) : (j, j') \in \bar{C}, e^j, e^{j'} \text{ are unit basis vectors}\}, \end{aligned}$$

where the last equality follows from (2.152). Hence, the result holds. In particular, by the definition of buffer communication, note that

$$\tilde{C} = \{\mu_j^* (e^i - e^{i'}) : \text{buffers } i \text{ and } i' \text{ communicate directly, } e^i, e^{i'} \in \mathbb{R}^I\}.$$

**Step 4:** We consider the matrix  $M$  defined in Lemma 6 (see (2.60)) and show that its rows form a basis for  $\mathcal{M}$ . To that end, let  $M^l$ ,  $l = 1, \dots, L$ , be the rows of the matrix  $M$  given in (2.60). Since the buffer pools partition the servers, the rows of  $M$  are linearly independent. Thus, to complete the proof, it suffices to show that  $\mathcal{M} = \text{span}(M^1, \dots, M^L)$ . Recalling that  $\mathcal{M} = \mathcal{N}^\perp$  and  $\mathcal{N} = \text{span}(\tilde{C})$ , it follows that  $a \in \mathcal{M}$  if and only if  $a \cdot z = 0$  for all  $z \in \tilde{C}$ . Moreover, since  $\mu_j^* > 0$  for all  $j \in \{1, \dots, b\}$ , it follows from Step 3 that

$$\mathcal{N} = \text{span}\left(\left\{e^i - e^{i'} : \text{buffers } i \text{ and } i' \text{ communicate directly, } e^i, e^{i'} \in \mathbb{R}^I\right\}\right).$$

Therefore,  $a \in \mathcal{M}$  if and only if  $a_i = a_{i'}$  for all buffers  $i$  and  $i'$  that communicate directly.

To prove that  $\mathcal{M} = \text{span}(M^1, \dots, M^L)$  we show inclusions of both sets. On the one hand, let  $a \in \mathcal{M}$ . Then  $a_i = a_{i'}$  for all buffers  $i$  and  $i'$  that communicate directly. By definition of buffer communication, it immediately follows that  $a_i = a_{i'}$  for all buffers  $i$  and  $i'$  that communicate. That is,  $a_i = a_{i'}$  for all buffers  $i$  and  $i'$  that are in the same buffer pool. Thus,  $a \in \text{span}(M^1, \dots, M^L)$ , implying that  $\mathcal{M} \subseteq \text{span}(M^1, \dots, M^L)$ . On the other hand, to show that  $\text{span}(M^1, \dots, M^L) \subseteq \mathcal{M}$ , it suffices to show that  $M^l \in \mathcal{M}$  for each  $l = 1, \dots, L$ . To that end, it is enough to show that  $M_i^l = M_{i'}^l$  for all buffers  $i$  and  $i'$  that communicate directly. However, this trivially holds by (2.60), since buffers  $i$  and  $i'$  that communicate directly are in the same buffer pool. Thus,  $\text{span}(M^1, \dots, M^L) \subseteq \mathcal{M}$ .  $\square$

*Proof of Lemma 7.* It is enough to show that  $(MR)_{lj} = (GA)_{lj}$  for all  $l = 1, \dots, L$  and  $j = 1, \dots, J$ , where  $G$  is given by (2.61). Indeed, by (2.2), (2.25), and (2.60),

$$(MR)_{lj} = \sum_{i=1}^I M_{li} R_{ij} = \sum_{i=1}^I \mathbf{1}_{\{i \in \mathcal{P}_l\}} \mu_j^* \mathbf{1}_{\{b(j)=i\}} = \mu_j^* \mathbf{1}_{\{b(j) \in \mathcal{P}_l\}}. \quad (2.153)$$

On the other hand, by (2.1) and (2.61),

$$(GA)_{lj} = \sum_{i=1}^I G_{li} A_{ij} = \sum_{i=1}^I \lambda_i^* \mathbf{1}_{\{i \in \mathcal{S}_l\}} \mathbf{1}_{\{s(j)=i\}} = \lambda_{s(j)}^* \mathbf{1}_{\{s(j) \in \mathcal{S}_l\}}. \quad (2.154)$$

Note that by (2.24) we have  $\mu_j^* = \lambda_{s(j)}^*$  and by (2.59) we have that  $b(j) \in \mathcal{P}_l$  if and only if  $s(j) \in \mathcal{S}_l$ . Thus, the desired result immediately follows by (2.153)–(2.154).  $\square$

*Proof of Lemma 8.* When  $L = 1$ , all buffers are in a single buffer pool. Thus, it follows immediately from (2.60) that  $M = e'$ . Furthermore, by definition of buffer communication and (2.59), there is a single server pool. It then follows from (2.61) that  $G = (\lambda^*)'$ .

To prove the first relationship in (2.70), note that

$$M\eta q = \eta Mq = \eta e' q = \eta \sum_{i=1}^I q_i = \eta,$$

where the second equality follows from  $M = e'$  and where the fourth equality follows from the fact that  $q$  is a probability vector. To prove the second relationship in (2.70), first note that  $MC = e' \in \mathbb{R}^J$ . This follows from  $M = e' \in \mathbb{R}^I$ , the definition of  $C$  in (2.2), and the fact that  $C$  has one nonzero element per column. Therefore,

$$MC \text{diag}(x^*) A' = e' \text{diag}(x^*) A' = (x^*)' A' = (Ax^*)' = e' \in \mathbb{R}^I,$$

the fourth equality following from the heavy traffic assumption (see (2.29)).  $\square$

*Proof of Lemma 9.* This is a straightforward convex optimization problem. Forming the Lagrangian

$$\mathcal{L}(\zeta, \nu) := \sum_{i=1}^I \alpha_i \zeta_i^2 - \nu \sum_{i=1}^I \zeta_i + \nu x,$$

where  $\nu$  is the Lagrange multiplier, the necessary first-order conditions then give

$$\zeta_i^* = \frac{\gamma}{2\alpha_i}, \quad i = 1, \dots, I.$$

Substituting this into the constant  $e' \zeta = x$  yields  $\nu = 2x/\hat{\alpha}$  and

$$\zeta_i^* = \frac{x}{\alpha_i \hat{\alpha}}, \quad i = 1, \dots, I. \tag{2.155}$$

The optimality of this solution follows from the convexity of the objective. Substituting (2.155) in the objective function yields  $c(x) = x^2/\hat{\alpha}$  as desired.  $\square$

*Proof of Proposition 3.* Let  $(Y, \zeta)$  be an admissible control for (2.48)–(2.55) with the corresponding state process  $Z$  and idleness process  $U$ . Letting  $W(t) = MZ(t)$  for  $t \geq 0$ , (2.48)

implies that (2.66) holds, and (2.65) holds by definition. Similarly, (2.67)–(2.68) follow from (2.54)–(2.55) whereas (2.69) follows from (2.53). Thus,  $(Z, U, \zeta)$  of the BCP formulation (2.48)–(2.55) is an admissible policy for the RBCP (2.64)–(2.69). Because the two formulations have the same process  $Z, U, \zeta$ , they have the same cost.

The converse follows exactly as in (the second part of) the proof of Theorem 1 in Harrison and Miqhem [1997] (see pages 753–754 of that paper) with the only substantive difference being (aside from the obvious notational differences) the process  $X$  on their equation (36) on page 755 is replaced with

$$B(t) - \eta q \int_0^t e' Z(s) ds - C \text{diag}(x^*) \int_0^t \kappa(s) ds$$

in our setting. Then following the same steps in their proof shows that the analogy of the process  $Y$  (in our setting) defined as in their equation (35) and  $\zeta$  is admissible for our BCP (2.48)–(2.55). Moreover, because  $(Y, \zeta)$  results in the same queue length process  $Z$ . Its cost is the same as that of the policy  $(Z, U, \zeta)$  for RBCP (2.64)–(2.69).  $\square$

*Proof of Proposition 4.* Given an admissible policy  $\theta$  for EWF and the corresponding processes  $W$  and  $L$ , we set  $Z_{i^*} = W$  and  $Z_i \equiv 0$  for  $i \neq i^*$ , as well as  $U_{k^*} = L$  and  $U_k \equiv 0$  for  $k \neq k^*$ . Moreover, we set  $\zeta_i(s) = \theta(s)/(\alpha_i \hat{\alpha}_i)$  for  $i = 1, \dots, I$  and  $s \geq 0$ , which results in  $\sum_{i=1}^I \alpha_i \zeta_i^2(s) = c(\theta(s))$  for  $s \geq 0$  by Lemma 9. Then it follows from (2.78)–(2.79) that  $(Z, U, \zeta)$  has the same cost in the RBCP as  $\theta$  does in the EWF.

To prove the converse, let  $(Z, U, \zeta)$  be an admissible policy for RBCP, and let

$$\theta(s) = e' \zeta(s), \quad W(s) = e' Z(s), \quad \text{and} \quad L(s) = (\lambda^*)' U(s), \quad s \geq 0.$$

It is easy to verify that  $\theta$  is admissible for EWF. Moreover,  $c(\theta(s)) \leq \sum_{i=1}^I \alpha_i \zeta_i^2(s)$  for all  $s \geq 0$  by Lemma 9, and, furthermore, we have  $hW(s) \leq \sum_{i=1}^I (h_i - h_0) Z_i(s)$  and  $rL(s) \leq c'U(s)$  for  $s \geq 0$ . Thus, the cost of  $\theta$  for the EWF is less than or equal to the cost of the policy  $(Z, U, \zeta)$



for the RBCP. □

*Proof of Proposition 5.* Consider the auxiliary stationary reflected diffusion on  $[0, \infty)$ , denoted by  $\{\tilde{W}(t), t \geq 0\}$ , associated with the drift rate function  $-(\eta y - a + \theta^*(y))$  and variance parameter  $\sigma^2$ . As noted on pages 470–471 of Browne and Whitt [1995] (also see Mandl [1968] and Karlin and Taylor [1981]) its probability density function  $\varphi$  is given as follows:

$$\varphi(x) = \frac{\exp\left\{-\int_0^x \frac{2}{\sigma^2}(\eta y - a + \theta^*(y)) dy\right\}}{\int_0^\infty \exp\left\{-\int_0^y \frac{2}{\sigma^2}(\eta s - a + \theta^*(s)) ds\right\} dy}, \quad x \in [0, \infty) \quad (2.156)$$

provided all integrals are finite, which we verify next. To this end, let  $k = \inf\{y \geq 0 : v(y) \geq 0\}$  where  $(v, \beta^*)$  solve the Bellman equation (2.90)–(2.91), and note from (2.91) that  $-r \leq v(y) \leq 0$  for  $y \leq k$  and  $0 \leq v(y) \leq h/\eta$  for  $y \geq k$ . In order to verify the integrals above are finite, using (2.92) note that

$$\begin{aligned} \exp\left\{-\int_0^y \frac{2}{\sigma^2}(\eta s - a + \theta^*(s)) ds\right\} &= \exp\left\{-\int_0^y \frac{2}{\sigma^2}\left(\eta s - a + \frac{\hat{\alpha}}{2}v(s)\right) ds\right\} \\ &= \exp\left\{-\frac{\eta y^2 - ay}{\sigma^2}\right\} \exp\left\{-\frac{\hat{\alpha}}{\sigma^2}\left[\int_0^k v(s) ds + \int_k^y v(s) dy\right]\right\} \\ &\leq \exp\left\{-\frac{\eta y^2 - ay}{\sigma^2}\right\} \exp\left\{\frac{\hat{\alpha}}{\sigma^2}rk\right\}, \end{aligned} \quad (2.157)$$

from which we also deduce that the integral in the denominator of the right-hand side of (2.156) is finite. Moreover, it follows from (2.157) that the stationary diffusion  $\tilde{W}$  has finite moments. In particular,

$$E[\tilde{W}(0)] = E[\tilde{W}(t)] < \infty, \quad t < \infty. \quad (2.158)$$

Next, we define another auxiliary stationary diffusion  $\widetilde{W}^*$  as follows:

$$\widetilde{W}^*(t) := W^*(0) + \widetilde{W}(t), \quad t \geq 0.$$

Noting  $W^*(0) < \widetilde{W}^*(0)$  almost surely, we define the stopping time  $\tau$  as follows:

$$\tau := \inf \{t \geq 0 : W^*(t) \geq \widetilde{W}^*(t)\}$$

and introduce the following process:

$$\widehat{W}^*(t) := \begin{cases} W^*(t), & t < \tau, \\ \widetilde{W}^*(t), & t > \tau. \end{cases}$$

By the strong Markov property of diffusions,  $\widehat{W}^*$  has the same distribution as  $W^*$ . Moreover,  $\widehat{W}^*(t) \leq \widetilde{W}^*(t)$  for all  $t \geq 0$ . Therefore, we conclude that

$$\begin{aligned} E[W^*(t)] &= E[\widehat{W}^*(t)] \leq E[\widetilde{W}^*(t)] = W^*(0) + E[\widetilde{W}(t)] \\ &= W^*(0) + E[\widetilde{W}(0)] < \infty, \end{aligned} \tag{2.159}$$

where the second equality follows from the definition of  $\widehat{W}^*$ , the third equality from the stationarity of  $\widetilde{W}$ , and the last equality from (2.158). Thus, we conclude from  $W^*(t) \geq 0$  for  $t \geq 0$  and (2.159) that

$$\lim_{t \rightarrow \infty} \frac{E[W^*(t)]}{t} \leq \lim_{t \rightarrow \infty} \frac{W^*(0) + E[\widetilde{W}(0)]}{t} = 0,$$

which completes the proof. □

The next lemma aids in the proof of Lemma 11. To state the result, it will be convenient

to rewrite (2.103)–(2.104) as follows:

$$v'(x) = q_2 v^2(x) + q_1(x)v(x) + q_0(x), \quad x \geq 0, \quad (2.160)$$

$$v(0) = -r, \quad (2.161)$$

where  $q_0(x) = \frac{2}{\sigma^2}(\beta - hx)$ ,  $q_1(x) = \frac{2}{\sigma^2}(\eta x - a)$ , and  $q_2 = \frac{\hat{\sigma}}{2\sigma^2} > 0$  for  $x \geq 0$ .

**Lemma 28.** *For each  $v \in C^1[0, \infty)$  satisfying (2.160)–(2.161),  $y(x) := \exp\{-q_2 \int_0^x v(t) dt\}$  for  $x \geq 0$  satisfies*

$$y''(x) - q_1(x)y'(x) + q_2q_0(x)y(x) = 0, \quad x \geq 0, \quad (2.162)$$

$$y(0) = 1, \quad y'(0) = rq_2. \quad (2.163)$$

*Conversely, for each  $y \in C^2[0, \infty)$  satisfying (2.162)–(2.163),  $v(x) := -y'(x)/(q_2y(x))$  for  $x \geq 0$  satisfies (2.160)–(2.161).*

*Proof of Lemma 28.* Suppose that  $v \in C^1[0, \infty)$  satisfies (2.160)–(2.161) and let  $y(x) = \exp\{-q_2 \int_0^x v(t) dt\}$  for  $x \geq 0$ . Then it follows that

$$\begin{aligned} & y''(x) - q_1(x)y'(x) + q_2q_0(x)y(x) \\ &= \left[ \exp\left\{-q_2 \int_0^x v(t) dt\right\} \cdot (-q_2v(x)) \right]' - q_1(x) \left[ \exp\left\{-q_2 \int_0^x v(t) dt\right\} \cdot (-q_2v(x)) \right] \\ & \quad + q_2q_0(x) \exp\left\{-q_2 \int_0^x v(t) dt\right\} \\ &= \left[ \exp\left\{-q_2 \int_0^x v(t) dt\right\} \cdot (-q_2v(x))^2 + \exp\left\{-q_2 \int_0^x v(t) dt\right\} \cdot (-q_2v'(x)) \right] \\ & \quad + q_2q_1(x)v(x) \exp\left\{-q_2 \int_0^x v(t) dt\right\} + q_2q_0(x) \exp\left\{-q_2 \int_0^x v(t) dt\right\} \\ &= q_2 \exp\left\{-q_2 \int_0^x v(t) dt\right\} [q_2v^2(x) - v'(x) + q_1(x)v(x) + q_0(x)] = 0. \end{aligned}$$

Moreover,  $y(0) = \exp\{-q_2 \cdot 0\} = 1$  and  $y'(0) = -q_2 \exp\{-q_2 \cdot 0\} v(0) = rq_2$ .

On the other hand, suppose  $y \in C^2[0, \infty)$  satisfies (2.162)–(2.163) and let  $v(x) = -y'(x)/(q_2y(x))$  for  $x \geq 0$ . Then it follows that

$$\begin{aligned} v'(x) &= \left[ -\frac{y'(x)}{q_2y(x)} \right]' = -\frac{1}{q_2} \left[ \frac{y''(x)}{y(x)} - \left( \frac{y'(x)}{y(x)} \right)^2 \right] = -\frac{y''(x)}{q_2y(x)} + q_2v^2(x) \\ &= -\frac{q_1(x)y'(x)}{q_2y(x)} + q_0(x) + q_2v^2(x) = q_2v^2(x) + q_1(x)v(x) + q_0(x). \end{aligned}$$

Moreover,  $v(0) = -y'(0)/(q_2y(0)) = -(rq_2)/(q_2 \cdot 1) = -r$ . This completes the proof.  $\square$

*Proof of Lemma 11.* It is known that (2.162)–(2.163) can be transformed into a degenerate hypergeometric equation known as a Kummer's equation; see Polyanin and Zaitsev [2003]. Such equations are known to have confluent hypergeometric function solutions; see Bateman and Erdélyi [1953] and Abramowitz and Stegun [2003]. It then follows from Lemma 28 that (2.160)–(2.161) have a solution  $v$ . To complete the proof, we must show that the solution  $v$  to (2.160)–(2.161) is unique. To this end, define the function  $f$  by

$$f(x, v) = q_2v^2 + q_1(x)v + q_0(x), \quad (x, v) \in [0, \infty) \times (-\infty, \infty).$$

We show that  $f$  is locally Lipschitz in  $v$ , i.e., that  $f$  is Lipschitz in  $v$  when restricted to the compact domain  $[0, N] \times [-M, M]$  where  $N, M > 0$ . More specifically, local Lipschitzness will demonstrate uniqueness on each compact interval, which can then be easily extended to the positive real line. To this end, for  $x \in [0, N]$  and  $v_1, v_2 \in [-M, M]$  we have that

$$\begin{aligned} |f(x, v_1) - f(x, v_2)| &= |q_2v_1^2 + q_1(x)v_1 - q_2v_2^2 - q_1(x)v_2| \leq q_2|v_1^2 - v_2^2| + |q_1(x)| \cdot |v_1 - v_2| \\ &= \left[ q_2|v_1 + v_2| + \frac{2}{\sigma^2}(\eta x + |a|) \right] \cdot |v_1 - v_2| \\ &\leq \left[ 2Mq_2 + \frac{2}{\sigma^2}(\eta N + |a|) \right] |v_1 - v_2|. \end{aligned}$$

Therefore,  $f$  is locally Lipschitz in  $v$ , which completes the proof.  $\square$

# CHAPTER 3

## A STOCHASTIC MODELING FRAMEWORK FOR ONLINE CONTENT MODERATION: MANAGING CONTENT VIA MACHINE LEARNING AND HUMAN LABELING

### 3.1 Introduction

Over the last decade, the use of online platforms, particularly social media, has continued to expand rapidly. Recent surveys by the Pew Research Center reveal that a significant portion of American adults and teenagers use social media; see Auxier and Anderson [2021] and Vogels et al. [2022]. The surge in the use of online platforms has consequently led to a dramatic increase in user-generated content (UGC), which includes text, images, videos, and audio. (Throughout the chapter, we use “content” and “UGC” interchangeably.) While the proliferation of UGC has enhanced our capability to share and exchange information, it has also introduced significant challenges in content moderation.

Online content moderation refers to the process of analyzing UGC to detect and remove content that violates platform policies. This task is particularly crucial for social media platforms such as Facebook and Instagram, where harmful content such as hate speech, misinformation, and explicit material poses risks to its users. The emergence of such content has compelled social media companies to deploy robust content moderation systems. One of the major challenges faced in online content moderation is the sheer volume of uploaded content that must be managed. For instance, as noted in Makhijani et al. [2021], large social media platforms like Facebook handle billions of posts per day, presenting significant operational challenges.

To address these challenges, many online platforms employ a combination of artificial intelligence (AI) and human moderators to moderate content. AI-enabled systems leverage machine learning and natural language processing algorithms to distinguish between harmful

and benign content. These systems excel at quickly processing large amounts of data, allowing for rapid pre-moderation of content before it goes live on the platform. However, they often struggle with nuanced and context-dependent decisions, particularly when interpreting the complexities of human language and cultural references. This can lead to false positives and false negatives in content classification.

Conversely, human moderators provide expertise and contextual insight, allowing for more accurate decision-making for content that challenges AI systems. For example, as noted by Makhijani et al. [2021], “Facebook still relies on human reviewers for moderating ‘ambiguous content’ that is harder for computers to classify reliably.” Nonetheless, human moderation has its limitations—it is inherently slower than AI moderation and can introduce human biases or inconsistencies. Additionally, human moderators risk exposure to disturbing or graphic content. Indeed, there is a growing body of evidence that human content moderators are at a considerable risk for significant psychological damage; see Chen [2017, 2014] and Ruberts [2017].

This paper introduces a stochastic modeling framework for online content moderation that integrates AI and human moderators to effectively manage UGC. Initially, content is pre-moderated by the AI system, which can then be sent to human review if necessary. The platform may choose to send multiple copies of the same piece of content to different reviewers, serving the following dual purposes. First, it reduces the expected misclassification cost associated with that content. Second, it produces more human-generated labels to enhance the AI system’s accuracy over time. Indeed, according to Nguyen et al. [2020], human labeling errors can be mitigated by multiple reviews, albeit at a cost of increased system time. Through a simulation study, we aim to explore the operational implications that arise from sending multiple copies of content to human reviewers. To the best of our knowledge, these issues have not yet been addressed in the existing literature.

## 3.2 Literature Review

The first stream of literature related to our work focuses on online content moderation. Two notable papers from Facebook’s Core Data Science team, Makhijani et al. [2021] and Nguyen et al. [2020], address various challenges encountered in analyzing large-scale online content moderation systems. First, Makhijani et al. [2021] describes QUEST (Queue Simulation for Content Moderation at Scale), which is a simulation tool employed at Facebook to address operational challenges in large-scale human content moderation systems. It is particularly valuable for testing the impact of potential decisions, like increasing staff or enhancing their skills. The paper details how the simulator supports decision-making in areas such as capacity planning and queue prioritization. Second, Nguyen et al. [2020] presents CLARA (Confidence of Labels and Raters), which is a statistical framework developed and deployed at Facebook that aggregates decisions from content reviewers and quantifies their uncertainty, recognizing the variability and potential biases of human judgment. A final relevant paper on online content moderation is Cen et al. [2023], which provides a framework for regulating and auditing a social media platform according to a “baseline feed” that represents content a user would see without algorithmic manipulation. The paper proposes a method to ensure that algorithmically curated feeds do not deviate significantly from these baselines, thereby enhancing user agency and regulatory compliance. Further discussions on the role of AI in content moderation are explored in the works of Cambridge Consultants [2019], Gillespie [2020], and Udupa et al. [2021]. These papers investigate the evolving capabilities of AI technologies in effectively moderating online content, highlighting both current applications and future prospects for AI-driven moderation systems.

The second stream of literature relevant to our work focuses on queueing with redundancy, as our model incorporates the sending of multiple copies of content for human review. Nageswaran and Scheller-Wolf [2022] examine the dynamics within service systems where

“redundant” customers join multiple queues to enhance their chances of receiving service, unlike others who join a single queue. The study discusses the implications of this system on fairness, particularly in the context of organ transplantation. Their findings highlight situations where redundancy improves system efficiency without compromising fairness, as well as conditions under which it may disadvantage single-queue customers. Similarly, Gardner et al. [2017] tackles server-side variability in queueing systems by implementing redundancy. In these systems, jobs replicate themselves and distribute these copies across multiple servers, with the completion of any single copy marking the end of the process. This approach strategically uses scheduling to balance low response times with fairness among job classes. Queueing with redundancy is also studied in various other papers, including Anton et al. [2022, 2024], Gardner [2022], and Guo and Hassin [2015].

A third stream of literature related to our work pertains to multiple listings in organ transplantation. Specifically, as waiting lists for organ transplants have grown, innovative strategies have been explored to expedite access to transplants. A notable approach is multiple listings, where patients can register at more than one transplant center to improve their chances of receiving an organ. This practice has been discussed and examined in studies such as those by Merion et al. [2004], Ashlagi and Roth [2012], Ardekani and Orlowski [2010], and Vagefi et al. [2014]. One particularly relevant paper from the operations management literature is Ata et al. [2017]. This paper addresses the stark disparities in waiting times for kidney transplants across different US regions. They propose an innovative solution using an affordable jet service called OrganJet, enabling them to register across multiple, geographically distant donation service areas. They model the decision-making process of patients as a selfish routing game where each patient aims to minimize their congestion costs. Their findings demonstrate that multiple listings facilitated by OrganJet could significantly enhance geographic equity in organ transplantation access.

Finally, our work is related to research on the interplay between learning and system



congestion, highlighted by two key papers: Alizamir et al. [2013] and Massoulié and Xu [2018]. Alizamir et al. [2013] explores the trade-offs between diagnostic accuracy and system congestion. This paper examines the impact of conducting additional diagnostic tests, which while improving diagnostic accuracy, also increase waiting times for other customers. They formulate this problem of balancing diagnostic accuracy against delays as a partially observed Markov decision process and characterize the optimal decision rule that maximizes long-term value by weighing correct diagnoses against the costs of misidentification and delays. Meanwhile, Massoulié and Xu [2018] considers modern processing systems designed for the purpose of learning and information extraction. In these systems, a group of experts (e.g., human agents) are used to perform inspections of a large number of jobs for the purpose of uncovering hidden features associated with each job up to a level of desirable accuracy. However, these experts possess a finite amount of service capacity. As such, this paper studies the minimum number of experts needed in order to maintain stability and learn a sufficient amount of information about every job in a stream of arrivals.

### 3.3 Stochastic Model for Online Content Moderation

This section introduces a stochastic model for an online content moderation system, as illustrated in Figure 3.1. The process begins with an ML classifier assessing UGC, assigning it a binary score that categorizes the content as either harmful or benign. Subsequently, the platform determines whether the content warrants further review by human moderators. If the decision is to bypass human review, the content’s fate—posting or removal—is determined partially based on the score assigned by the ML classifier. Conversely, if the content is flagged for human review, multiple copies of the content may be sent to various human reviewers for further assessment. These human moderators then generate labels that serve dual purposes: they inform the final decision on the content’s fate and enhance the ML classifier’s accuracy over time through feedback. The mathematical details of this model are discussed further in

the subsequent sections.

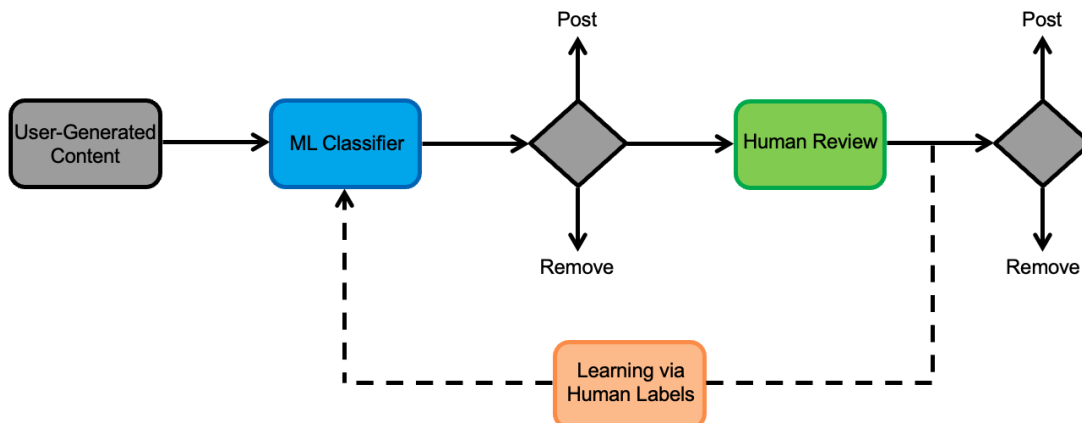


Figure 3.1: A process flow diagram of an online content moderation system featuring a machine learning classifier and human moderators. The dashed line symbolizes the feedback mechanism where human-generated labels refine the machine learning classifier’s accuracy over time via supervised learning.

### 3.3.1 User-Generated Content

Each piece of user-generated content is characterized by an associated class  $x \in \mathcal{X}$  and a virality level  $y \in \mathcal{Y}$ . The components are critical to the model and are further elaborated below:

- **Content Class:** The class of a piece of content encompasses tangible characteristics, including text, images, videos, and metadata, that can in principle be harmful. A content’s class is a key determiner of its underlying violation status. The set of all content classes is denoted by  $\mathcal{X}$ .
- **Content Virality:** Virality quantifies the anticipated reach of content, essentially predicting the number of users likely to be exposed. Although virality itself does not determine the content’s compliance with platform standards, incorrectly assessing highly viral content can result in greater consequences for the platform and its user base. The set of virality levels is denoted by  $\mathcal{Y}$ .

Next, we discuss the probabilistic elements of our model, and we assume throughout that all random quantities are defined on a fixed probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ . The inherent randomness of a piece of content arises from two primary sources: its arrival time and its underlying violation status. We discuss both sources below.

Content is continuously generated by users over time. We model this process by a renewal process, where for each pair  $(x, y) \in \mathcal{X} \times \mathcal{Y}$ , the arrivals of class  $x$  content with virality  $y$  follow a homogeneous Poisson process  $M(x, y) = (M_t(x, y) : t \geq 0)$  with a rate of  $\Lambda(x, y) > 0$  per unit time. Therefore, the arrival time of the  $k$ th piece of class  $x$  content with virality  $y$  is defined as

$$\tau_k(x, y) := \inf \{t \geq 0 : M_t(x, y) = k\}, \quad k = 1, 2, \dots \quad (3.1)$$

The violation status of a piece of content is influenced by its class and specific idiosyncratic features that differentiate it from other members of its class. The platform does not have direct knowledge of a content's true violation status; such knowledge would enable flawless decision-making regarding content posting or removal. However, we assume that the platform has access to the distribution of content within each class that violates the posting standards. To formalize this, for each  $x \in \mathcal{X}$ , we introduce independent sequences of i.i.d. Bernoulli random variables  $\{\theta_k(x), k \geq 1\}$  and  $\{\epsilon_k(x), k \geq 1\}$  with parameters

$$\bar{\theta}(x) := \mathbb{P}(\theta_1(x) = 1) \quad \text{and} \quad \bar{\epsilon}(x) := \mathbb{P}(\epsilon_1(x) = 1), \quad (3.2)$$

respectively. Here,  $\theta_k(x)$  signifies the class-level violation status for class  $x$  content, with  $\theta_k(x) = 1$  indicating compliance with the posting standards, and  $\theta_k(x) = 0$  indicating noncompliance. Similarly,  $\epsilon_k(x)$  signifies the violation status of the idiosyncratic features for class  $x$  content. Content is considered compliant with the platform's posting standards only if both its class-level violation status and idiosyncratic features meet these standards. Hence,

the overall violation status of the  $k$ th piece of class  $x$  content is captured by the following Bernoulli random variable:

$$\theta_k(x, \epsilon) := \theta_k(x) \epsilon_k(x), \quad k = 1, 2, \dots \quad (3.3)$$

In the proceeding discussion, we will often just refer to an individual piece of content. Thus, for ease of discussion and notational simplicity, we will often write  $\theta(x, \epsilon)$ ,  $\theta(x)$  and  $\epsilon(x)$  instead of  $\theta_k(x, \epsilon)$ ,  $\theta_k(x)$  and  $\epsilon_k(x)$ , respectively.

### 3.3.2 Content Classifiers

As previously mentioned, social media companies utilize both ML classifiers and human moderators to manage UGC. Each component in the moderation system evaluates content, assigning scores to determine if it is harmful or benign. The functioning of these classifiers is discussed in further detail below.

#### Machine Learning Classifier

Upon arrival at time  $t$ , class  $x$  content is evaluated by the ML classifier, which assigns a binary score  $\theta_t^{\text{ML}}(x) \in \{0, 1\}$ , where  $\theta_t^{\text{ML}}(x) = 1$  signifies the classifier's assessment of the content as benign, and  $\theta_t^{\text{ML}}(x) = 0$  otherwise. For each class  $x \in \mathcal{X}$ , the accuracy of the ML classifier at time  $t$  is represented by a time-dependent confusion matrix as follows:

$$m_{ii}^{\text{ML}}(x, t) := \mathbb{P}(\theta_t^{\text{ML}}(x) = i | \theta(x) = i, \mathcal{F}_{t-}), \quad i = 0, 1, \quad (3.4)$$

where  $\mathcal{F}_{t-}$  denotes the information available to the platform up to time  $t$ . As the platform collects human-generated labels (see Section 3.3.2), the accuracy of the ML classifier improves over time. This motivates the following assumption.

**Assumption 8.** For  $x \in \mathcal{X}$ , there exist constants  $\alpha(x) \in [0, 1]$  and  $\beta(x) \in (0, \infty)$  such that

$$m_{ii}^{\text{ML}}(x, t) = 1 - \alpha(x) \exp\{-\beta(x) N_{t-}(x)\}, \quad i = 0, 1, \quad (3.5)$$

where  $N_{t-}(x) \in \mathbb{Z}_{\geq 0}$  the count of human-assigned labels for class  $x$  content up to time  $t$ .

This assumption implies that the accumulation of labels improves the accuracy of the ML classifier, i.e.,  $m_{ii}^{\text{ML}}(x, t) \rightarrow 1$  as  $N_{t-}(x) \rightarrow \infty$  for  $i = 0, 1$ . The assumption of nondecreasing accuracy, captured by (3.5), is a deliberate modeling simplification that implicitly assumes that all labels are informative. While the actual classifier accuracy may exhibit variability due to the stochasticity of noisy labels (see, e.g., Frénay and Verleysen [2014]), this simplification streamlines the analysis without compromising the fundamental relationship between classifier accuracy and the number of labels. Furthermore, the analysis presented below is robust to the precise functional form of  $m_{ii}^{\text{ML}}(x, t)$ .

Although the ML classifier's ability to predict class-level violation statuses improves incrementally with each label, the classifier's true accuracy is limited by the presence of the idiosyncratic features. The following lemma makes this statement precise.

**Lemma 29.** *For each  $x \in \mathcal{X}$ , we have that*

$$\mathbb{P}(\theta_t^{\text{ML}}(x) = \theta(x, \epsilon) | \mathcal{F}_{t-}) \leq 1 - \bar{\theta}(x) (1 - \bar{\epsilon}(x)), \quad t \geq 0.$$

*Proof.* By the law of total probability, we have that  $\mathbb{P}(\theta_t^{\text{ML}}(x) = \theta(x, \epsilon) | \mathcal{F}_{t-}) = \mathbb{P}(\theta_t^{\text{ML}}(x) = \theta(x, \epsilon) = 1 | \mathcal{F}_{t-}) + \mathbb{P}(\theta_t^{\text{ML}}(x) = \theta(x, \epsilon) = 0 | \mathcal{F}_{t-})$  for each  $t \geq 0$ . The result then follows by applying Bayes' rule to each term and using the independence of  $\epsilon(x)$ .  $\square$

## Human Moderation System

Content not immediately posted or removed undergoes human moderation. This subsystem consists of a single queue where content (jobs) await review and a pool of  $N$  homogeneous agents who manually review content. The service times are modeled as i.i.d. random variables  $\{\nu_k, k \geq 1\}$  with  $\mathbb{E}[\nu_1] \in (0, \infty)$ , and are assumed to be independent of all other stochastic primitives. We restrict to non-idling, first-come-first-serve (FCFS) service disciplines.

Upon evaluation, a human moderator assigns a binary label  $\theta^H(x) \in \{0, 1\}$  to class  $x$  content, with the same interpretation as before. Unlike the ML classifier, human moderators are presumed capable of discerning both the class-level and specific idiosyncratic attributes of the content. To that end, for each  $x \in \mathcal{X}$ , the accuracy of human-assigned labels is represented by a confusion matrix as follows:

$$m_{ij}^H(x) := \mathbb{P}(\theta^H(x) = 1 \mid \theta(x) = i, \epsilon = j), \quad i, j = 0, 1. \quad (3.6)$$

We introduce an assumption that says the human-generated labels are informative.

**Assumption 9.** For each  $x \in \mathcal{X}$ , it holds that  $m_{11}^H(x) > 0.5$  and  $m_{ij}^H(x) < 0.5$  for all  $(i, j) \neq (1, 1)$ .

The platform has the option to send multiple copies of a single piece of content to different human reviewers. However, given the sheer volume of content that requires moderation, it is impractical for human reviewers to examine every piece of content. As such, the platform must prioritize the use of its human reviewers, balancing the expected misclassification costs associated with different types of content against potential delays incurred as content awaits a posting decision. This operational challenge motivates the following definition.

**Definition 5.** A *labeling policy*  $\pi = (n(x, y) : (x, y) \in \mathcal{X} \times \mathcal{Y})$  is a non-anticipating stochastic process, where for each pair  $(x, y) \in \mathcal{X} \times \mathcal{Y}$ , the process  $n(x, y) = \{n_t(x, y), t \geq 0\}$  is such that  $n_t(x, y) \in \mathbb{Z}_{\geq 0}$  for all  $t \geq 0$ .

A labeling policy thus determines the number of human reviews for class  $x$  content with virality  $y$  at any given time  $t$ , including the option  $n_t(x, y) = 0$ , in which case the content's fate is solely based on the ML classifier's assessment. The labeling policy impacts the queue length within the human moderation system. When service times are exponentially distributed, it is straightforward to analytically describe this process due to the memoryless property. In this case, given a labeling policy  $\pi$ , the queue length process  $Q^\pi = \{Q^\pi(t), t \geq 0\}$  for the human moderation system is defined recursively as follows:

$$Q^\pi(t) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \int_0^t n_s(x, y) dM_s(x, y) - N\left(\mu \int_0^t \min\{N, Q^\pi(s)\} ds\right), \quad t \geq 0, \quad (3.7)$$

where  $N = \{N(t), t \geq 0\}$  is a unit-rate Poisson process and  $\mu := (\mathbb{E}[v_1])^{-1}$ . However, the analytical expression of the queue length process becomes significantly more complex with generally distributed (i.e., non-exponential) service times due to the necessity of dynamically tracking service start times, departure times, and server assignments over time.

## Posting and Removal Decisions

The decision to post or remove content inevitably leads to misclassification errors. Type I errors, or false positives, occur when the system incorrectly removes benign content. Type II errors, or false negatives, occur when the system erroneously posts harmful content. For each content class  $x$  and virality  $y$ , we denote the costs associated with type I and type II errors by  $c^I(x, y) \geq 0$  and  $c^{II}(x, y) \geq 0$ , respectively. The expected cost of these errors, conditional on a labeling decision  $\theta \in \{0, 1\}$  at time  $t$ , is denoted by  $C_t(x, y, n_t(x, y), \theta)$ . We next discuss the formal definition of this cost function and how the platform makes the final posting and removal decisions.

Consider class  $x$  content with virality  $y$  that arrives at time  $t$ . When the platform does not send the content for human review, i.e.,  $n_t(x, y) = 0$ , the ML classifier immediately decides

on posting or removal. This decision is based on the classifier's current accuracy, the type I and type II error costs, and the score assigned to the content. Specifically, if  $\theta_t^{\text{ML}}(x) = 0$ , then the system's accuracy is contingent on  $\theta(x, \epsilon) = 0$ , with inaccuracies arising if  $\theta(x, \epsilon) = 1$ . Thus, adhering to the classifier's suggestion to remove the content leads to a type I error with probability  $\mathbb{P}(\theta(x, \epsilon) = 1 | \theta_t^{\text{ML}}(x) = 0)$ . Conversely, choosing to post the content against the classifier's advice results in a type II error with probability  $\mathbb{P}(\theta(x, \epsilon) = 0 | \theta_t^{\text{ML}}(x) = 0)$ . A similar reasoning applies for  $\theta_t^{\text{ML}}(x) = 1$ . That is, assuming  $\theta_t^{\text{ML}}(x) = 1$ , then the system is inaccurate if  $\theta(x, \epsilon) = 0$ . Thus, adhering to the classifier's suggestion to post the content leads to the type II error with probability  $\mathbb{P}(\theta(x, \epsilon) = 0 | \theta_t^{\text{ML}}(x) = 1)$ . On the other hand, choosing to remove the content against the classifier's advice leads to a type I error with probability  $\mathbb{P}(\theta(x, \epsilon) = 1 | \theta_t^{\text{ML}}(x) = 1)$ . Assuming that the platform makes the posting or removal decision that results in the lowest expected cost, it follows that

$$C_t(x, y, 0, \theta_t^{\text{ML}}(x)) := \begin{cases} \min\{\bar{m}_{10}^{\text{ML}}(x, t) c^{\text{I}}(x, y), \bar{m}_{00}^{\text{ML}}(x, t) c^{\text{II}}(x, y)\}, & \theta_t^{\text{ML}}(x) = 0, \\ \min\{\bar{m}_{11}^{\text{ML}}(x, t) c^{\text{I}}(x, y), \bar{m}_{01}^{\text{ML}}(x, t) c^{\text{II}}(x, y)\}, & \theta_t^{\text{ML}}(x) = 1, \end{cases} \quad (3.8)$$

where  $\bar{m}_{ij}^{\text{ML}}(x, t) := \mathbb{P}(\theta(x, \epsilon) = i | \theta_t^{\text{ML}}(x) = j)$  for  $i, j = 0, 1$ .

When the platform decides to send the content for human review, a crucial decision is to determine the number of copies  $n_t(x, y) > 0$  to distribute among the reviewers. This step is important because the aggregated score derived from these reviews significantly impacts the final decision on whether to post or remove the content. (In particular, this choice involves a trade-off between enhancing the ML classifier's accuracy through additional labels and the posting delays that occur due to the queue of content waiting to be reviewed.) Specifically, we denote by  $\theta^{\text{H}}(x, n) \in \{0, 1\}$  the aggregate score assigned to class  $x$  content when reviewed by  $n$  human moderators. The confusion matrix associated with this aggregate score is given by  $m_{ij}^{\text{H}}(x, n) := \mathbb{P}(\theta^{\text{H}}(x, n) = 1 | \theta(x) = i, \epsilon(x) = j)$ , where  $i, j = 0, 1$ . As before, assuming that the platform makes the posting or removal decision that results in the lowest expected cost,



it follows that

$$\begin{aligned}
& C_t(x, y, n_t(x, y), \theta^H(x, n_t(x, y))) \\
& := \begin{cases} \min\{\bar{m}_{10}^H(x, n_t(x, y)) c^I(x, y), \bar{m}_{00}^H(x, n_t(x, y)) c^H(x, y)\}, & \theta^H(x, n_t(x, y)) = 0, \\ \min\{\bar{m}_{11}^H(x, n_t(x, y)) c^I(x, y), \bar{m}_{01}^H(x, n_t(x, y)) c^H(x, y)\}, & \theta^H(x, n_t(x, y)) = 1, \end{cases} \quad (3.9)
\end{aligned}$$

where  $\bar{m}_{ij}^H(x, n_t(x, y)) := \mathbb{P}(\theta(x, \epsilon) = i \mid \theta^H(x, n_t(x, y)) = j)$  for  $i, j = 0, 1$ . Aggregating human-assigned labels into a singular decision is an important question. For example, Nguyen et al. [2020] proposes a more sophisticated method called CLARA to aggregate human reviewer decisions, but there are other known methods within the crowdsourcing literature. However, our primary focus lies in understanding the operational effects of various labeling policies. Therefore, for simplicity and to maintain our focus on the operational issues, we assume label aggregation via majority vote. Specifically, majority voting is given by letting

$$\theta^H(x, n) = \mathbb{1}\left\{n^{-1} \sum_{k=1}^n \theta_k^H(x) \geq \frac{1}{2}\right\}, \quad (x, n) \in \mathcal{X} \times \{0, 1, 2, \dots\}, \quad (3.10)$$

with the aggregated confusion matrix being calculated as follows:

$$\begin{aligned}
m_{ij}^H(x, n) &= \mathbb{P}(\theta^H(x, n) = 1 \mid \theta(x) = i, \epsilon(x) = j) \\
&= \mathbb{P}\left(n^{-1} \sum_{k=1}^n \theta_k^H(x) \geq \frac{1}{2} \mid \theta(x) = i, \epsilon(x) = j\right) \\
&= \sum_{k=\lceil n/2 \rceil}^n \binom{n}{k} (m_{ij}^H(x))^k (1 - m_{ij}^H(x))^{n-k}. \quad (3.11)
\end{aligned}$$

Observe that the third equality in (3.11) follows from the fact that the human-generated label  $\theta_k^H(x)$  follows a Bernoulli distribution under the probability measure  $\mathbb{P}(\cdot \mid \theta(x) = i, \epsilon(x) = j)$  for each  $x \in \mathcal{X}$  and  $k = 1, 2, \dots$ .

### 3.4 Simulation Study

In this section, we conduct a simulation study on the stochastic model presented in the previous section, specifically examining the operational impact of implementing different labeling policies. Our primary goal is to compare dynamic labeling policies against static ones. The following two questions guide our analysis:

1. Do dynamic labeling policies offer better performance than static labeling policies, or are static policies sufficient?
2. How do static and dynamic labeling policies affect learning speed and system congestion?

Before addressing these questions, let us define what we mean by a static labeling policy and describe the specific dynamic labeling policy under consideration. Static labeling policies consistently send a predetermined number of copies of a particular type of content, regardless of the state of the system such as the ML accuracy and the queue length at the human moderation system. Specifically, a labeling policy  $\pi$  is said to be a static if

$$n_t(x, y) = c(x, y), \quad x \in \mathcal{X}, \quad y \in \mathcal{Y}, \quad t \geq 0, \quad (3.12)$$

where  $c(x, y) \in \{0, 1, 2, \dots\}$  are fixed integers for all  $(x, y) \in \mathcal{X} \times \mathcal{Y}$ . If  $c(x, y) = 0$ , the policy does not send any copies of class  $x$  content of virality  $y$ ; if  $c(x, y) > 0$  it sends  $c(x, y)$  copies to the human moderation system.

We next outline a dynamic labeling policy designed to minimize expected misclassification costs while considering the current accuracy of the ML classifier. Specifically, given the choice of sending  $n$  copies of class  $x$  content with virality  $y$  at time  $t$ , the expected type I and type

II error cost is given as follows:

$$\bar{C}_t(x, y, n) := \begin{cases} \mathbb{P}(\theta_t^{\text{ML}}(x) = 0 | \mathcal{F}_t) C_t(x, y, 0, 0) + \mathbb{P}(\theta_t^{\text{ML}}(x) = 1 | \mathcal{F}_t) C_t(x, y, 0, 1), & n = 0, \\ \mathbb{P}(\theta^{\text{H}}(x, n) = 0) C_t(x, y, n, 0) + \mathbb{P}(\theta^{\text{H}}(x, n) = 1) C_t(x, y, n, 1), & n \geq 1. \end{cases} \quad (3.13)$$

This cost function, which is non-convex with respect to the number of labels, highlights the complexity of the decision-making process. The following equation, for  $j = 0, 1$ , explicitly characterizes the terms in the confusion matrix that reflect the ML classifier's accuracy:

$$\bar{m}_{1j}^{\text{ML}}(x, t) = \frac{\bar{\theta}(x) \bar{\epsilon}(x) m_{11}^{\text{ML}}(x, t)^j (1 - m_{11}^{\text{ML}}(x, t))^{1-j}}{\bar{\theta}(x) m_{11}^{\text{ML}}(x, t)^j (1 - m_{11}^{\text{ML}}(x, t))^{1-j} + (1 - \bar{\theta}(x)) m_{00}^{\text{ML}}(x, t)^{1-j} (1 - m_{00}^{\text{ML}}(x, t))^j}.$$

The following equation, for  $j = 0, 1$ , explicitly characterizes the terms in the confusion matrix that reflect the accuracy of the human moderators:

$$\bar{m}_{1j}^{\text{H}}(x, n) = \frac{\bar{\theta}(x) \bar{\epsilon}(x) m_{11}^{\text{H}}(x, n)^j (1 - m_{11}^{\text{H}}(x, n))^{1-j}}{\mathbb{1}_{\{j=0\}} + (-1)^{1-j} \sum_{k=0}^1 \sum_{i=0}^1 m_{ik}^{\text{H}}(x, n) \bar{\theta}(x)^i (1 - \bar{\theta}(x))^{1-i} \bar{\epsilon}(x)^k (1 - \bar{\epsilon}(x))^{1-k}}.$$

By the law of total probability, for  $j = 0, 1$ , the complementary probabilities for the ML classifier and human moderators are:

$$\bar{m}_{0j}^{\text{ML}}(x, t) = 1 - \bar{m}_{1j}^{\text{ML}}(x, t) \quad \text{and} \quad \bar{m}_{0j}^{\text{H}}(x, n) = 1 - \bar{m}_{1j}^{\text{H}}(x, n).$$

Consider the following deterministic optimization problem at each time  $t$ :

$$\underset{\{n_t(x, y) : x \in \mathcal{X}, y \in \mathcal{Y}\}}{\text{minimize}} \quad \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \Lambda(x, y) \bar{C}_t(x, y, n_t(x, y)) \quad (3.14)$$

$$\text{subject to:} \quad \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \Lambda(x, y) n_t(x, y) \leq \delta N \mu \quad (3.15)$$

$$n(x, y) \geq 0, \quad x \in \mathcal{X}, \quad y \in \mathcal{Y}, \quad (3.16)$$

where  $\delta \in [0, 1]$  is a scaling parameter that adjusts the amount of human capacity available to the dynamic decision maker. The solution to this optimization problem (3.14)–(3.16) yields a dynamic labeling policy, which we refer to as the dynamic fluid policy. The objective function represents the expected type I and type II error costs based on the incoming content flow. The first constraint ensures the stability of the content queue for human review, and the second constraint is a natural physical restriction requiring that the number of labels collected be nonnegative.

We conducted a simulation study to compare the efficacy of static labeling policies against dynamic fluid policies, using Python in a Google Colaboratory notebook. The core of our simulation is the “ContentModerationSystemSimulation” class, which models an online content moderation system based on the stochastic framework outlined in Section 3.3. This class includes various functions that accommodate content arrivals, service completions, and labels generated by both the ML classifier and human moderators. We ensure the accuracy of our simulation by validating the average number of jobs in waiting and in service against the theoretical average derived from the Erlang C formula, under exponential service times and under a static labeling policy. Our simulation is designed to flexibly handle various configurations, including an arbitrary number of content classes, arbitrary number of virality levels, arbitrary number of human moderators, arbitrary class and idiosyncratic distribution parameters, arbitrary ML classifier and human moderator confusion matrix parameters, arbitrary cost parameters, arbitrary service time distributions, and an arbitrary time horizon.

For  $\delta \in [0, 1]$ , we consider two distinct policies: (1) “Fluid  $\delta$ ” policy, which is the dynamic fluid policy using  $\delta$  as the scaling parameter, and (2) “Static  $10 \times \delta$ ” policies, which are static policies such that  $\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} n(x, y) \Lambda(x, y) = \delta N \mu$ . In particular, the Fluid  $\delta$  policy makes available  $(10 \times \delta)\%$  of the human moderation capacity, mirroring the Static  $10 \times \delta$  policies which utilize on average  $(10 \times \delta)\%$  of the capacity. Therefore, the Fluid  $\delta$  policy and the Static  $10 \times \delta$  policies are directly comparable. In subsequent figures, we will illustrate the comparative

performance of these policies in terms of learning improvement and system congestion. “Learning improvement” under each policy is quantified as the total accuracy improvement of the ML system throughout the simulation, calculated as  $\sum_{x \in \mathcal{X}} [m_{11}^{\text{ML}}(x, T) - 1 + \alpha(x)]$ , over the time horizon  $[0, T]$ . “System congestion” under a policy is quantified as the time average queue length of the human moderation system over the course of the simulation.

We test our policies within a simulated system characterized by three content classes and two virality levels: nonviral and viral. Our simulation is based on several assumptions and simplifications. First, we assume that labels for an individual piece of content are aggregated via majority vote in the human moderation system. Second, we assume that the service times are distributed as i.i.d. Gamma random variables. Third, to manage computational complexity, our simulation is designed to process thousands of content arrivals within a set time horizon. In future research, we aim to refine the simulation to handle billions of content arrivals—akin to those on large-scale social media platforms—while maintaining a reasonable runtime. Figure 3.2 presents a base case scenario where the learning improvement rate is set at  $\beta(x) = 0.001$  for all  $x \in \mathcal{X}$ , and where the static policies only send content from one content class. All other model parameters remain constant throughout the discussion.




Figure 3.2: A graph illustrating the relationship between learning improvement and system congestion in a system with three content classes and two virality levels, utilizing static policies that label content exclusively from one class.

We observe from the figure that dynamic fluid policies lead to significantly shorter queue lengths compared to the corresponding static policies. Additionally, over time, dynamic fluid policies surpass static policies in terms of learning improvement. These advantages are consistent across all scenarios involving one-class static policies. Figure 3.3 presents a scenario identical to the base case, with the exception that the learning improvement rate has been increased to  $\beta(x) = 0.01$  for all content classes  $x \in \mathcal{X}$ , up from 0.001.

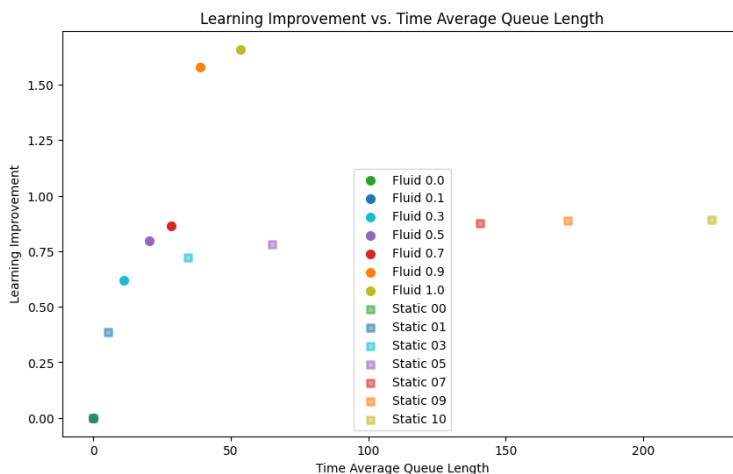


Figure 3.3: A graph illustrating the relationship between learning improvement and system congestion in a system with three content classes and two virality levels, but with a faster learning improvement rate.

Similar to observations in the system with the slower learning improvement rate, this figure demonstrates that dynamic fluid policies continue to outperform static policies in terms of system congestion and learning improvement. Notably, there is a dramatic increase in learning improvement as we transition from the Fluid 0.7 policy to the Fluid 0.9 policy. This is attributable to the increased capacity at when  $\delta = 0.9$ , which allows for labeling content from a second class—unlike at  $\delta = 0.7$ , where capacity limitations restrict labeling to only one high-cost class. The faster learning improvement rate makes this effect more pronounced within the same simulation time horizon, suggesting a “bang-bang” structure for the dynamic fluid policy. Finally, Figure 3.4 examines a system with the same parameters as in the base case, except that the static policies now label content from two different content classes.




Figure 3.4: A graph illustrating the relationship between learning improvement and system congestion in a system with three content classes and two virality levels, utilizing static policies that label content exclusively from two classes.

Observations from this figure align with behaviors noted in the previous two figures. These results continue to hold true across various two-class static policies and are also upheld when considering three-class labeling policies. The simulation study demonstrates that dynamic fluid policies consistently outperform static policies in terms of system congestion and eventually surpass them in learning improvement. These findings underscore the value of adopting dynamic labeling policies over static ones in online content moderation systems.

### 3.5 A Related Model via a Bayesian Learning Framework

In this section, we propose a related stochastic modeling framework for online content moderation via Bayesian learning. Notably, the model we present does not incorporate the idiosyncratic features that were present in the previous model. Ultimately, we provide a weak convergence result reflecting the learning dynamics of the platform. This result paves the way for studying a suitable diffusion control problem. We leave this as a direction for future research.

### 3.5.1 *User-Generated Content and Content Scores*

As before, we consider an online platform that dynamically moderates user-generated content. We again let  $\mathcal{X}$  denote set of content classes and  $\mathcal{Y}$  denote the set of virality levels. We assume that  $\mathcal{X}$  and  $\mathcal{Y}$  are countably infinite metric space with corresponding Borel  $\sigma$ -algebras  $\mathcal{B}(\mathcal{X})$  and  $\mathcal{B}(\mathcal{Y})$ , respectively. Because content with higher virality generates higher misclassification costs, we require  $\mathcal{Y}$  to be a totally ordered and bounded set, with larger values corresponding to high levels of virality.

Notably, the underlying metric on the space of content classes  $\mathcal{X}$  captures how similar two pieces of content are in terms of their class attributes, and we assume that content with similar attributes are also similar with respect to how harmful they are.<sup>1</sup> However, developing a reliable and robust metric to capture the differences between content remains an ongoing challenge. This is because the nature of online content is constantly evolving, and new forms of harmful content can emerge at any time. Additionally, the platform’s community standards may change over time in response to societal and cultural shifts, adding further complexity to the problem. Therefore, it is important for researchers to continue exploring new approaches and techniques for measuring and categorizing online content, in order to improve the accuracy and effectiveness of content moderation systems. In this model, however, we focus on moderating content under a fixed metric and we leave these other important questions open for future research.

The platform assigns a score to each piece of content based on a predefined (finite) set of rules that reflect their community standards. These scores measure various aspects of the content, such as the presence of sexually explicit material or hate speech. It is essential for the platform to clearly communicate its community standards to ensure that the rules consistently enforce these standards. The set of rules is given by  $\mathcal{S} := \{1, \dots, S\}$ , and the

---

1. The metric space structure of  $\mathcal{X}$  offers the potential for modeling transfer learning. However, at present, we have not incorporated this aspect into this model. Exploring the integration of transfer learning is a potential avenue for future research.



violation status of content  $x$  is encoded by a vector of scores:

$$v(x) = (v_1(x), \dots, v_S(x)) \in \{0, 1\}^S, \quad x \in \mathcal{X}, \quad (3.17)$$

where  $v_s(x) = 1$  if class  $x$  content meets the posting standards imposed by rule  $s$ , and  $v_s(x) = 0$  otherwise. The platform does not know the true vector of scores  $v(x)$  but learns it over time based on the review decisions of the human moderators. The platform uses a Bayesian learning framework to learn the true violation status of the content. The prior belief that class  $x$  content satisfies the posting standards imposed by rule  $s$  is given by:

$$\bar{v}_s(x) := \mathbb{P}(v_s(x) = 1) \in (0, 1), \quad x \in \mathcal{X}, \quad s \in \mathcal{S}. \quad (3.18)$$

That is,  $v_s(x)$  to be a Bernoulli random variable with parameter  $\bar{v}_s(x)$ . We assume that the underlying scores are independent across all content classes  $x$  and rules  $s$ . Therefore, the prior belief that class  $x$  content satisfies the posting standards of the platform is given by<sup>2</sup>

$$\bar{\theta}(x) := \mathbb{P}(v_1(x) = v_2(x) = \dots = v_S(x) = 1) = \prod_{s \in \mathcal{S}} \bar{v}_s(x), \quad x \in \mathcal{X}. \quad (3.19)$$

### 3.5.2 Content Arrival Process

The content moderation system receives content from three primary sources: (1) newly uploaded user content, (2) previously posted content that has been flagged by users as harmful, and (3) user appeals from previously removed content. Although this model will not differentiate between these sources of content, they are implicitly incorporated into the

---

2. The precise probabilistic framework is as follows. We assume that all random elements are defined on a probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ . For each pair  $(x, s) \in \mathcal{X} \times \mathcal{S}$ , we define the probability measures

$$\mathbb{P}_{s0}^x(\cdot) := \mathbb{P}(\cdot | v_s(x) = 0), \quad \mathbb{P}_{s1}^x(\cdot) := \mathbb{P}(\cdot | v_s(x) = 1).$$

Then, for each  $\delta \in [0, 1]$ , we define the probability measure  $\mathbb{P}_{s\delta}^x(\cdot) := (1 - \delta)\mathbb{P}_{s0}^x(\cdot) + \delta\mathbb{P}_{s1}^x(\cdot)$ , and we denote by  $\mathbb{E}_{s\delta}^x[\cdot]$  the corresponding expectation operator (cf. pages 334–336 of Peskir and Shiryaev [2006]).

model primitives that follow.

Users on the platform continuously generate new pieces of content over time, and we model this content generation process as a Poisson process  $M = (M_t : t \geq 0)$  with a rate of  $\Lambda > 0$  pieces of content per unit of time. Therefore, the arrival epoch of the  $k$ th piece of content to the system is defined as:

$$\tau_k := \inf \{t \geq 0 : M_t = k\}, \quad k = 1, 2, \dots$$

For content arriving at time  $t$ , we denote by  $X_t$  the content's class and  $Y_t$  the content's virality score. We assume that the random variables  $X_t$  and  $Y_t$  arriving at time  $t$  are drawn independently across time according to a probability distribution  $\mathbb{Q}_t : \mathcal{B}(\mathcal{X}) \otimes \mathcal{B}(\mathcal{Y}) \rightarrow [0, 1]$ . Thus, the probability that content arriving in the system at time  $t$  is of class  $x$  with virality level  $y$  is given as follows:

$$q_t(x, y) := \mathbb{Q}_t(\{x\} \times \{y\}) = \mathbb{P}(X_t = x, Y_t = y). \quad (3.20)$$

We note that the dynamic nature of the content generation distributions captures the phenomenon of concept drift that is present in temporally changing environments.

### 3.5.3 Automated Content Moderation

The role of the ML classifier in this Bayesian learning model of an online content moderation system is to automatically report a vector of scores that reflect the platform's perceived violation status of content over time. To that end, for  $x \in \mathcal{X}$  and  $s \in \mathcal{S}$ , we define

$$r_{st}(x) := \mathbb{P}(r_s(x) = 1 \mid \mathcal{F}_t), \quad t \geq 0, \quad (3.21)$$

where  $\mathcal{F}_t$  is the platform’s information available at time  $t$ . In words,  $r_{st}(x)$  represents the conditional probability that class  $x$  content satisfies the posting standards imposed by rule  $s$  given the platform’s information available at time  $t$ .

We define the belief process for class  $x$  content to be the  $S$ -dimensional stochastic process  $r(x) = (r_t(x), t \geq 0)$ , where  $r_t(x) = (r_{st}(x), s \in \mathcal{S})$  is given by (3.21). Thus, the belief process captures the platform’s evolving understanding of the violation status of the content as new information becomes available to the platform. Since the vector of scores  $r(x)$  given by (3.17) is independent across its components, the platform’s estimated probability that class  $x$  content simultaneously satisfies all the posting guidelines at time  $t$  is given as follows:

$$\theta_t(x) := \mathbb{P}(r_{1t}(x) = r_{2t}(x) = \dots = r_{St}(x) = 1) = \prod_{s \in \mathcal{S}} r_{st}(x) \in [0, 1], \quad x \in \mathcal{X}. \quad (3.22)$$

A higher value of  $\theta_t(x)$  indicates a stronger belief that the content satisfies the guidelines, while a lower value implies a lower confidence level. This probability-based measure assists in assessing the compliance of the content with the platform’s standards. For simplicity, we also refer to the process  $\theta(x) = (\theta_t(x), t \geq 0)$  as the belief process for class  $x$  content.

#### 3.5.4 Human-Based Content Moderation

As in the previous model, the human content moderation system comprises human reviewers (agents) who manually review content (jobs) sent from the ML system. Each piece of content that is evaluated by a human reviewer receives a label indicating whether or not it should be posted on the platform. As before, we assume that the human reviewers are inherently noisy decision-makers and are susceptible to error. Furthermore, we assume that the human reviewers vary in their ability to evaluate content of different types and the various rules enforced by the platform. To account for this heterogeneity, we assume that there is a finite set  $\mathcal{J} := \{1, \dots, J\}$  of pools of human reviewers, where humans in the same pool have the

same skill set in evaluating content. More specifically, when a piece of class  $x$  content is evaluated by a human reviewer in pool  $j$ , it is assigned a label  $y_s^j(x) \in \{0, 1\}$  for each rule  $s$ . Agents in server in pool  $j$  label content  $x$  according to the following confusion matrix:

$$p_{si}^j(x) := \mathbb{P}(y_s^j(x) = i \mid \bar{r}_s(x) = i), \quad s \in \mathcal{S}, \quad i \in \{0, 1\}, \quad (3.23)$$

where  $0.5 < p_{si}^j(x) < 1$ . The confusion matrix captures the classification accuracy of the agents in the various server pools. It is important to note that the entries of the confusion matrix depend not only on the content type and the rule being evaluated, but also on the underlying true label of the content. For instance, if the rule pertains to hate speech, the probability of observing a removal decision may be higher if the content contains more explicit or aggressive language. Conversely, if the rule relates to nudity, the likelihood of observing a removal decision may be greater if the content includes more explicit or graphic images. We assume that the labels produced by a human reviewer are independent across rules and across time, as well as across the reviewers themselves.

As mentioned above, the platform employs a Bayesian learning framework to update its beliefs about the true violation status of content over time. Specifically, it utilizes Bayes' rule to calculate the posterior probability of content satisfying the platform's posting standards, given the observed labels from human moderators. This approach helps the system learn the true labels and reduce content misclassification errors while simultaneously improving the accuracy of the ML-classifier over time. To further reduce misclassification errors, we assume that the platform can send multiple copies of the same piece of content to multiple human reviewers. As discussed in the initial model, enabling the platform to send multiple copies of the same piece of content to different reviewers provides control over the rate at which the true labels of various content types are learned over time, which is a crucial aspect of the human moderation system's operations. To formalize these concepts, we introduce the following definition:

**Definition 6.** A labeling policy for class  $x$  content is a  $J$ -dimensional stochastic process

$$n(x) = ((n_t(x, j) : j \in \mathcal{J}), t \geq 0), \quad x \in \mathcal{X}, \quad (3.24)$$

where  $n_t(x, j) \in \{0, 1, 2, \dots\}$  represents the number of copies of class  $x$  content arriving at time  $t$  that are assigned to human reviewers in pool  $j$ .

It is worth noting that when  $n_t(x, j) = 0$  for all  $j \in \mathcal{J}$ , the content is automatically posted or removed after being evaluated by the ML classifier. On the other hand, if  $n_t(x, j) > 0$  for some  $j \in \mathcal{J}$ , then the content is sent for human review. Different labeling policies can have a significant impact on the performance of the overall moderation system. For example, assigning too few copies of content to human reviewers may result in delayed learning of the true label, while sending too many copies may lead to redundant and unnecessary evaluations, wasting the platform's resources. The goal is to strike a balance between the speed of learning the true labels and the efficient utilization of resources.

We model the human moderation system as a queueing system with  $J$  server pools, where each server pool  $j \in \mathcal{J}$  consists of  $N_j$  servers, an infinite-capacity buffer, and an exponential service rate of  $\mu_j > 0$  at each of its servers. The service rate  $\mu_j$  reflects the speed at which the human moderators in pool  $j$  can review content, and is assumed to be the same across content types. We denote by  $Q_j(t)$  the amount of content (jobs) in pool  $j$  that are either waiting or being reviewed at time  $t$ . Under a labeling policy  $\pi = (n(x) : x \in \mathcal{X})$ , the dynamics of the process  $Q_j = \{Q_j(t), t \geq 0\}$  are given as follows:

$$Q_j(T) = Q_j(0) + \int_0^T n_t(X_t, j) dM_t - D_j \left( N_j \mu_j \int_0^T \min\{N_j, Q_j(t)\} dt \right), \quad (3.25)$$

where  $Q_j(0) \in \{0, 1, 2, \dots\}$  is the initial number of pool  $j$  jobs and  $D_j = \{D_j(t), t \geq 0\}$  is a unit-rate Poisson process independent of all other stochastic primitives. The second (resp., third) term in (3.25) corresponds to the total amount of content that has arrived to (resp.,

departed from) pool  $j$  by time  $T$ .

Embedded in the process dynamics (3.25) is the implicit assumption of a first-come-first-serve (FIFO) and non-idling service discipline within each server pool. On the one hand, the FIFO assumption is without loss of optimality because we will assume a fixed holding cost structure across content types.<sup>3</sup> On the other hand, the non-idling assumption is innocuous because the platform can induce idling by simply adjusting its labeling policy to restrict the flow of content to the human reviewers. This can be beneficial for the system in certain circumstances, such as when the system experiences a surge in incoming content, especially content that it does not know much about. Furthermore, once the platform has enough information to make accurate predictions, it may make sense to restrict the flow of content to the human reviewers to avoid unnecessary costs associated with manual moderation.

### 3.5.5 *The Stochastic Control Problem*

When the system posts and removes content, misclassification errors can occur. Type I errors occur when when the system removes innocent content, while type II errors occur when the system posts harmful content. Let  $f^I(x, y, \theta)$  denote the expected type I error cost incurred for removing class  $x$  content with virality  $y$  given the current belief  $\theta$ . Similarly, let  $f^{II}(x, y, \theta)$  denote the expected type II error cost incurred for posting class  $x$  content with virality  $y$  given the current belief  $\theta$ .

Observe that the type I and type II cost functions depend on the virality of a piece of content. This because misclassifying high virality content can result in greater consequences than misclassifying non-viral context. To be more specific, if a piece of viral content that violates platform standards is mistakenly classified as safe, the number of users exposed to the harmful content increases. Conversely, if viral benign content is incorrectly classified as

---

3. When holding costs vary across content types, the chosen service discipline (or scheduling protocol) at each of the server pools is left as a matter for dynamic decision-making. This is left as an important direction for future research.

unsafe, the platform may suffer greater reputational damage.

**Assumption 10.** The type I and type II error cost functions  $f^I, f^{II} : \mathcal{X} \times \mathcal{Y} \times [0, 1] \rightarrow [0, \infty)$  satisfy the following conditions:

- (1)  $f^I(x, y, 0) = f^{II}(x, y, 1) = 0$  for all  $x \in \mathcal{X}$  and  $y \in \mathcal{Y}$ ,
- (2)  $f^I(x, y, \cdot)$  is continuously increasing on  $[0, 1]$  for all  $x \in \mathcal{X}$  and  $y \in \mathcal{Y}$ ,
- (3)  $f^{II}(x, y, \cdot)$  is continuously decreasing on  $[0, 1]$  for all  $x \in \mathcal{X}$  and  $y \in \mathcal{Y}$ .
- (4)  $f^I(x, y_1, \cdot) \leq f^I(x, y_2, \cdot)$  and  $f^{II}(x, y_1, \cdot) \leq f^{II}(x, y_2, \cdot)$  for all  $x \in \mathcal{X}$  and  $y_1, y_2 \in \mathcal{Y}$  such that  $y_1 \leq y_2$ .

Conditions (1)–(3) in Assumption 10 reflect that the expected type I cost for removing class  $x$  content increases as the platform’s confidence in the content’s compliance with posting guidelines increases. They also reflect that the expected type II cost for posting content  $x$  increases as the platform’s confidence in the content’s violation of posting rules increases. Condition (4) in Assumption 10 reflects the fact that misclassifying content with higher virality is more costly.<sup>4</sup> The minimum expected misclassification cost associated to posting or removing content  $x$  with current belief  $\theta$  is given by

$$f(x, y, \theta) := \min \left\{ f^I(x, y, \theta), f^{II}(x, y, \theta) \right\}. \quad (3.26)$$

We next formulate a stochastic control problem where the platform’s control is the labeling policy  $\pi$ . To that end, the expected infinite-horizon discounted cost under a labeling policy  $\pi$  is given as follows:

$$J_\pi := \mathbb{E} \left[ \int_0^\infty e^{-rt} \sum_{j=1}^J h_j Q_j(t) dt + \int_0^\infty e^{-rt} f(X_t, Y_t, \theta_t(X_t)) dM_t \right], \quad (3.27)$$

---

4. For example,  $f^I(x, y, \theta) = c^I(x, y) \cdot \theta$  and  $f^{II}(x, y, \theta) = c^{II}(x, y) \cdot (1 - \theta)$  satisfy the conditions of Assumption 10, where  $c^I(x, y)$  and  $c^{II}(x, y)$  are positive constants for  $x \in \mathcal{X}$  that are increasing in  $y$ .

where  $r > 0$  is a discount factor and  $h_j > 0$  is a holding cost rate for  $j \in \mathcal{J}$ . The first term in (3.27) captures the operational costs of sending content to the human moderators,<sup>5</sup> while the second term accounts for misclassification costs of posting and removing content. Therefore, the platform's problem is the following:

$$\underset{\pi}{\text{minimize}} \int_0^\infty e^{-rt} \sum_{j=1}^J h_j \mathbb{E}Q_j(t) dt + \int_0^\infty \Lambda e^{-rt} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} q_t(x, y) \mathbb{E}[f(x, y, \theta_t(x))] dt \quad (3.28)$$

$$\text{subject to (3.21)–(3.22) and (3.25).} \quad (3.29)$$

### 3.5.6 Weak Convergence of the Belief Process

Before presenting the weak convergence result and the the appropriate asymptotic regime, we express the dynamics of the belief processes more explicitly using Bayes' rule. To that end, denote by  $n(x) = \{n_1(x), n_2(x), \dots\}$  the labeling policy for class  $x$  content at each of its arrival epochs, where

$$n_k(x) = (n_k(x, j) : j \in \mathcal{J}), \quad k = 1, 2, \dots, \quad (3.30)$$

that is,  $n_k(x, j)$  corresponds to the number of labels assigned by pool  $j$  reviewers to the  $k$ th content  $x$  arrival. Given a labeling policy  $n(x)$ , we denote by  $y(x) = \{y_1(x), y_2(x), \dots\}$  the sequence of (random) labels collected for class  $x$  content, where

$$y_k(x) = (y_{sk}(x, l) \in \{0, 1\} : s \in \mathcal{S}, l \in [n_k(x, j)], j \in \mathcal{J}), \quad k = 1, 2, \dots, \quad (3.31)$$

that is,  $y_{sk}(x, l)$  is the (random) label assigned regarding the violation status of rule  $s$  to the  $l$ th copy of the  $k$ th class  $x$  content arrival sent to reviewers in pool  $j$ . We define the

---

5. Sending content to the human moderators increases their workload and potentially exposes them to harmful content that can result in psychological distress. But sending content to the human moderators enables the platform to collect more information on the true violation status of the content. Thus, we can interpret the holding cost as a cost of learning.



likelihood ratios for a server in pool  $j$  to be

$$\mathcal{L}_{s0}(x, j) := \frac{p_{s0}^j(x)}{1 - p_{s1}^j(x)}, \quad \mathcal{L}_{s1}(x, j) := \frac{1 - p_{s0}^j(x)}{p_{s1}^j(x)}, \quad (x, s) \in \mathcal{X} \times \mathcal{S}. \quad (3.32)$$

We next define the (random) likelihood ratios associated to the labeling policy  $n(x)$  as follows:

$$\mathcal{L}_s(y(x), x, n(x)) = \prod_{j=1}^J \prod_{k=1}^{n_j(x)} \left( \mathcal{L}_{s0}(x, j) \right)^{1 - y_{sk}(x, j)} \left( \mathcal{L}_{s1}(x, j) \right)^{y_{sk}(x, j)}, \quad s \in \mathcal{S}. \quad (3.33)$$

We will use the shorthand notation  $\mathcal{L}_s(x)$  for  $\mathcal{L}_s(y(x), x, n(x))$ . Also, let us denote by  $\mathcal{L}_s(x)$  the range of possible values for  $\mathcal{L}_s(x)$  and by  $Q_s(L, r_s) := \mathbb{P}(\mathcal{L}_s(x) = L | r_s(x))$  the conditional probability distribution of  $\mathcal{L}_s(x)$  given a vector of beliefs  $r_s(x)$ . We also define the vector  $\mathcal{L}(x) = (\mathcal{L}_s(x) : s \in \mathcal{S})$  and denote by  $\mathcal{L}(x) = (\mathcal{L}_s(x) : s \in \mathcal{S})$  its range and by  $Q(\mathcal{L}, r)$  its conditional probability distribution given  $r$ .

**Lemma 30.** *Let  $M_t(x)$  denote the Poisson process with intensity  $\Lambda(x)$  counting the arrivals of content  $x$  that are labeled. Then the belief process  $r_{st}(x)$  for  $x \in \mathcal{X}$  and  $s \in \mathcal{S}$  evolves according to the following stochastic differential equation:*

$$dr_{st}(x) = \eta_s(r_{st-}(x), \mathcal{L}_{st}(x)) dN_t(x), \quad t \geq 0, \quad (3.34)$$

where  $\mathcal{L}_{st}(x) = \mathcal{L}_s(y_t(x), x, n(x))$  and  $y_t(x)$  denotes the random labels of a piece of content  $x$  that arrives at time  $t$ .

*Proof.* By Bayes' rule, the belief process can be rewritten as

$$r_{st}(x) = \frac{r_{st-}(x)}{r_{st-}(x) + (1 - r_{st-}(x)) \mathcal{L}_s(x)} = r_{st-}(x) + \eta(r_{st-}(x), \mathcal{L}_s(x)),$$

where  $\eta(r, \mathcal{L}) := [r(1-r)(1-\mathcal{L})] / [r + (1-r)\mathcal{L}]$ . Therefore, to complete the proof, we can proceed exactly as in the proofs of Lemmas 1 and 2 in Araman and Caldentey [2022].  $\square$

In what follows, we simplify the notation by dropping the dependence on  $x$ . To establish weak convergence of  $r_t$  to a diffusion process, it suffices to establish the convergence of its instantaneous drift vector and volatility matrix. To this end, the following lemma will prove useful. To state it, we first define

$$Q(\mathcal{L}, r) := \prod_{s \in \mathcal{S}} Q_s(\mathcal{L}_s, r_s) = \prod_{s \in \mathcal{S}} [r_s Q_s(\mathcal{L}_s, 1) + (1 - r_s) Q_s(\mathcal{L}_s, 0)].$$

**Lemma 31.** *For all  $s \in \mathcal{S}$  and  $r_s \in [0, 1]$ , we have that*

$$\sum_{\mathcal{L}_s \in \mathcal{L}_s} \eta(r_s, \mathcal{L}_s) Q_s(\mathcal{L}_s, r_s) = 0.$$

*Proof.* From the definition of the likelihood ratio  $\mathcal{L}_s$  we have

$$\begin{aligned} \mathcal{L}_s(y(x), x, n(x)) &= \prod_{j=1}^J \prod_{k=1}^{n_j(x)} (\mathcal{L}_{s0}(x, j))^{1-y_{sk}(x, j)} (\mathcal{L}_{s1}(x, j))^{y_{sk}(x, j)} \\ &= \prod_{j=1}^J \prod_{k=1}^{n_j(x)} \left( \frac{p_{s0}(x, j)}{1 - p_{s1}(x, j)} \right)^{1-y_{sk}(x, j)} \left( \frac{1 - p_{s0}(x, j)}{p_{s1}(x, j)} \right)^{y_{sk}(x, j)} \\ &= \frac{\prod_{j=1}^J \prod_{k=1}^{n_j(x)} (p_{s0}(x, j))^{1-y_{sk}(x, j)} (1 - p_{s0}(x, j))^{y_{sk}(x, j)}}{\prod_{j=1}^J \prod_{k=1}^{n_j(x)} (1 - p_{s1}(x, j))^{1-y_{sk}(x, j)} (p_{s1}(x, j))^{y_{sk}(x, j)}} \\ &= \frac{Q_s(\mathcal{L}_s, 0)}{Q_s(\mathcal{L}_s, 1)}. \end{aligned}$$

From this and the definition of  $\eta(r, \mathcal{L})$  it follows that

$$\begin{aligned} &\sum_{\mathcal{L}_s \in \mathcal{L}_s} \eta(r_s, \mathcal{L}_s) Q_s(\mathcal{L}_s, r_s) \\ &= \sum_{\mathcal{L}_s \in \mathcal{L}_s} \frac{r_s(1-r_s)(1-\mathcal{L}_s)}{r_s + (1-r_s)\mathcal{L}_s} [r_s Q_s(\mathcal{L}_s, 1) + (1-r_s) Q_s(\mathcal{L}_s, 0)] \\ &= \sum_{\mathcal{L}_s \in \mathcal{L}_s} \frac{r_s(1-r_s)(Q_s(\mathcal{L}_s, 1) - Q_s(\mathcal{L}_s, 0))}{r_s Q_s(\mathcal{L}_s, 1) + (1-r_s) Q_s(\mathcal{L}_s, 0)} [r_s Q_s(\mathcal{L}_s, 1) + (1-r_s) Q_s(\mathcal{L}_s, 0)] \end{aligned}$$

$$= r_s (1 - r_s) \sum_{\mathcal{L}_s \in \mathcal{L}_s} (Q_s(\mathcal{L}_s, 1) - Q_s(\mathcal{L}_s, 0)) = 0,$$

which completes the proof.  $\square$

We next consider an asymptotic regime where the amount of content arriving to the human moderation system is large and where the information that the platform received per unit of labeled content is low. This is analogous to the “high frequency vs. low informativeness” regime formalized in Araman and Caldentey [2022] in the context of learning via sequential experimentation.

**Assumption 11.** We consider a sequence of problem instances indexed by  $n$  in which the arrival rate  $\Lambda^n$  and the probabilities  $Q_s^n(\mathcal{L}_s, 0)$  and  $Q_s^n(\mathcal{L}_s, 1)$  satisfy the following:

- (a)  $\Lambda^n = n \Lambda$ .
- (b)  $\lim_{n \rightarrow \infty} \sqrt{n} [Q_s^n(\mathcal{L}_s, 1) - Q_s^n(\mathcal{L}_s, 0)] = \alpha(\mathcal{L}_s)$  for  $i = 0, 1$ .

The next proposition provides a weak convergence result for the belief process to a diffusion. Prior to stating it, we first note that the  $S$ -dimensional process  $r_t(x)$  evolves as a pure-jump Markov process whose infinitesimal generator is given by

$$\mathcal{A}f(r) := \int_w [f(r+w) - f(r)] K(r, dw), \quad r \in [0, 1],$$

where the transition kernel  $K(r, w)$  satisfies

$$\int_w f(w) K(r, dw) = \Lambda \sum_{\mathcal{L} \in \mathcal{L}} f(\eta(r, \mathcal{L})) Q(\mathcal{L}, r).$$

**Proposition 6.** *Under Assumption 11 and under a static labeling policy  $\pi$ , we have that  $r^{\pi, n}(x) \Rightarrow \tilde{r}^\pi(x)$  as  $n \rightarrow \infty$ , where  $\tilde{r}_s(x) = (\tilde{r}_{st}(x), t \geq 0)$  is a diffusion process that satisfies*

the following stochastic differential equation:

$$d\tilde{r}_{st}^\pi(x) = \sqrt{2\Lambda} \tilde{r}_{st}^\pi(x) (1 - \tilde{r}_{st}^\pi(x)) \sigma_s dW_{st}(x),$$

where  $W_s(x) = (W_{st}(x), t \geq 0)$  is a standard Brownian motion, independent across rules and content classes.

*Proof.* Fix a static labeling policy  $\pi$  and a content class  $x$ . It suffices to establish the convergence of its instantaneous drift vector and volatility matrix. On the one hand, by Lemma 31, it follows that the instantaneous drift  $b(r)$  of the process  $r_t(x)$  satisfies

$$b(r) = \int_w w K(r, dw) = \Lambda \sum_{\mathcal{L} \in \mathcal{L}} \eta(r, \mathcal{L}) Q(\mathcal{L}, r) = 0.$$

On the other hand, the instantaneous volatility matrix  $c(r) = [c_{ss'}(r) : s, s' \in \mathcal{S}]$  satisfies

$$\begin{aligned} c_{ss'}(r) &= \int_w w_s w_{s'} K(r, dw) = \Lambda \sum_{\mathcal{L} \in \mathcal{L}} \eta(r_s, \mathcal{L}_s) \eta(r_{s'}, \mathcal{L}_{s'}) Q(\mathcal{L}, r) \\ &= \Lambda \sum_{\mathcal{L}_s \in \mathcal{L}_s} \sum_{\mathcal{L}_{s'} \in \mathcal{L}_{s'}} \eta(r_s, \mathcal{L}_s) \eta(r_{s'}, \mathcal{L}_{s'}) Q_s(\mathcal{L}_s, r_s) Q_{s'}(\mathcal{L}_{s'}, r_{s'}), \quad s, s' \in \mathcal{S}. \end{aligned}$$

If  $s \neq s'$ , then Lemma 31 implies that

$$c_{ss'}(r) = \Lambda \left( \sum_{\mathcal{L}_s \in \mathcal{L}_s} \eta(r_s, \mathcal{L}_s) Q_s(\mathcal{L}_s, r_s) \right) \left( \sum_{\mathcal{L}_{s'} \in \mathcal{L}_{s'}} \eta(r_{s'}, \mathcal{L}_{s'}) Q_{s'}(\mathcal{L}_{s'}, r_{s'}) \right) = 0.$$

If, however,  $s = s'$  then we have that

$$\begin{aligned} c_{ss}(r) &= \Lambda \sum_{\mathcal{L}_s \in \mathcal{L}_s} \eta^2(r_s, \mathcal{L}_s) Q_s(\mathcal{L}_s, r_s) \\ &= \Lambda r_s^2 (1 - r_s)^2 \sum_{\mathcal{L}_s \in \mathcal{L}_s} \frac{(Q_s(\mathcal{L}_s, 1) - Q_s(\mathcal{L}_s, 0))^2}{r_s Q_s(\mathcal{L}_s, 1) + (1 - r_s) Q_s(\mathcal{L}_s, 0)}. \end{aligned}$$

Thus,  $c(r)$  is a diagonal matrix. Next, observe that Assumption 11 implies that

$$\sum_{\mathcal{L}_s \in \mathcal{L}_s} \alpha(\mathcal{L}_s) = 0 \quad \text{and} \quad \lim_{k \rightarrow \infty} Q_s^n(\mathcal{L}_s, i) = \frac{1}{2}, \quad i = 0, 1,$$

and define  $\sigma_s^2 := \sum_{\mathcal{L}_s \in \mathcal{L}_s} \alpha^2(\mathcal{L}_s)$  for  $s \in \mathcal{S}$ . Therefore, the instantaneous volatility  $c_{ss}^n(r)$  for  $s \in \mathcal{S}$  converges asymptotically to

$$\begin{aligned} \lim_{n \rightarrow \infty} c_{ss}^n(r) &= \lim_{n \rightarrow \infty} \Lambda^n r_s^2 (1 - r_s)^2 \sum_{\mathcal{L}_s \in \mathcal{L}_s} \frac{(Q_s^n(\mathcal{L}_s, 1) - Q_s^n(\mathcal{L}_s, 0))^2}{r_s Q_s^n(\mathcal{L}_s, 1) + (1 - r_s) Q_s^n(\mathcal{L}_s, 0)} \\ &= 2\Lambda r_s^2 (1 - r_s)^2 \sigma_s^2. \end{aligned}$$

Finally, by invoking Theorem 4.21 in Chapter IX in Jacod and Shiryaev [2003] on the convergence of pure-jump Markov processes to diffusion processes, we conclude that  $\tilde{r}_{st}(x)$  satisfies the following stochastic differential equation:

$$d\tilde{r}_{st}^\pi(x) = \sqrt{2\Lambda} \tilde{r}_{st}^\pi(x) (1 - \tilde{r}_{st}^\pi(x)) \sigma_s(x) dW_{st}(x), \quad s \in \mathcal{S},$$

where  $W_s(x)$  is a Weiner process, such that  $W_s(x)$  and  $W_{s'}(x)$  are independent for  $s \neq s'$ .  $\square$

### 3.6 Concluding Remarks

This chapter proposed a stochastic model for online content moderation that integrates both machine learning classifiers and human moderators. Our model accommodates multiple content classes, various levels of content virality, and unique idiosyncratic features specific to individual pieces of content. Current research efforts are focused on conducting a simulation study to examine the operational impacts of various labeling policies. Future directions to this study will include exploring a broader array of labeling policies, such as various state-dependent policies, and conducting sensitivity analyses to determine the effects of

different model parameters.

Another future direction includes analytically studying the corresponding optimal control problem. One approach within this context involves considering the system as it approaches heavy traffic conditions, where the number of content arrivals and the number of human moderators increase grow without bound. For instance, in Section 3.5, we demonstrated how to approximate the platform’s “belief process” using a diffusion process within a suitable asymptotic regime. However, the study of the optimal control problem remains a formidable challenge due to the inherent multidimensionality of this problem.

## REFERENCES

- Milton Abramowitz and Irene A. Stegun. *Handbook Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publications, New York, 2003.
- Kranthi M. Adusumilli and John J. Hasenbein. Dynamic admission and service rate control of a queue. *Queueing Systems*, 66(2):131–154, 2010.
- Philipp Afèche, Zhe Liu, and Costis Maglaras. Surge pricing and dynamic matching for hotspot demand shock in ride-hailing networks. *Working Paper*, 2021.
- Philipp Afèche, Zhe Liu, and Costis Maglaras. Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. *Manufacturing & Service Operations Management*, 25(5):1890–1908, 2023.
- Saed Alizamir, Francis de Véricourt, and Peng Sun. Diagnostic accuracy under congestion. *Management Science*, 59(1):157–171, 2013.
- Amir A. Alwan and Barış Ata. A diffusion approximation framework for ride-hailing with travel delays. *Working Paper*, 2020.
- Amir A. Alwan, Barış Ata, and Yuwei Zhou. A queueing model of dynamic pricing and dispatch control for ride-hailing systems incorporating travel times. *Queueing Systems*, 106(1–2):1–66, 2024a.
- Amir A. Alwan, René A. Caldentey, Tingrui Shi, and Amy R. Ward. A simulation model for online content moderation: Managing content via machine learning and human labeling. *Working Paper*, 2024b.
- American Psychological Association. Health advisory on social media use in adolescence. <https://www.apa.org/topics/social-media-internet/health-advisory-adolescent-social-media-use.pdf>, 2023.
- Elene Anton, Rhonda Righter, and Ina M. Verloop. Scheduling under redundancy. *Performance Evaluation Review*, 50(2):30–32, 2022.
- Elene Anton, Rhonda Righter, and Ina M. Verloop. Efficient scheduling in redundancy systems with general service times. *Queueing Systems*, 2024.
- Victor F. Araman and René A. Caldentey. Diffusion approximations for a class of sequential experimentation problems. *Management Science*, 68(8):5958–5979, 2022.
- Mohammad S. Ardekani and Janis M. Orlowski. Multiple listing in kidney transplantation. *American Journal of Kidney Diseases*, 55(4):717–725, 2010.
- Itai Ashlagi and Alvin E. Roth. New challenges in multihospital kidney exchange. *American Economic Review*, 102(3):354–359, 2012.

- Barış Ata. Dynamic power control in a wireless static channel subject to a quality-of-service constraint. *Operations Research*, 53(5):842–851, 2005.
- Barış Ata. Dynamic control of a multiclass queue with thin arrival streams. *Operations Research*, 54(5):876–892, 2006.
- Barış Ata and Nasser Barjesteh. An approximate analysis of dynamic pricing, outsourcing, and scheduling policies for a multiclass make-to-stock queue in the heavy traffic regime. *Operations Research*, 2023.
- Barış Ata and Sunil Kumar. Heavy traffic analysis of open processing systems with complete resource pooling: Asymptotic optimality of discrete review policies. *The Annals of Applied Probability*, 15(1A):331–391, 2005.
- Barış Ata and Wuqin Lin. Heavy traffic analysis of maximum pressure policies for stochastic processing networks with multiple bottlenecks. *Queueing System*, 59:191–235, 2008.
- Barış Ata and Tava L. Olsen. Near-optimal dynamic lead-time quotation and scheduling under convex-concave customer delay costs. *Operations Research*, 57(3):753–768, 2009.
- Barış Ata and Tava L. Olsen. Congestion-based leadtime quotation and pricing for revenue maximization with heterogeneous customers. *Queueing Systems*, 73(1):35–78, 2013.
- Barış Ata and Shiri Shneorson. Dynamic control of an M/M/1 service system with adjustable arrival and service rates. *Management Science*, 52(11):1778–1791, 2006.
- Barış Ata and Mustafa H. Tongarlak. On scheduling a multiclass queue with abandonments under general delay costs. *Queueing Systems*, 74(1):65–104, 2013.
- Barış Ata and Konstantinos E. Zachariadis. Dynamic power control in a fading downlink channel subject to an energy constraint. *Queueing Systems*, 55(1):41–69, 2007.
- Barış Ata, J. Michael Harrison, and Lawrence A. Shepp. Drift rate control of a Brownian processing system. *The Annals of Applied Probability*, 15(2):1145–1160, 2005.
- Barış Ata, Anton Skaro, and Sridhar Tayur. Organjet: Overcoming geographical disparities in access to deceased donor kidneys in the united states. *Management Science*, 63(9):2776–2794, 2017.
- Barış Ata, Deishin Lee, and Erkut Sönmez. Dynamic volunteer staffing in multicrop gleaning operations. *Operations Research*, 67(2):295–314, 2019.
- Barış Ata, Nasser Barjesteh, and Sunil Kumar. Dynamic matching and centralized relocation in ridesharing platforms. Working Paper, 2020a.
- Barış Ata, Nasser Barjesteh, and Sunil Kumar. Spatial pricing: An empirical analysis of taxi rides in New York City. Working Paper, 2020b.



- Bariş Ata, Mustafa H. Tongarlak, Deishin Lee, and Joy Field. A dynamic model for managing volunteer engagement. *Operations Research*, 2024.
- Brooke Auxier and Monica Anderson. Social media use in 2021. Technical report, Pew Research Center, 2021. URL <https://www.pewresearch.org/internet/2021/04/07/social-media-use-in-2021/>.
- Siddhartha Banerjee, Carlos Riquelme, and Ramesh Johari. Pricing in ride-share platforms: A queueing-theoretic approach. Working Paper, 2016.
- Siddhartha Banerjee, Yash Kanoria, and Pengyu Qian. State dependent control of closed queueing networks with application to ride-hailing. Working Paper, 2018.
- Siddhartha Banerjee, Yash Kanoria, and Pengyu Qian. Dynamic assignment control of a closed queueing network under complete resource pooling. Working Paper, 2020.
- Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. Pricing and optimization in shared vehicle systems: An approximation framework. *Operations Research*, 70(3):1783–1805, 2022.
- Harry Bateman and Arthur Erdélyi. *Higher Transcendental Functions, Volume I*. McGraw-Hill, New York, 1953.
- Steven L. Bell and Ruth J. Williams. Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: Asymptotic optimality of a threshold policy. *The Annals of Applied Probability*, 11(3):608–649, 2001.
- Steven L. Bell and Ruth J. Williams. Dynamic scheduling of a parallel server system in heavy traffic with complete resource pooling: Asymptotic optimality of a threshold policy. *Electronic Journal of Probability*, 10:1044–1115, 2005.
- Dimitris Bertsimas, Patrick Jaillet, and Sébastien Martin. Online vehicle routing: The edge of optimization in large-scale applications. *Operations Research*, 67(1):143–162, 2019.
- Omar Besbes, Francisco Castro, and Ilan Lobel. Surge pricing and its spatial supply response. *Management Science*, 67(3):1350–1367, 2021.
- Omar Besbes, Francisco Castro, and Ilan Lobel. Spatial capacity planning. *Operations Research*, 70(2):1271–1291, 2022.
- Patrick Billingsley. *Convergence of Probability Measures*. John Wiley & Sons, Inc., New York, 2nd edition, 1999.
- Kostas Bimpikis, Ozan Candogan, and Daniela Saban. Spatial pricing in ride-sharing networks. *Operations Research*, 67(3):744–769, 2019.
- Maury Bramson and Jim G. Dai. Heavy traffic limits for some queueing networks. *The Annals of Applied Probability*, 11(1):49–90, 2001.

- Anton Braverman, Jim G. Dai, Xin Liu, and Lei Ying. Empty-car routing in ridesharing systems. *Operations Research*, 67(5):1437–1452, 2019.
- Sid Browne and Ward Whitt. Piecewise-linear diffusion processes. *Advances in Queueing: Theory, Methods, and Open Problems*, pages 463–480, 1995.
- Amarjit Budhiraja and Arka P. Ghosh. A large deviations approach to asymptotically optimal control of crisscross network in heavy traffic. *The Annals of Applied Probability*, 15(3): 1887–1935, 2005.
- Amarjit Budhiraja, Xin Liu, and Subhamay Saha. Construction of asymptotically control for crisscross network from a free boundary problem. *Stochastic Systems*, 6(2):459–518, 2016.
- G erard P. Cachon, Kaitlin M. Daniels, and Ruben Lobel. The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, 19(3):337–507, 2017.
- Cambridge Consultants. Use of AI in online content moderation. Technical report, Office of Communications (Ofcom), London, England, United Kingdom, 2019. URL <https://www.ofcom.org.uk/research-and-data/online-research/online-content-moderation>.
- Juan C. Castillo, Dan Knoepfle, and E. Glen Weyl. Surge pricing solves the wild goose chase. Working Paper, 2018.
- Sabri  elik and Costis Maglaras. Dynamic pricing and lead-time quotation for a multiclass make-to-order queue. *Management Science*, 54(6):1132–1146, 2008.
- Sarah H. Cen, Aleksander Madry, and Devavrat Shah. A user-driven framework for regulating and auditing social media. Working Paper, 2023.
- Adrian Chen. The laborers who keep dick pics and beheadings out of your Facebook feed. <https://www.wired.com/2014/10/content-moderation/>, 2014.
- Adrian Chen. The human toll of protecting the internet from the worst of humanity. <https://www.newyorker.com/tech/annals-of-technology/the-human-toll-of-protecting-the-internet-from-the-worst-of-humanity>, 2017.
- Hong Chen, Ping Yang, and David D. Yao. Control and scheduling in a two-station queueing network: Optimal policies and heuristics. *Queueing Systems*, 18(3–4):301–331, 1994.
- M. Keith Chen and Michael Sheldon. Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 2016.
- Qi Chen, Yanzhe Lei, and Stefanus Jasin. Real-time spatial-intertemporal dynamic pricing for balancing supply and demand in a network. Working Paper, 2020.

- Jim G. Dai and Wuqin Lin. Maximum pressure policies in stochastic processing networks. *Operations Research*, 53(2):197–218, 2005.
- Jim G. Dai and Wuqin Lin. Asymptotic optimality of maximum pressure policies in stochastic processing networks. *The Annals of Applied Probability*, 18(6):2239–2299, 2008.
- Stewart N. Ethier and Thomas G. Kurtz. *Markov Processes: Characterization and Convergences*. John Wiley & Sons, Inc., New York, 2005.
- Lin Fan and Peter W. Glynn. Diffusion approximations for thompson sampling. Working Paper, 2021.
- Benoit Frénay and Michel Verleysen. Classification in the presence of label noise: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 25(5):845–869, 2014.
- Kristen Gardner. Correlation in redundancy systems. *Queueing Systems*, 100(3–4):197–199, 2022.
- Kristen Gardner, Mor Harchol-Balter, Esa Hyyti a, and Rhonda Righter. Scheduling for efficiency and fairness in systems with redundancy. *Performance Evaluation*, 116:1–25, 2017.
- Nikhil Garg and Hamid Nazerzadeh. Driver surge pricing. *Management Science*, 68(5):3219–3235, 2021.
- Jennifer M. George and J. Michael Harrison. Dynamic control of a queue with adjustable service rate. *Operations Research*, 49(5):720–731, 2001.
- Arka P. Ghosh and Ananda P. Weerasinghe. Optimal buffer size for a stochastic processing network in heavy traffic. *Queueing Systems*, 55(3):147–159, 2007.
- Arka P. Ghosh and Ananda P. Weerasinghe. Optimal buffer size and dynamic rate control for a queueing system with impatient customers in heavy traffic. *Stochastic Processes and Their Applications*, 120(11):2103–2141, 2010.
- Tarleton Gillespie. Content moderation, AI, and the question of scale. *Big Data & Society*, 7(2), 2020.
- Peter W. Glynn. Diffusion approximations. *Handbooks in OR & MS, Vol. 2, Stochastic Models*, pages 145–198, 1990.
- Bilal Gokpinar and Cemil Selcuk. Selection of prices and commissions in a spatial model of ride-hailing. Working Paper, 2019.
- Harish Guda and Upender Subramanian. Your Uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. *Management Science*, 65(5):1995–2014, 2019a.

- Harish Guda and Upender Subramanian. Your Uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. *Management Science*, 65(5):1995–2014, 2019b.
- Pengfei Guo and Refael Hassin. Equilibrium strategies for placing duplicate orders in a single server queue. *Operations Research Letters*, 43(3):343–348, 2015.
- J. Michael Harrison. Brownian models of queueing networks with heterogeneous customer populations. *Stochastic Differential Systems, Stochastic Control Theory and Applications*, IMA Volumes in Mathematics and its Applications(10):147–186, 1988.
- J. Michael Harrison. *The BIGSTEP Approach to Flow Management in Stochastic Processing Networks*. Oxford University Press, 1996.
- J. Michael Harrison. Heavy traffic analysis of a system with parallel servers: Asymptotic optimality of discrete-review policies. *The Annals of Applied Probability*, 8(3):822–848, 1998.
- J. Michael Harrison. Brownian models of open processing networks: Canonical representation of workload. *The Annals of Applied Probability*, 10(1):75–103, 2000.
- J. Michael Harrison. A broader view of Brownian networks. *The Annals of Applied Probability*, 13(3):1119–1150, 2003.
- J. Michael Harrison. *Brownian Models of Performance and Control*. Cambridge University Press, Cambridge, UK, 2013.
- J. Michael Harrison and Marcel J. López. Heavy traffic resource pooling in parallel-server systems. *Queueing Systems*, 33(4):339–368, 1999.
- J. Michael Harrison and Jan A. Van Mieghem. Dynamic control of Brownian networks: State space collapse and equivalent workload formulations. *The Annals of Applied Probability*, 7(3):747–771, 1997.
- J. Michael Harrison and Lawrence M. Wein. Scheduling networks of queues: Heavy traffic analysis of a simple open network. *Queueing Systems*, 5(4):265–280, 1989.
- Long He, Zhenyu Hu, and Meilin Zhang. Robust repositioning for vehicle sharing. *Manufacturing & Service Operations Management*, 22(2):241–256, 2020.
- Mahsa Hosseini, Joseph Milner, and Gonzalo Romero. Dynamic relocations in car-sharing networks. Working Paper, 2021.
- Wei-Kang Hsu, Jiaming Xu, Xiaojun Lin, and Mark R. Bell. Integrated online learning and adaptive control in queueing systems with uncertain payoffs. *Operations Research*, 70(2):1166–1181, 2022.

- Bin Hu, Ming Hu, and Han Zhu. Surge pricing and two-sided temporal responses in ride-hailing. *Management & Service Operations Management*, 24(1):91–109, 2022.
- Ming Hu and Yun Zhou. Price, wage and fixed commission in on-demand matching. Working Paper, 2020.
- Ming Hu and Yun Zhou. Dynamic type matching. *Management & Service Operations Management*, 24(1):125–142, 2021.
- Jagan Jacob and Ricky Roet-Green. Ride solo or pool: Designing price-service menus for a ride-sharing platform. *European Journal of Operations Research*, 295(3):1008–1024, 2021.
- Jean Jacod and Albert N Shiryaev. *Limit Theorems for Stochastic Processes*. Springer-Verlag, 2nd edition, 2003.
- Shaler Stidham Jr. and Richard R. Weber. Monotonic and insensitive optimal policies for control of queues with undiscounted costs. *Operations Research*, 37(4):611–625, 1989.
- Yash Kanoria and Pengyu Qian. Near optimal control of a ride-hailing platform via mirror backpressure. Working Paper, 2019.
- Samuel Karlin and Howard M. Taylor. *A Second Course in Stochastic Processes*. Academic Press, New York, 1981.
- Jeunghyun Kim and Amy R. Ward. Dynamic scheduling of a GI/GI/1+GI queue with multiple customer classes. *Queueing Systems*, 75(2–4):339–384, 2013.
- Yaakov Kogan and Robert Sh. Liptser. Limit non-stationary behavior of large closed queueing networks with bottlenecks. *Queueing Systems*, 14(1–2):33–55, 1993.
- Yaakov Kogan, Robert Sh. Liptser, and A.V. Smorodinsky. Gaussian diffusion approximation of closed Markov models of computer networks. *Problems of Information Transmission*, 22(1):38–51, 1986.
- A.A. Krichagina and E.V. Puhalskii. A heavy-traffic analysis of a closed queueing system with a GI/ $\infty$  service center. *Queueing Systems*, 25(1–4):235–280, 1997.
- Ravi Kumar, Mark E. Lewis, and Huseyin Topaloglu. Dynamic service rate control for a single-server queue with Markov-modulated arrivals. *Naval Research Logistics*, 60(8):661–677, 2013.
- Sunil Kumar. Two-server closed networks in heavy traffic: Diffusion limits and asymptotic optimality. *The Annals of Applied Probability*, 10(3):930–961, 2000.
- Harold J. Kushner and L. Felipe Martins. Heavy traffic analysis of a controlled multi-class queueing network via weak convergence methods. *SIAM Journal on Control and Optimization*, 34(5):1781–1797, 1996.

- Ricardo Lagos. An alternative approach to search frictions. *Journal of Political Economy*, 108(5):851–873, 2000.
- Chungsang T. Lam and Meng Liu. Demand and consumer surplus in the on-demand economy: The case of ride sharing. Working Paper, 2017.
- Alice Lu, Peter Frazier, and Oren Kislev. Surge pricing moves Uber’s driver partners. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 2018.
- Rahul Makhijani, Parikshit Shah, Vashist Avadhanula, Caner Gocmen, Nicol’as E Stier-Moses, and Juli’an Mestre. QUEST: Queue simulation for content moderation at scale. <https://arxiv.org/abs/2103.16816>, 2021.
- Petr Mandl. *Analytic Treatment of One-Dimensional Markov Processes*. Springer-Verlag, New York, 1968.
- L. Felipe Martins, Steven E. Shreve, and H. Mete Soner. Heavy traffic convergence of a controlled multiclass queueing system. *SIAM Journal on Control and Optimization*, 34(6): 2133–2171, 1996.
- Laurent Massoulié and Kuang Xu. On the capacity of information processing systems. *Operations Research*, 66(2):568–586, 2018.
- Robert M. Merion, Mary K. Guidinger, John M. Newmann, Mary D. Ellison, Friedrich K. Port, and Robert A. Wolfe. Prevalence and outcomes of multiple-listing for cadaveric kidney and liver transplantation. *American Journal of Transplantation*, 4(1):94–100, 2004.
- Meta Platforms. Community standards enforcement report. Facebook Transparency, 2023. URL <https://transparency.fb.com/reports/community-standards-enforcement/>. Accessed: 2024-03-17.
- Leela Nageswaran and Alan Scheller-Wolf. Queues with redundancy: Is waiting in multiple lines fair? *Manufacturing & Service Operations Management*, 24(4):1959–1976, 2022.
- Viet-An Nguyen, Peibei Shi, Jagdish Ramakrishnan, Udi Weinsberg, Henry C. Lin, Steve Metz, Neil Chandra, Jane Jing, and Dimitris Kalimeris. CLARA: Confidence of labels and raters. *KDD ’20: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2542–2552, 2020.
- Erhun Özkan. Joint pricing and matching in ride-sharing systems. *European Journal of Operational Research*, 287(3):1149–1160, 2020.
- Erhun Özkan and Amy R. Ward. Dynamic matching for real-time ridesharing. *Stochastic Systems*, 10(1):29–70, 2020.
- Guodong Pang, Rishi Talreja, and Ward Whitt. Martingale proofs of many-server heavy-traffic limits for Markovian queues. *Probability Surveys*, 4:193–267, 2007.

- Goran Peskir and Albert Shiryaev. *Optimal Stopping and Free-Boundary Problems*. Birkhäuser Verlag, Basel, Switzerland, 2006.
- Wiebe R. Pestman. Measurability of linear operators in the Skorokhod topology. *Bull. Belg. Math. Soc. Simon Stevin*, 2(4):381–388, 1995.
- Andrei D. Polyanin and Valentin F. Zaitsev. *Handbook of Exact Solutions for Ordinary Differential Equations*. Chapman & Hall/CRC, Boca Raton, FL, 2nd edition, 2003.
- Josh Reed and Amy R. Ward. A diffusion approximation for a generalized Jackson network with reneging. In *Proceedings of the 42nd Annual Conference on Communication, Control, and Computing*, 2004.
- Josh Reed and Amy R. Ward. Approximating the GI/GI/1+GI queue with a nonlinear drift diffusion: Hazard rate scaling in heavy traffic. *Mathematics of Operations Research*, 33(3): 606–644, 2008.
- Sarah T. Ruberts. Social media’s silent filter. <https://www.theatlantic.com/technology/archive/2017/03/commercial-content-moderation/518796/>, 2017.
- Melanie Rubino and Barış Ata. Dynamic control of a make-to-order, parallel-server system with cancellations. *Operations Research*, 57(1):94–108, 2009.
- A.V. Smorodinskii. Asymptotic distribution of the queue length of one service system. *Avtomatika i Telemekhanika*, 2:92–99, 1986. In Russian.
- Alexander L. Stolyar. Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *The Annals of Applied Probability*, 14(1):1–53, 2004.
- Kalyan R. Talluri and Garrett van Ryzin. *The Theory and Practice of Revenue Management*. Springer, New York, NY, 2004.
- Sahana Udupa, Elonnai Hickok, Antonis Maronikolakis, Hinrich Schuetze, Laura Csuka, Axel Wisioerek, and Leah Nann. Artificial intelligence, extreme speech, and the challenges of online content moderation. Technical report, AI4 Dignity Project, 2021. URL <https://doi.org/10.5282/ubm/epub.76087>.
- Parsia A. Vagefi, Sandy Feng, Jennifer L. Dodge, James F. Markmann, and John P. Roberts. Multiple listings as a reflection of geographic disparity in liver transplantation. *Journal of the American College of Surgeons*, 219(3):496–504, 2014.
- Sushil M. Varma, Pornpawee Bumpensanti, Siva T. Maguluri, and He Wang. Dynamic pricing and matching for two-sided queues. *Operations Research*, 71(1):83–100, 2022.
- Emily A. Vogels, Risa Gelles-Watnick, and Navid Massarat. Teens, social media and technology 2022. Technical report, Pew Research Center, 2022. URL <https://www.pewresearch.org/internet/2022/08/10/teens-social-media-and-technology-2022/>.

- Stefan Wager and Kuang Xu. Diffusion asymptotics for sequential experiments. Working Paper, 2022.
- Xing Wang, Niels Agatz, and Alan Erera. Stable matching for dynamic ride-sharing systems. *Transportation Science*, 52(4):850–867, 2017.
- Ward Whitt. *Stochastic-Process Limits*. Springer-Verlag, New York, 2002.
- Ruth J. Williams. Diffusion approximations for open multiclass queueing networks: Sufficient conditions involving state space collapse. *Queueing Systems*, 30(1–2):27–88, 1998.
- Chiwei Yan, Helin Zhu, Nikita Korolko, and Dawn Woodard. Dynamic pricing and matching in ride-hailing platforms. *Naval Research Logistics*, 67(8):705–724, 2020.
- Pu Yang, Krishnamurthy Iyer, and Peter Frazier. Mean field equilibria for resource competition in spatial settings. *Stochastic Systems*, 8(4):307–334, 2018.
- Rick Zhang and Marco Pavone. Control of robotic mobility-on-demand systems: A queueing-theoretical perspective. *The International Journal of Robotics Research*, 35(1-3):186–203, 2016.