Dynamic Pricing of Relocating Resources in Large Networks

Santiago R. Balseiro[1], David B. Brown[2], and Chen Chen[2]

[1]Graduate School of Business, Columbia University
[2]Fuqua School of Business, Duke University
srb2155@columbia.edu, dbbrown@duke.edu, cc459@duke.edu

First version: April 2, 2019
This version: May 23, 2020

Abstract

Motivated by applications in shared vehicle systems, we study dynamic pricing of resources that relocate over a network of locations. Customers with private willingness-to-pay sequentially request to relocate a resource from one location to another, and a revenue-maximizing service provider sets a price for each request. This problem can be formulated as an infinite horizon stochastic dynamic program, but is quite difficult to solve, as optimal pricing policies may depend on the locations of all resources in the network. We first focus on networks with a hub-and-spoke structure, and we develop a dynamic pricing policy and a performance bound based on a Lagrangian relaxation. This relaxation decomposes the problem over spokes and is thus far easier to solve than the original problem. We analyze the performance of the Lagrangian-based policy and focus on a *supply-constrained large network* regime in which the number of spokes (n) and the number of resources grow at the same rate. We show that the Lagrangian policy loses no more than $O\big(\sqrt{\ln n/n}\big)$ in performance compared to an optimal policy, thus implying asymptotic optimality as n grows large. We also show that no static policy is asymptotically optimal in the large network regime. Finally, we extend the Lagrangian relaxation to provide upper bounds and policies to general networks with multiple, interconnected hubs and spoke-to-spoke connections, and to incorporate relocation times. We also examine the performance of the Lagrangian policy and the Lagrangian relaxation bound on some numerical examples, including examples based on data from RideAustin.

Subject classifications: Dynamic pricing, resource relocation, hub-and-spoke networks, Lagrangian relaxations, asymptotic optimality.

1 Introduction

Motivated by the growing popularity of shared vehicle systems (e.g., "ride-sharing" or "ride-hailing" platforms such as Lyft and Uber, as well as many bicycle-sharing platforms), there has been an increased study of algorithms for efficiently managing resources that circulate over a network of locations. A number of other applications related to transportation networks or logistics also involve resources that change locations over time. In this paper, we study a dynamic pricing problem faced by a service provider managing a finite number of resources over a potentially large network. Customers (e.g., riders) with a private willingness-to-pay sequentially arrive with known rates and request to relocate one resource (e.g., driver) from an origin location to a destination location. When a request arrives, the service provider selects a price based on the origin and destination, as well as potentially the location (i.e., "state") of all resources in the system. The request is fulfilled if and only if the origin location contains resources and the customer's private value is at least equal to the selected price; a fulfilled request generates revenue (equal to the selected price) for the provider and leads to one unit of resource relocating from the origin to the destination. If the origin has no resource, the request is lost. The problem is to find a dynamic pricing policy that maximizes the provider's average revenue over an infinite horizon. With many locations, this problem is difficult to solve, as optimal pricing policies may depend on the locations of all resources in the system.

The model we study closely follows the models studied in Waserhole and Jost (2016) and Banerjee et al. (2016): both of these papers are motivated by ride-sharing and study the role that pricing plays in controlling the flow of resources while attempting to maximize a particular objective (e.g., revenue, throughput, or social welfare). The key methodology in Waserhole and Jost (2016) and Banerjee et al. (2016) is the use of fluid relaxations that provide upper bounds on the performance of an optimal pricing policy and lead to a static pricing policy (i.e., prices depend on the origin-destination pairs but not on the number of resources across the network). Both Waserhole and Jost (2016) and Banerjee et al. (2016) show that the fluid-based static policy is within a factor of $\frac{m}{m+n-1}$ of the fluid upper bound, where m is the number of resources and n is the number of locations. This result holds for general network structures and implies that the fluid-based static policy is asymptotically optimal in the *large supply* regime where the number of resources grows faster than the number of locations.

Although the large supply regime may be appropriate when resources are concentrated within a few locations (e.g., compact cities with high population density), in this paper we consider instead a setting when resources are distributed over a large number of locations (e.g., sprawling cities with many neighborhoods and heterogeneous population density). This motivates us to study a *supply-constrained large network* regime in which the total supply of resources grows at the same rate as the number of locations. In Section 2 we show that when the total arrival rate of requests grows proportionally with n, resource utilization (measured in average relocations per resource per unit time - e.g., rides per driver per hour in ride-sharing) is bounded from below by a constant in the large network regime we consider.[1] Thus, resources are ensured to be "frequently" in use, and the regime we consider can be relevant in some practical settings.

The large network regime also appears to require significantly different policies and analysis than the large supply regime. In particular, in the large supply regime, the system behavior becomes essentially deterministic as the network becomes congested in the large supply limit, and thus fluid relaxations provide a very good approximation and fluid-based static policies provide a good performance. In contrast, when the number of resources and size of the network grow at the same rate, the limiting behavior of the system retains a stochastic character, and it is essential to adjust prices dynamically based on the locations of resources to attain good performance.

Whereas fluid relaxations have been shown to perform well in the large supply regime for general network structures, our initial theoretical analysis is specialized to networks with a hub-and-spoke structure (see Figure 1) and a large number of spokes. Hub-and-spoke network structures are popular in the transportation industry (perhaps most famously by airlines) because they minimize the number of routes connecting every location. In shared vehicle systems, a hub-and-spoke network provides a reasonable approximation to a metropolitan area in which commuters travel back and forth from suburbs (the spokes) to a densely populated urban core (the hub). In addition to the transportation industry, hub-and-spoke structures are used widely in communications and logistics (see e.g., Pirkul and Schilling 1998 and references therein).

Our approach involves a Lagrangian relaxation of the constraint that the number of resources in the hub be non-negative. With this relaxation, the problem simplifies considerably in that the dynamic pricing problem decouples across spokes. With the optimal Lagrange multiplier, the Lagrangian relaxation has the interpretation of ensuring that the hub has non-negative number of resources in expectation. A key innovation in our approach is to then consider a *perturbed*

[1]This is consistent with findings in the popular press on how companies like Lyft and Uber may "optimize" the total number of drivers: for a given urban area, there may be a "target" number of requests per driver per hour on average. See, e.g., `https://www.nytimes.com/2018/08/10/upshot/uber-lyft-taxi-ideal-number-per-city.html`. Moreover, large cities like NYC are considering capping the number of drivers on road. See, e.g., `https://www.nytimes.com/2018/08/08/nyregion/uber-vote-city-council-cap.html`.

Lagrangian relaxation that ensures that some relatively small amount $\delta = o(n)$ of resources remains in the hub in expectation. We develop a dynamic pricing policy based on a dual formulation of the perturbed Lagrangian relaxation; this policy involves precise, state-dependent control of the resources at each spoke. We show that the performance of this policy in the original (i.e., fully constrained) system is "close" to the performance in the relaxation by relating the performance gap to the probability that the hub is depleted in the original system. Finally, by choosing δ as a particular function of n and analyzing the properties of the stationary distributions of resources in the spokes, we show via concentration inequalities that the hub depletion probability is small; this result implies that the Lagrangian-based policy is asymptotically optimal for large hub-and-spoke networks in which the number of spokes n and resources m grow at the same rate.

We then extend our policies and performance bounds to general networks with multiple, interconnected hubs, and spoke-to-spoke connections. For systems with multiple hubs, we further dualize the flow balance constraint for each hub, i.e., the constraint that, for each hub, the average inflow of resources is equal to the average outflow of resources: this ensures the in-flow and out-flow of each hub is balanced in expectation (with optimal Lagrange multipliers). To handle spoke-to-spoke connections, we further dualize the relocation constraint associated with these requests. In our base model, we assume relocations occur instantaneously. We also develop a tractable approach to incorporating relocation times into the policies and bounds. In this paper, we restrict most of the theoretical analysis to single-hub networks with no spoke-to-spoke transitions, and instantaneous relocations (we also analyze multiple hub systems under some restrictive assumptions on the hubs). Extending the analysis of Lagrangian-based policies to general networks remains an open challenge.

Although the fluid-based static policy is easy to implement and provides provably good performance for general networks, we show that the Lagrangian relaxations we consider provide tighter performance bounds, and no static policy is asymptotically optimal in the large network regime we consider. In some sense, the Lagrangian policy can be viewed as a refinement of the fluid-based static policy in which we price dynamically for requests involving a node with fewer resources (i.e., a spoke), but retain static pricing for requests in which both the origin and destination tend to have a large number of resources (i.e., hubs). Our dynamic pricing policies may depend on the number of resources either at the origin or the destination location. When all the nodes in the network are treated as hubs, the Lagrangian relaxation and policy are identical to the fluid relaxation and the fluid-based static policy. In Section 7.2, we show the Lagrangian policy performs well and outperforms both the fluid-based static policy and another Lagrangian-based static policy in general networks using a model calibrated on empirical data from RideAustin.

The paper is organized as follows. Section 1.1 reviews some related work. In Section 2 we formulate the problem, and in Section 3 we discuss the Lagrangian relaxation and corresponding Lagrangian policies. Section 4 presents the main results on performance analysis and an overview of the key steps of the proof. In Section 5 we compare to fluid relaxations and the optimal static policies, and provide an analytical example showing that the optimal static policy is strictly suboptimal in the large network regime we consider. Section 6 extends the Lagrangian relaxation and policy to more general network structures with multiple hubs, spoke-to-spoke connections, and nonzero relocation times. In Section 7 we evaluate the performance of the Lagrangian policies on single hub synthetic examples and a realistic example based on data from RideAustin. Section 8 concludes. The main proofs are in Appendix A; additional discussions and derivations are in Appendices B through G.

1.1 Related Literature

In addition to Waserhole and Jost (2016) and Banerjee et al. (2016), a number of other researchers have studied stochastic control problems involving relocating resources. Braverman et al. (2016) also study ride-sharing problems and focus on how to optimally reposition resources after each service to under-supplied locations to better meet future demands; they consider both arrival rates and number of resources growing large for a fixed network of locations, but incorporate positive transition times. Their approach also involves the study of fluid limits: they show the fluid-based upper bounds and their induced static policies are asymptotically optimal in the heavy traffic limit, whereas we show these fail to be asymptotically optimal in the large network regime we study. Ozkan and Ward (2017) consider instead the assignment problem of matching requests with nearby resources and show that a fluid-based static (matching) policy achieves asymptotic optimality in the heavy traffic regime. Banerjee et al. (2018) consider a similar assignment problem and develop a state-dependent assignment policy that achieves the optimal decay rate of demand-dropping probability in the large supply regime. Kanoria and Qian (2019) study joint matching and pricing and develop a state-dependent assignment policy that does not require prior knowledge of the demand arrival rates and is asymptotically optimal in the large supply regime. Variations of these models involving equilibrium considerations (e.g., strategic drivers in ride-sharing networks) have also recently been studied: Bimpikis et al. (2016) study spatial pricing in ride-sharing networks with strategic drivers operating in an infinite-horizon fluid model[2], and Besbes et al. (2018) study

[2]Although Bimpikis et al. (2016) focus on spatial pricing with strategic drivers, they too study networks related to hub-and-spoke networks; they refer to these as "star" networks. Bimpikis et al. (2016) argue that spatial dis-

spatial pricing in a static equilibrium model with demand-supply imbalances. Finally, Besbes et al. (2019) study the problem of determining the optimal number of drivers in ride-hailing systems to balance service utilization and customer waiting times while accounting for pickup and travel times, under a heavy traffic regime.

Other researchers have studied similar relocation problems for managing logistics and transportation networks. For example, Adelman (2007) develops a price-directed control policy to manage a network of shipping containers. This policy is based on approximate dynamic programming and the problem studied involves accept-or-reject decisions for requests (rather than a continuum of prices). A number of other papers (e.g., Du and Hall 1997; George and Xia 2011; Song and Carter 2008) consider designing the optimal fleet size and/or redistribution policy to control the flow of equipment from over-supplied locations to under-supplied locations. Du and Hall (1997) and Song and Carter (2008) also consider hub-and-spoke networks and a decomposition over spokes for equipment and vehicle redistribution problems; these papers do not consider Lagrangian penalties in the relaxation. Several of these papers incorporate other realistic features of the problem (e.g., nonzero transit times for relocations), but none of these papers provide theoretical analysis or guarantees on the performance of the various heuristic policies studied. Finally, a number of researchers have used Lagrangian relaxations for variations of hub-and-spoke design problems. For example, Pirkul and Schilling (1998) and An et al. (2015) consider static, discrete optimization problems aimed at optimizing the location of multiple hubs given the possibility that some hubs may fail.

The resource relocation problem we study can be viewed as a control problem for a closed queueing network. For example, Banerjee et al. (2016) use the Gordon-Newell theorem (Gordon and Newell 1967) to analyze the stationary distribution of their static pricing policy, and George and Xia (2011) use results from the study of BCMP networks (Baskett et al. 1975) to analyze optimal fleet sizing problems. In general, fluid limits are a widely used tool in the study of closed queueing networks (see, e.g., Harrison and Wein 1990 and Kumar and Kumar 1996). Typically, an optimal solution to the fluid relaxation problem can be implemented in the original problem as a static policy and proven to be asymptotically optimal in a regime where the number of resources grows faster than the number of locations. These static policies often result in stationary distributions that are separable across the network (i.e., the joint distribution of the system has a product form across nodes in the networks); this separability need not hold in the state-dependent pricing policies

crimination in pricing in their setting confers more profits in networks with demand that is less "balanced," with hub-and-spoke (or star) networks being a prominent example of this.

we study.

For hub-and-spoke networks, the dynamic pricing problem can also be viewed as a special case of a weakly coupled stochastic dynamic program (DP), where the linking constraint is the capacity constraint that the total number of resources in the spokes can not exceed the total number of resources m (or equivalently, the number of resources in the hub must be non-negative). Hawkins (2003), Adelman and Mersereau (2008) and Bertsimas and Mišić (2016) study DPs that are linked through global resource constraints and they all consider Lagrangian relaxations of the linking constraints. Lagrangian relaxations of stochastic DPs have also been used in many applications including network revenue management (e.g., Talluri and van Ryzin 1998; Topaloglu 2009; Kunnumkal and Talluri 2016), inventory allocation (e.g., Marklund and Rosling 2012; Miao et al. 2020), marketing (e.g., Bertsimas and Mersereau 2007 and Caro and Gallien 2007), and restless bandit problems (e.g, Brown and Smith 2017, Hu and Frazier 2017, and Zayas-Cabán et al. 2019).

1.2 Notation and Terminology

We let $\mathbb{N}$ denote the set of nonnegative integers and $\mathbb{N}_+$ the set of strictly positive integers. For any $x, y \in \mathbb{R}$, we let $x \wedge y = \min\{x, y\}$ denote the minimum of x and y. For any $a, b \in \mathbb{N}$ with $a \leq b$, we let $[a : b] = \big\{a, a+1, \ldots, b-1, b\big\}$ denote a sequence of integers starting from a and ending with b and we denote $[n] = [1 : n]$ for any $n \in \mathbb{N}_+$. We let $\boldsymbol{e}_i$ denote a vector with the i-th element being one and all the other elements being zeros; the dimension of $\boldsymbol{e}_i$ is clear from the context.

2 Problem Formulation

We study a dynamic pricing problem with m resources relocating in a network. We consider an infinite horizon continuous-time model. Customers with private willingness-to-pay sequentially request to relocate one resource from one location to another. We assume relocations are instantaneous and after each relocation the resource remains in the destination until relocated again by another request. Requests of each type (i, j) — i.e., those requests to relocate one resource from location i to location j — arrive following independent Poisson processes with rates η_{ij}. We focus on the embedded discrete-time model where in each time period a request arrives and with probability $q_{ij} \triangleq \frac{\eta_{ij}}{\sum_{k,l} \eta_{kl}}$ the request type is (i, j), independent of requests in earlier time periods.[3] We can

[3] The continuous-time model and its embedded discrete-time model are equivalent because, under any stationary policy, the limiting distributions of the continuous-time and discrete-time Markov chains converge to the same

describe the network topology by a directed graph with each node representing a location and there is an edge from location i to location j if $q_{ij} > 0$. Throughout the paper, we assume that the network topology is strongly connected.

Upon the arrival of each request, a service provider selects a price p based on the type of the request and the locations of the resources. If the number of resources at the origin location is positive and the private value for the request is greater than the price, the request is fulfilled, the service provider earns p, and one resource transits from the origin to the destination; otherwise the request is lost. If there is no resource in the origin, the service provider simply drops the request. The problem is to find a feasible policy that maximizes the average revenue per time period over an infinite horizon.

Throughout the paper, we assume that private values are independent. Moreover, the private value for a request of type (i, j) follows a known continuous cumulative distribution F_{ij}. Thus given a price p and assuming that the origin location contains resources, the probability that a request of type (i, j) is fulfilled is $d_{ij}(p) = 1 - F_{ij}(p)$. We let $G_{ij}(d) \triangleq \sup\left\{p : 1 - F_{ij}(p) \geq d\right\}$ be the generalized inverse demand function for $d \in [0, 1]$. The service provider equivalently selects a demand level d in the interval $[0, 1]$ every time a request arrives and charges the price $G_{ij}(d)$. We let $r_{ij}(d) = d \cdot G_{ij}(d)$ be the one-period expected revenue functions, and we assume the following properties on these functions.

Assumption 2.1. For all request types (i, j), the one-period expected revenues $r_{ij}(d) = d \cdot G_{ij}(d)$ are strictly concave in the interval $[0, 1]$ with $r_{ij}(0) = 0$; the maximum values are uniformly bounded by some constant $\bar{r} > 0$, i.e., $\max_{d\in[0,1]} r_{ij}(d) \leq \bar{r}$; the derivatives are uniformly bounded by some $\bar{\omega} > 0$, i.e., $\max_{d\in[0,1]} \left|r'_{ij}(d)\right| \leq \bar{\omega}$; and the unique maximum points $d^*_{ij} \triangleq \mathrm{argmax}_{d\in[0,1]} r_{ij}(d)$ lie in the interior of interval $[0, 1]$.

Strict concavity simplifies our analysis, as this leads to unique solutions of the resulting optimization problems and rules out optimality of randomized prices (e.g., mixed policies). The assumption that the maximizers of the revenue functions are strictly interior is helpful in ensuring the Markov chains under various policies we study have helpful properties (e.g., irreducibility).

2.1 Large Network Regime

We consider a large network regime in which the number of resources m scales at the same rate as the number of locations. When the demand rate per location is constant, this scaling of resources

stationary distribution (which may depend on the initial state if the Markov chain has multiple recurrent classes).

ensures that resources are in use “frequently” under any pricing policy, in a sense we now make precise.

Let n denote the number of locations and π denote any feasible policy; formally, π is a mapping from the state – which is the origin and destination of the current request as well as the vector $\mathbf{x} = (x_i)_{i\in[n]} \in \mathbb{N}^n$, describing how many resources are in each location – to a demand level $d \in [0,1]$. The next result bounds the resource utilization of a policy in terms of the average revenue per period.

Proposition 2.1. *Suppose the total arrival rate satisfies $\underline{\eta} n \leq \sum_{i,j} \eta_{ij} \leq \bar{\eta} n$ for some $\underline{\eta}, \bar{\eta} > 0$. For any policy π, let V^π denote the average revenue per period (i.e., request) using π. Then resource utilization as measured in the average number of relocations (e.g., rides) per resource per unit time*[4]*, which we denote by Φ^π, satisfies*

$$\frac{\underline{\eta} V^\pi}{\bar{\omega}} \cdot \frac{n}{m} \leq \Phi^\pi \leq \bar{\eta} \cdot \frac{n}{m}.$$

We prove the resource utilization inequalities under an assumption on the total arrival rate, which can be justified when locations have a constant demand for relocations (for example, when each location covers a geographical area with a similar population) and we consider increasing the number of locations (for example, by extending the geographical area of coverage). This assumption implies that the expected time between requests is on the order of $1/n$. We prove Proposition 2.1 in Appendix A.1 and note that Proposition 2.1 also holds with relocation times (see the proof for a justification of this).

When m grows at the same rate as n, Proposition 2.1 implies that the resource utilization is bounded from below by a positive constant for any pricing policy π (under the mild condition that V^π is bounded from below by a positive constant). In this sense, resources are ensured to be “frequently” in use, irrespective of the size of the system. In the remainder of this paper we focus on the large network regime and aim to develop policies that perform well in this regime. We begin with networks with a hub-and-spoke structure; Section 6 discusses how to extend our approach to more general networks.

2.2 Hub-and-Spoke Networks: Optimal Control

We first focus on a hub-and-spoke network as illustrated in Figure 1. There is one hub at the center denoted by location 0 and n spokes around the hub denoted by locations 1 to n. In this

[4] By “time” here we refer to the underlying continuous time model. For example, in ride-sharing, how many rides per hour is completed by each driver, on average.

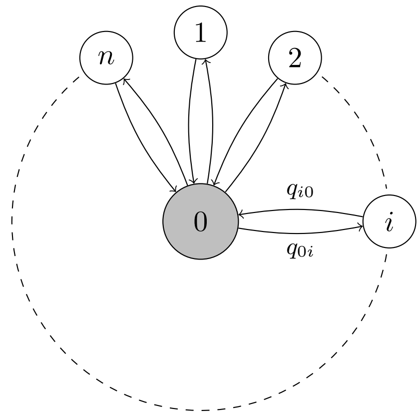

Figure 1: A hub-and-spoke network with one hub (grey) and n spokes (white).

model, the only requests are between the hub and a particular spoke. The request type is $(i, 0)$ with probability $q_{i0} > 0$ and is $(0, i)$ with probability $q_{0i} > 0$ for each spoke i. We let $q_i = q_{i0} + q_{0i}$ be the probability that spoke i is involved in a request; these probabilities q_i sum up to one.

We let x_i denote the number of resources in location i for $i \in [0\!:\!n]$ and let vector $\mathbf{x} = (x_i)_{i\in[n]} \in \mathbb{N}^n$ denote the number of resources in each of the spokes. The number of resources in the hub is uniquely determined by the number of resources in the spokes because $x_0 + \sum_{i\in[n]} x_i = m$. Thus the set $\mathcal{X} = \left\{\mathbf{x} = (x_i)_{i\in[n]} \in \mathbb{N}^n : \sum_{i\in[n]} x_i \leq m\right\}$ incorporates all feasible states of the resources. We can represent the system state by a tuple $(\mathbf{x}, s)$, with $\mathbf{x} \in \mathcal{X}$ being the state of resources and s being the type of the arriving request (i.e., either $(i, 0)$ or $(0, i)$ for some $i \in [n]$). Sometimes we will also represent the system state by $(\mathbf{x}, i, 0)$ or $(\mathbf{x}, 0, i)$ given that the request type is either $(i, 0)$ or $(0, i)$. The average revenue optimization problem is

$$\begin{aligned}
V^{\mathrm{OPT}} = \max_{\pi\in\Pi} \quad & \lim_{T\to\infty} \frac{1}{T} \cdot \mathbb{E}\left\{ \sum_{t=1}^{T} \sum_{i\in[n]} \left(y_{i0,t} \cdot r_{i0}\big(d^{\pi}_{i0,t}\big) + y_{0i,t} \cdot r_{0i}\big(d^{\pi}_{0i,t}\big) \right) \right\} \\
\text{s.t.} \quad & x^{\pi}_{i,t+1} = x^{\pi}_{i,t} - y_{i0,t} \cdot \mathbb{1}\left[\xi_t \leq d^{\pi}_{i0,t}\right] + y_{0i,t} \cdot \mathbb{1}\left[\xi_t \leq d^{\pi}_{0i,t}\right], \ \forall\, i \in [n], t \geq 1, \\
& x^{\pi}_{i,t} \geq 0, \ \forall\, i \in [0:n], t \geq 1, \\
& \sum_{i\in[0:n]} x^{\pi}_{i,t} = m, \ \forall\, t \geq 1,
\end{aligned} \tag{1}$$

where Π is the set of all non-anticipative policies (that is, policies that can only depend on the observed history), $d^{\pi}_{ij,t}$ is the demand level employed by policy π at time t when the request type is (i, j), $x^{\pi}_{i,t}$ is the resource level of location i at the beginning of time t under policy π, $y_{ij,t}$ is a binary random variable with $y_{ij,t} = 1$ if the request type is (i, j) at time t, and $\{\xi_t\}_{t\geq 1}$ is a sequence of independent and identically distributed random variables following a uniform distribution with support $[0,1]$. The first constraint models the dynamics of the number of resources in the spokes; the events $\{\xi_t \le d^{\pi}_{i0,t}\}$ and $\{\xi_t \le d^{\pi}_{0i,t}\}$ capture whether the request is fulfilled and a resource relocates from a spoke to the hub and the hub to a spoke, respectively. The second constraint ensures that locations have non-negative amount of resources while the third constraint determines the number of resources in the hub. We let V^{OPT} be the optimal value of (1). In Proposition 2.2 we show V^{OPT} does not depend on the initial state of the system and provide Bellman equations for the optimal control problem.

Proposition 2.2. *The optimal value V^{OPT} of (1) does not depend on the initial state of the system. Moreover, V^{OPT} together with the differential value functions $v(\mathbf{x}, s)$ for each state $(\mathbf{x}, s)$ satisfies*

$$\begin{aligned} V^{\mathrm{OPT}} + v(\mathbf{x}, 0, i) &= \max_{d \in [0, 1 \wedge x_0]} \Big\{ r_{0i}(d) + d \cdot v(\mathbf{x} + \boldsymbol{e}_i) + (1-d) \cdot v(\mathbf{x}) \Big\}, \\ V^{\mathrm{OPT}} + v(\mathbf{x}, i, 0) &= \max_{d \in [0, 1 \wedge x_i]} \Big\{ r_{i0}(d) + d \cdot v(\mathbf{x} - \boldsymbol{e}_i) + (1-d) \cdot v(\mathbf{x}) \Big\}, \end{aligned} \tag{2}$$

for all $\mathbf{x} \in \mathcal{X}$ and $i \in [n]$, where $v(\mathbf{x}) = \mathbb{E}_s\big[v(\mathbf{x}, s)\big] = \sum_{i \in [n]} \Big\{ q_{0i} \cdot v(\mathbf{x}, 0, i) + q_{i0} \cdot v(\mathbf{x}, i, 0) \Big\}$ are the average differential value functions over request types. Finally, an optimal stationary policy exists; any policy that selects $d^*(\mathbf{x}, s)$ in every state $(\mathbf{x}, s)$, where $d^*(\mathbf{x}, s)$ is optimal in (2), is an optimal stationary policy.

Even though the control space is infinite (i.e., $d \in [0,1]$) in our problem, concavity of the revenue functions $r_{0i}(d)$ and $r_{i0}(d)$ guarantees that there exists a solution to (2). We provide more details in Appendix A.2.

3 Lagrangian Relaxations

The Bellman equation (2) can be difficult to solve due to the "curse of dimensionality": the number of states is in general exponential in the number of spokes n and the number of resources m. For example, if $n = m$, the number of all feasible states of resources is $|\mathcal{X}| = \binom{2n}{n} \approx \frac{4^n}{\sqrt{\pi n}}$. Intuitively, a service provider following an optimal policy has to balance the overall spatial distribution of resources. For example, suppose a request of type $(i, 0)$ arrives. Although the service provider would like to accept the request and collect an immediate revenue, an optimal value for the demand level d depends on the overall spatial distribution of the resources: if the number of resources in spoke i is low (high) relative to the number of resources at other spokes in the network, the service provider should accept the request with a relatively low (high) probability. More generally, the service provider may try to decrease the probability if accepting the request will cause the overall distribution of resources to deviate further from some "ideal" distribution of resources, and vice versa. This behavior may be particularly important in systems in which the number of resources and the number of locations are comparable. In this section we study Lagrangian relaxations that decompose the problem over spokes. The Lagrangian relaxation can then be used to generate a feasible policy as well as an upper bound on the performance V^{OPT} of an optimal policy.

3.1 The Lagrangian Relaxation

We consider a Lagrangian relaxation that relaxes the constraint that the number of resources in the hub be non-negative or, equivalently, the capacity constraint $\sum_{i\in[n]} x_i \le m$, and uses a dual variable $\lambda \ge 0$ to penalize violations of this constraint. The Lagrangian relaxation allows the number of resources in the hub to be negative but in every time period, every resource in the spokes incurs a cost λ. The set of all feasible states of resources in the Lagrangian relaxation is $\bar{\mathcal{X}} = \left\{\mathbf{x} = (x_i)_{i\in[n]} \in \mathbb{N}^n : x_i \le m, \ \forall\, i \in [n]\right\}$; note that here we still require that the number of resources in each spoke not to exceed the total number of resources m. The Lagrangian relaxation can be written as

$$\begin{aligned}
\bar{V}^{\lambda} = \max_{\pi\in\Pi} \quad & \lim_{T\to\infty} \frac{1}{T} \cdot \mathbb{E}\left\{ \sum_{t=1}^{T} \sum_{i\in[n]} \Big(y_{i0,t} \cdot r_{i0}\big(d^{\pi}_{i0,t}\big) + y_{0i,t} \cdot r_{0i}\big(d^{\pi}_{0i,t}\big) \Big) + \lambda \sum_{t=1}^{T} \Big(m - \sum_{i\in[n]} x^{\pi}_{i,t} \Big) \right\} \\
\text{s.t.} \quad & x^{\pi}_{i,t+1} = x^{\pi}_{i,t} - y_{i0,t} \cdot \mathbb{1}\left[\xi_t \le d^{\pi}_{i0,t}\right] + y_{0i,t} \cdot \mathbb{1}\left[\xi_t \le d^{\pi}_{0i,t}\right], \ \forall\, i \in [n], t \ge 1, \qquad (3) \\
& 0 \le x^{\pi}_{i,t} \le m, \ \forall\, i \in [n], t \ge 1.
\end{aligned}$$

following the same notation as in (1). After the relaxation, the number of resources $x^{\pi}_{0,t}$ in the hub is no longer part of the state; however, we can still determine its value, which might be negative, by $x^{\pi}_{0,t} = m - \sum_{i\in[n]} x^{\pi}_{i,t}$. We let $\bar{V}^{\lambda}$ denote the optimal value of (3). In Proposition 3.1 we show that $\bar{V}^{\lambda}$ does not depend on the initial state of the system and it provides an upper bound on V^{OPT}. Moreover, the Lagrangian relaxation decomposes over spokes into a sequence of spoke problems that can be solved independently.

Proposition 3.1. *For any $\lambda \ge 0$, the average revenue $\bar{V}^{\lambda}$ of an optimal policy to the Lagrangian relaxation is independent of the initial state and satisfies $\bar{V}^{\lambda} \ge V^{\mathrm{OPT}}$. Moreover, $\bar{V}^{\lambda}$ decomposes* *over spokes as*

$$\bar{V}^{\lambda} = m\lambda + \sum_{i=1}^{n} h_i^{\lambda}, \tag{4}$$

where h_i^{λ} *is the average revenue of an optimal policy to each spoke problem, independent of the initial state of the spoke. Finally,* h_i^{λ} *together with the differential value functions* $v_i^{\lambda}(x, i, 0)$*,* $v_i^{\lambda}(x, 0, i)$*, and* $v_i^{\lambda}(x, \varnothing)$ *satisfies*

$$\begin{aligned} h_i^{\lambda} + v_i^{\lambda}(x, i, 0) &= \max_{d \in [0, 1 \wedge x]} \left\{ r_{i0}(d) + d \cdot \left(v_i^{\lambda}(x-1) - v_i^{\lambda}(x) \right) \right\} + v_i^{\lambda}(x) - \lambda \cdot x, \\ h_i^{\lambda} + v_i^{\lambda}(x, 0, i) &= \max_{d \in [0, 1 \wedge (m-x)]} \left\{ r_{0i}(d) + d \cdot \left(v_i^{\lambda}(x+1) - v_i^{\lambda}(x) \right) \right\} + v_i^{\lambda}(x) - \lambda \cdot x, \\ h_i^{\lambda} + v_i^{\lambda}(x, \varnothing) &= v_i^{\lambda}(x) - \lambda \cdot x, \end{aligned} \tag{5}$$

for each resource level $x \in [0 : m]$ *of spoke* i*, where* $v_i^{\lambda}(x, i, 0)$*,* $v_i^{\lambda}(x, 0, i)$*, and* $v_i^{\lambda}(x, \varnothing)$ *are the differential value functions for spoke* i *with* x *resources and the request type being* $(i, 0)$*,* $(0, i)$*, or one of any other types respectively, and* $v_i^{\lambda}(x) = q_{i0} \cdot v_i^{\lambda}(x, i, 0) + q_{0i} \cdot v_i^{\lambda}(x, 0, i) + (1 - q_i) \cdot v_i^{\lambda}(x, \varnothing)$ *are the average differential value functions over request types.*

The decomposition in Proposition 3.1 is intuitive because once we allow the number of resources in the hub to be negative, the hub has infinite supply of resources and, as a result, decisions in different spokes no longer affect each other. We prove Proposition 3.1 in Appendix A.3.

Following standard results in dynamic programming, we can solve (5) as an linear program (LP), with variables representing the value functions in every state. Rather than working with this LP, we will instead solve a dual formulation of this LP, with variables representing stationary distributions of the resources and the optimal pricing decisions in every state, as we now describe.

Proposition 3.2. *The spoke-specific optimal revenue satisfies:*

$$\begin{aligned} h_i^{\lambda} = \max_{\substack{d_i(x,i,0) \in [0,1], \\ d_i(x,0,i) \in [0,1], \\ p_i(x) \geq 0}} \quad & \sum_{x=0}^{m} p_i(x) \left[q_{i0} \cdot r_{i0}\Big(d_i(x, i, 0)\Big) + q_{0i} \cdot r_{0i}\Big(d_i(x, 0, i)\Big) \right] - \lambda \cdot \sum_{x=0}^{m} x \cdot p_i(x) \\ \text{s.t.} \quad & \sum_{x=0}^{m} p_i(x) = 1, \\ & p_i(x) \cdot q_{0i} \cdot d_i(x, 0, i) = p_i(x+1) \cdot q_{i0} \cdot d_i(x+1, i, 0), \ \forall \, x \in [0 : m-1], \\ & d_i(0, i, 0) = 0, \quad d_i(m, 0, i) = 0. \end{aligned} \tag{6}$$

We prove Proposition 3.2 in Appendix A.4. We can interpret $p_i(x)$ in problem (6) as the stationary probability that spoke i has x resources, and $d_i(x,i,0)$ and $d_i(x,0,i)$ as the optimal controls when there are x resources and the request type is $(i,0)$ and $(0,i)$, respectively. Problem (6) thus provides an optimal stationary distribution for each spoke and a set of optimal controls to maximize the average revenue net of the Lagrangian penalty. We refer to the set of demand values from a solution to (6) as the *Lagrangian policy*. Once we fix a policy, the dynamics in each spoke follow a birth-and-death chain process, and hence the stationary distribution is reversible, as indicated by the second constraint in (6). Finally, we can convert (6) into the convex optimization problem (7)

$$\begin{aligned} h_i^\lambda = \max_{p_i(x)\geq 0} \quad & \sum_{x=0}^{m-1} p_i(x)\cdot\gamma_i\left(\frac{p_i(x+1)}{p_i(x)}\right) - \lambda\cdot\sum_{x=0}^{m} x\cdot p_i(x) \\ \text{s.t.} \quad & \sum_{x=0}^{m} p_i(x) = 1, \end{aligned} \tag{7}$$

where

$$\begin{aligned} \gamma_i(\beta) = \max_{d_{i0},d_{0i}\in[0,1]} \quad & q_{0i}\cdot r_{0i}(d_{0i}) + \beta\cdot q_{i0}\cdot r_{i0}(d_{i0}) \\ \text{s.t.} \quad & q_{0i}\cdot d_{0i} = \beta\cdot q_{i0}\cdot d_{i0}, \end{aligned} \tag{8}$$

and we set $x\cdot\gamma_i\left(\frac{y}{x}\right) = 0$ if $x=0$. $\gamma_i(\beta)$ gives the optimal revenue between the hub and spoke i when both have infinite capacity and the rate to the spoke is scaled by β. When β is one, this coincides with the optimal revenue of the fluid problem. We prove the equivalence between (6) and (7) in Appendix A.4.2 by partitioning the optimization problem. For fixed values of $p_i(x)$, we can find the optimal values of the controls $d_i(x,0,i)$ and $d_i(x+1,i,0)$ by setting $\beta = p_i(x+1)/p_i(x)$ and solving the problem $\gamma_i(\beta)$. Eliminating the controls gives a simpler problem in terms of the probabilities $p_i(x)$ that can be formulated as (7), and convexity follows because the perspective of a convex function is convex ($\gamma_i(\beta)$ is concave in β by Lemma A.3). In Appendix A.4.3 we provide a specialized algorithm for solving (7) by exploiting its first-order optimality conditions.

3.2 The Lagrangian Policy in the Relaxation

In this section, we formally define the Lagrangian policy derived from an optimal solution to (6). The policies we describe in Section 3.4 and analyze in Section 4 correspond to Lagrangian policies with specific choices of dual variable λ. We discuss how to choose λ in Section 3.4.

We let $p_i(x)$, $d_i(x,i,0)$ and $d_i(x,0,i)$ be an optimal solution to (6). Let $I_i = \{x\in[0:m] : p_i(x)>0\}$ be the support of the probability distribution $p_i(x)$; the set I_i is non-empty because $\sum_{x=0}^{m} p_i(x) = 1$. In Lemma B.1 we show that the set I_i takes the form of $I_i = [0:H_i]$ for some non-negative integer $0 \le H_i \le m$. We use this fact first in the definition of the Lagrangian policy.

Definition 3.1 (Lagrangian Policy)**.** We construct the Lagrangian policy in the following two steps:

1. For a given $\lambda \ge 0$, solve the dual problem (6) for each spoke and let $p_i(x)$, $d_i(x, i, 0)$ and $d_i(x, 0, i)$ denote an optimal solution.
2. Store the values of $d_i(x, i, 0)$ and $d_i(x, 0, i)$ for all resource levels $x \in I_i = [0 : H_i]$, and set $d_i(x, i, 0) = 1$ and $d_i(x, 0, i) = 0$ for any $x > H_i$.[5]

In the Lagrangian relaxation, the *Lagrangian policy* selects $d_i(x, i, 0)$ if the request type is $(i, 0)$ and the resource level of spoke i is x, and selects $d_i(x, 0, i)$ if the request type is $(0, i)$ and the resource level of spoke i is x. The Lagrangian policy in the original problem is defined analogously, with the difference that the policy drops the request if the request type is $(0, i)$ and the hub has no resources.

Lemma B.1 provides several other useful results about the Lagrangian policy for each spoke-specific DP: specifically, we show that the states in I_i form a single positive recurrent class and the corresponding Markov chain is aperiodic; thus the limiting distribution converges to a unique stationary distribution, which is $p_i(x)$, independent of the initial state. Finally, the Lagrangian policy is optimal to each spoke-specific DP.

In Proposition 3.3, we show a key property of the Lagrangian policy: namely, the controls $d_i(x, i, 0)$ and $d_i(x, 0, i)$ are monotone in the resource level x for $x \in I_i$.

Proposition 3.3. *For each spoke $i \in [n]$, the controls $d_i(x, i, 0)$ and $d_i(x, 0, i)$ are increasing and decreasing in x for $x \in I_i$, respectively.*

Figure 2 illustrates the monotonicity property for the controls $d_i(x, i, 0)$ and $d_i(x, 0, i)$. Specifically, when a request $(i, 0)$ arrives, if there are many resources in spoke i, the service provider following the Lagrangian policy will increase the acceptance probability (and hence decrease the price) to encourage resources to relocate out of i, and vice versa.

The monotonicity property in Proposition 3.3 implies that the stationary distribution $p_i(x)$ in (6) is (discrete) log-concave, which is useful for analyzing the performance of the Lagrangian policies (see Proposition B.9 for a formal statement). Intuitively, this suggests that the stationary

[5] We have flexibility to set the controls for any unsupported state $x \in I_i^c$, where I_i^c denotes the complement of I_i, because we only require that the states $x \in I_i^c$ be transient under the Lagrangian policy. Lemma B.1 shows that set I_i is a positive recurrent class under the Lagrangian policy, so it suffices to set $d_i(x, i, 0) > 0$ for any state $x > H_i$.

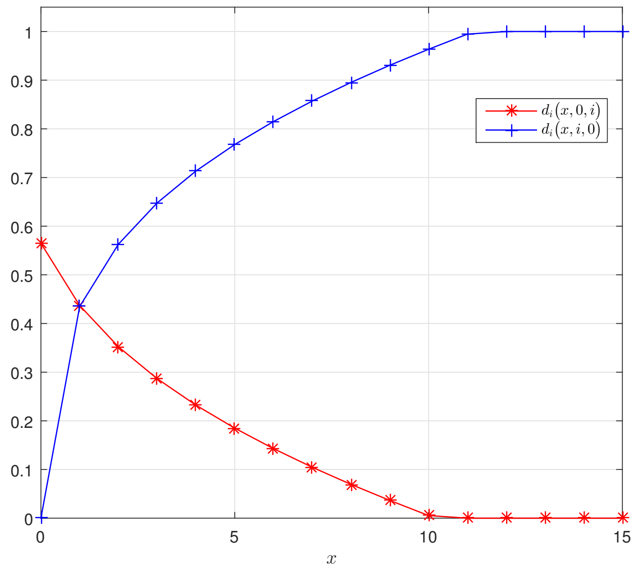

Figure 2: The optimal control from solving (6) for the spoke problem with $m = 20$, $q_{i0} = q_{0i} = 0.05$, and dual variable $\lambda = 0.003$. All private values follow uniform distributions with support $[0, 1]$.

distribution for each spoke under the Lagrangian policy will peak around some "intermediate" states (see Figure 3(a)).

Finally, we show that when considering the full Lagrangian relaxation (across all spokes), the set $\prod_{i=1}^n I_i$ forms a positive recurrent class that is aperiodic: hence, the Lagrangian policy is a unichain policy in the Lagrangian relaxation and all initial states share the same average revenue. Moreover, in the Lagrangian relaxation, the stationary distribution of resources are independent across spokes, which will also be helpful in our ensuing analysis. These facts are presented formally in Corollary B.10.

3.3 The Lagrangian Dual Problem

From Proposition 3.1, we can use the Lagrangian relaxation as an upper bound on the performance of an optimal policy. Although any $\lambda \geq 0$ provides an upper bound, we want to choose $\lambda \geq 0$ to provide the best possible bound. We can write the Lagrangian dual problem as

$$V^{\mathrm{R}} \triangleq \min_{\lambda \geq 0} \bar{V}^{\lambda}, \tag{9}$$

where $\bar{V}^{\lambda}$ is given in (4). We let λ^* denote an optimal solution to (9). The objective $\bar{V}^{\lambda}$ in (9) is convex in λ and (9) can be solved efficiently (e.g., using bisection). We provide more details in Appendix C on different formulations of (9) and on solving (9).

Proposition 3.4 provides necessary and sufficient optimality conditions for λ.

Proposition 3.4. *The dual variable λ^* is optimal to (9) if and only if*

$$(\lambda^* \geq 0) \;\perp\; \left(\sum_{i=1}^{n}\sum_{x=0}^{m} x \cdot p_i(x) \leq m\right), \tag{10}$$

where $p_i(x)$ is optimal to (6) with λ^* and $\perp$ denotes that at least one of these conditions is binding.

It follows that the Lagrangian dual problem (9) is equivalent to the problem of maximizing the average revenue subject to the constraint that the *expected* number of resources in the hub is non-negative.

3.4 Lagrangian Policy in the Original Problem and Perturbed Lagrangian Relaxation

The Lagrangian policy can be implemented in the original (i.e., fully constrained) problem as well, with the only difference that requests from the hub are lost when the hub contains zero resources. Analogous to Corollary B.10, we can show that the Lagrangian policy is a unichain policy in the original problem, and hence with this policy all initial states lead to the same average revenue. This is formalized in Corollary B.11.

One may consider using the Lagrangian policy with an optimal dual variable λ^* as a policy in the original system. However, according to (10), if $\lambda^* > 0$, the capacity constraint $\sum_{i=1}^{n}\sum_{x=0}^{m} x \cdot p_i(x) \leq m$ is binding and using this Lagrangian policy may leave too few resources in the hub: in the Lagrangian relaxation, there are zero resources in the hub in expectation when $\lambda^* > 0$. Instead, we develop policies that attempt to leave the hub with some amount $\delta = o(n)$ of resources on average by solving a "perturbed" Lagrangian relaxation

$$V^{\mathrm{R}}(\delta) = \min_{\lambda \geq 0} \bar{V}^{\lambda} - \delta\lambda. \tag{11}$$

Comparing (11) to (9), and noting the optimality conditions (10), an optimal policy to (11) ensures the expected number of resources in the hub is at least δ. Let $\lambda^*(\delta)$ denote an optimal solution to (11). We consider the Lagrangian policy with dual variable $\lambda^*(\delta)$. We denote this Lagrangian policy by $\pi(\delta)$ and denote its performance in the original problem by $V^{\pi}(\delta)$. Note that although $V^{\mathrm{R}}(0) = V^{\mathrm{R}}$ is the tightest upper bound from the Lagrangian relaxation, $V^{\mathrm{R}}(\delta)$ is not necessarily an upper bound on V^{OPT} when $\delta > 0$.

The Lagrangian policy only depends on the resource level of the spoke involved in the request and the request type, and is not necessarily optimal to the original problem. In Section 4, we analyze the performance of the policy $\pi(\delta)$ and we show that, with a specific choice of δ, the performance gap $V^{\mathrm{OPT}} - V^{\pi}(\delta)$ converges to zero under a scaling where the number of spokes increases and the number of resources per location remains fixed.

4 Performance Analysis

In this section we provide our main results. Theorem 4.1 first bounds the performance gap of our policy by a term proportional to the probability the hub runs out of resources plus a perturbation term. In Sections 4.3 and 4.4 we bound the probability that the hub runs out of resources when spokes satisfy some regularity conditions.

Theorem 4.1. *For any $\lambda^*(\delta)$ optimal to (11), the corresponding Lagrangian policy $\pi(\delta)$ satisfies*

$$V^{\pi}(\delta) \le V^{\mathrm{OPT}} \le V^{\mathrm{R}} \le V^{\pi}(\delta) + \bar{r} \cdot \frac{\delta}{m-\delta} + (\bar{r} + \bar{\omega}) \cdot \mathbb{P}\Big[X_0(\delta) = 0\Big],$$

where $\mathbb{P}\big[X_0(\delta) = 0\big]$ is the stationary probability that the hub runs out of resources in the original problem under the policy $\pi(\delta)$.

The first inequality of the theorem follows because our policy is feasible and the second from Proposition 3.1. The final inequality in Theorem 4.1 then follows from two key steps:

1. We bound from above the Lagrangian relaxation bound V^{R} in terms of the optimal value $V^{\mathrm{R}}(\delta)$ of the perturbed problem (11) plus a term that is proportional to $\frac{\delta}{m-\delta}$ (Lemma 4.2). The analysis in this step is deterministic in nature and uses results from sensitivity analysis for convex optimization.
2. We bound from below the performance $V^{\pi}(\delta)$ of the policy $\pi(\delta)$ in terms of the optimal value $V^{\mathrm{R}}(\delta)$ of the perturbed problem (11) minus a term that is proportional to the stationary probability that the hub runs out of resources in the original problem (Lemma 4.3). The analysis in this step is probabilistic in nature and involves studying the dynamics induced by the Lagrangian policy in the relaxed and original systems.

Theorem 4.1 then follows by combining the upper bound from step 1 and the lower bound from step 2.

4.1 Bounding $V^{\mathrm{R}} - V^{\mathrm{R}}(\delta)$

In this section we bound from above the gap $V^{\mathrm{R}} - V^{\mathrm{R}}(\delta)$ between the Lagrangian bound and the optimal value of the perturbed problem (11). Recall that we denote by $\lambda^*(\delta)$ an optimal solution to (11). Since $-\lambda^*(\delta)$ is a super-gradient of $V^{\mathrm{R}}(\delta)$, we have

$$V^{\mathrm{R}}(\delta) \le V^{\mathrm{R}}(0) = V^{\mathrm{R}} \le V^{\mathrm{R}}(\delta) + \lambda^*(\delta) \cdot \delta, \tag{12}$$

where the first inequality is because $V^{\mathrm{R}}(\delta)$ is decreasing in δ. It is not hard to show that for any $\delta < m$, the optimal dual variable $\lambda^*(\delta)$ satisfies $\lambda^*(\delta) \le \frac{\bar{r}}{m-\delta}$. This implies that when $\frac{m}{n}$ is a fixed ratio, for any $\delta = o(n)$, the term $\lambda^*(\delta) \cdot \delta$ in (12) goes to zero and as a result $V^{\mathrm{R}}(\delta)$ will converge to the Lagrangian upper bound V^{R}.

Putting these facts together leads to the following result.

Lemma 4.2. *The Lagrangian upper bound V^{R} and the optimal value $V^{\mathrm{R}}(\delta)$ of (11) satisfy*

$$V^{\mathrm{R}}(\delta) \le V^{\mathrm{R}}(0) = V^{\mathrm{R}} \le V^{\mathrm{R}}(\delta) + \bar{r} \cdot \frac{\delta}{m-\delta}.$$

4.2 Bounding $V^{\mathrm{R}}(\delta) - V^{\pi}(\delta)$

In this section we bound the gap $V^{\mathrm{R}}(\delta) - V^{\pi}(\delta)$ between the optimal value of the perturbed problem (11) and the average revenue of the policy $\pi(\delta)$ in the original problem.

We first consider a relaxation in which we drop the constraint $x_0 \ge 0$ and we retain all other constraints. We refer to this as the *relaxed system*, and this is equivalent to the Lagrangian relaxation without the terms corresponding to the Lagrangian penalty (i.e., with $\lambda = 0$). The average revenue of the Lagrangian policy $\pi(\delta)$ in the relaxed system equals $V^{\mathrm{R}}(\delta)$; this follows from complementary slackness as given in (10) with m replaced by $m - \delta$. Thus obtaining a bound on $V^{\mathrm{R}}(\delta) - V^{\pi}(\delta)$ is equivalent to obtaining a bound on the difference between the average revenues of the Lagrangian policy in the relaxed and the original systems.

According to Corollaries B.10 and B.11, in both systems, the limiting distribution converges to a unique stationary distribution starting from any initial state. We let the random variables $X_i(\delta)$ and $\tilde{X}_i(\delta)$ denote the number of resources in location $i \in [0:n]$ under the stationary distributions of the Lagrangian policy $\pi(\delta)$ in the original and the relaxed systems, respectively. In Lemma 4.3 we bound the gap $V^{\mathrm{R}}(\delta) - V^{\pi}(\delta)$ in terms of $\mathbb{P}\big[X_0(\delta) = 0\big]$, the stationary probability that the hub runs out of resources in the original problem.

Lemma 4.3. *The average revenues $V^{\mathrm{R}}(\delta)$ and $V^{\pi}(\delta)$ of the policy $\pi(\delta)$ in the relaxed and the* *original systems satisfy*

$$V^{\mathrm{R}}(\delta) - V^{\pi}(\delta) \le (\bar{r} + \bar{\omega}) \cdot \Big(\sum_{i\in[n]} q_{0i}\Big) \cdot \mathbb{P}\big[X_0(\delta) = 0\big] \le (\bar{r} + \bar{\omega}) \cdot \mathbb{P}\big[X_0(\delta) = 0\big].$$

We prove Lemma 4.3 in Appendix A.8 by first showing that the value function of the Lagrangian policy in the relaxed system approximately solves the Bellman equation of the original system along the path induced by the Lagrangian policy in the original system. We then use a verification theorem to bound the total loss between these two systems. The first step follows because the Lagrangian policy takes different actions in these two systems at the same state $(\mathbf{x}, s)$ only when $x_0 = 0$ and $s = (0, i)$ for some $i \in [n]$, i.e., when the hub is depleted and there is a request originating from the hub. By using the monotonicity property of the Lagrangian policy in Proposition 3.3, we show that every time the Lagrangian policy differs in the two systems, the difference in continuation value functions is at most $\bar{r} + \bar{\omega}$. Hence the difference of the average revenues in the two systems can be bounded from above by $\bar{r} + \bar{\omega}$ times the probability that the hub runs out of resources in the original system. Our approach is similar in spirit to the "compensated coupling" argument used in Vera and Banerjee (2018).

Lemma 4.2 and Lemma 4.3 imply Theorem 4.1. In Lemma 4.4 we show that for each spoke $i \in [n]$, the number of resources $\tilde{X}_i(\delta)$ in the relaxed system first-order stochastically dominates the number of resources $X_i(\delta)$ in the original system.

Lemma 4.4. *The number of resources $\tilde{X}_i(\delta)$ in each spoke $i \in [n]$ of the relaxed system first-order stochastically dominates the number of resources $X_i(\delta)$ in the original system. The number of resources $\tilde{X}_0(\delta)$ in the hub of the relaxed system is first-order stochastically dominated by the number of resources $X_0(\delta)$ in the original system.*

We prove Lemma 4.4 in Appendix A.9 by coupling $X_i(\delta)$ and $\tilde{X}_i(\delta)$ based on the same arrival sequence of requests. Intuitively, since the hub in the relaxed system has infinite supply of resources (because the number of resources in the hub can be negative), requests to the spokes can always be fulfilled in the relaxed system, which implies that the number of resources in each spoke in the relaxed system remains no smaller than that in the original system. Figure 3 illustrates the stochastic dominance relationship.

Lemma 4.4 implies that the depletion probability $\mathbb{P}\big[X_0(\delta) = 0\big]$ of the original system is bounded from above by the depletion probability $\mathbb{P}\big[\tilde{X}_0(\delta) \le 0\big]$ of the relaxed system.

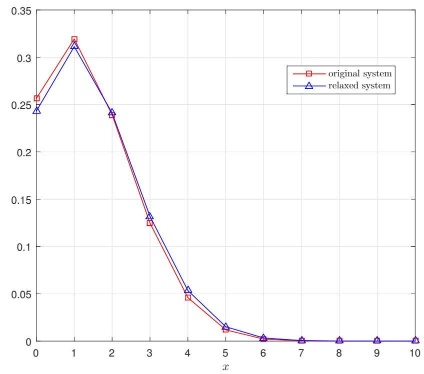

(a) A spoke - density form

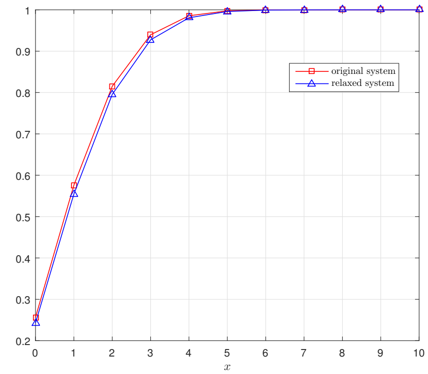

(b) A spoke - cumulative form

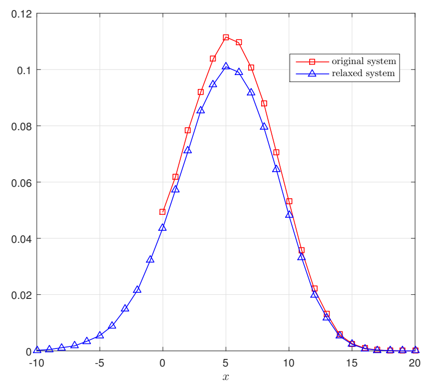

(c) The hub - density form

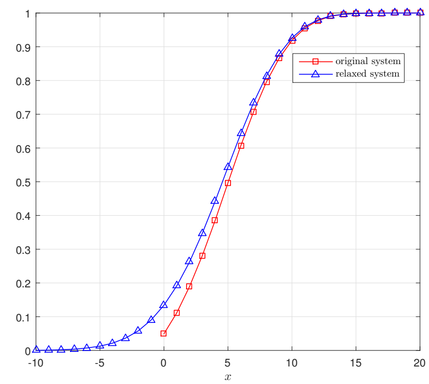

(d) The hub - cumulative form

Figure 3: The stationary distribution of the number of resources in a spoke and the hub under the Lagrangian policy in the original and relaxed systems, respectively. We set $n = 10$, $m = 20$, $q_{i0} = q_{0i} = 0.05$ for all spokes $i \in [n]$, and use the dual variable $\lambda = 0.003$. All private values follow standard uniform distributions with support $[0, 1]$. The dual variable $\lambda = 0.003$ is optimal to the perturbed Lagrangian relaxation (11) with $\delta = \sqrt{n \ln n} = 4.80$; thus δ is the mean number of resources in the hub of the relaxed problem.

Corollary 4.5. *The depletion probability* $\mathbb{P}\big[X_0(\delta) = 0\big]$ *of the original system is bounded from above by the depletion probability* $\mathbb{P}\big[\tilde{X}_0(\delta) \le 0\big]$ *of the relaxed system, i.e.,*

$$\mathbb{P}\Big[X_0(\delta) = 0\Big] \le \mathbb{P}\Big[\tilde{X}_0(\delta) \le 0\Big].$$

Therefore, in order to control the probability that the hub runs out of resources in the original system, it suffices to control the probability that the hub has a non-positive amount of resources in the relaxed system. The latter is easier to analyze because the number of resources in each spoke are independent in the relaxed system (by Corollary B.10). In the next section we bound this probability from above if the expected number resources in each spoke is uniformly bounded.

4.3 Bounding the Hub Depletion Probability

In this section we show that if the expected number of resources $\mathbb{E}\big[\tilde{X}_i(\delta)\big]$ in each spoke of the relaxed system is uniformly bounded by some constant (Assumption 4.1), then with a particular choice of the parameter δ, the depletion probability $\mathbb{P}\big[\tilde{X}_0(\delta) \leq 0\big]$ of the relaxed system shrinks to zero as the number of spokes n increases and the ratio $\frac{m}{n}$ remains fixed. Thus from Theorem 4.1 and Corollary 4.5, the Lagrangian policy is asymptotically optimal in this regime. We provide sufficient conditions for Assumption 4.1 to hold in Section 4.4.

Assumption 4.1. $\mathbb{E}\big[\tilde{X}_i(\delta)\big] \leq c$ for all spokes $i \in [n]$ and $\delta \geq 0$, and some constant $c > 0$ that is independent of n.

Proposition 4.6. *Suppose Assumption 4.1 holds. Letting $b = \frac{1}{1+c}$, the depletion probability of the relaxed system satisfies*

$$\mathbb{P}\Big[\tilde{X}_0(\delta) \leq 0\Big] \leq \exp\left(-\frac{b}{2} \cdot \frac{\delta^2}{m+n}\right).$$

We prove Proposition 4.6 in Appendix A.10 by developing a concentration inequality for a sum of independent random variables with discrete log-concave distributions and uniformly bounded means. The random variables $\tilde{X}_i(\delta)$ for $i \in [n]$ are independent and log-concave by Proposition B.9 and Corollary B.10; their mean values are uniformly bounded by assumption. Intuitively, the variance of $\tilde{X}_0(\delta)$ is linear in n because $\tilde{X}_0(\delta) = m - \sum_{i\in[n]} \tilde{X}_i(\delta)$ and the variances of $\tilde{X}_i(\delta)$ for $i \in [n]$ are uniformly bounded (this follows from the means being uniformly bounded and log-concavity). From the discussion in Section 3.4, the mean value of $\tilde{X}_0(\delta)$ is at least δ. Thus if δ grows faster than the standard deviation of $\tilde{X}_0(\delta)$, which is of order $\sqrt{n}$, the probability that the hub has a non-positive amount of resources in the relaxed system goes to zero. Although this intuition is based on a Chebyshev-like analysis, we can in fact obtain an exponential rate of convergence by bounding the moment generating functions of $\tilde{X}_i(\delta)$ by a geometric distribution with the same mean.

Putting Theorem 4.1, Corollary 4.5, and Proposition 4.6 together, we obtain the following result.

Corollary 4.7. *Under Assumption 4.1 and letting $b = \frac{1}{1+c}$, the Lagrangian policy $\pi(\delta)$ with $0 \leq \delta < m$ satisfies*

$$V^{\pi}(\delta) \leq V^{\mathrm{OPT}} \leq V^{\mathrm{R}} \leq V^{\pi}(\delta) + \bar{r} \cdot \frac{\delta}{m-\delta} + (\bar{r} + \bar{\omega}) \cdot \exp\left(-\frac{b}{2} \cdot \frac{\delta^2}{m+n}\right).$$

In particular, if we set $\delta = \sqrt{\frac{1}{b} \cdot (m+n) \cdot \ln n}$*, we have*

$$V^{\mathrm{OPT}} - V^{\pi}(\delta) \le O\left(\sqrt{\frac{\ln n}{n}}\right)$$

when m *and* n *grow at the same rate.*

4.4 Sufficient Conditions for Uniformly Bounded Spoke Resources

In this section, we provide sufficient conditions for Assumption 4.1 to hold. Intuitively, when spokes are not too different from each other, resources would tend to distributed evenly across the network and the expected number of resources in each spoke would be uniformly bounded. In the extreme case when all spokes are identical, since the total number of resources at the spokes in the relaxed system is no larger than $m - \delta$, we have $\mathbb{E}\big[\tilde{X}_i(\delta)\big] \le \frac{m}{n}$ for all $i \in [n]$ by symmetry. A similar reasoning applies to a more general high multiplicity model in which spokes are partitioned into a fixed number of types and the number of spokes of each type s is a fixed proportion α_s of n; we can take $c = \frac{m}{\underline{\alpha} \cdot n}$ with $\underline{\alpha} = \min_s \alpha_s$ being the smallest proportion. Lemma 4.8 further generalizes the high multiplicity assumption and shows that under some regularity conditions on the problem primitives, Assumption 4.1 holds; in particular, these conditions imply that our theoretical guarantees hold even when all spokes are different from each other.

Lemma 4.8. *Suppose that for the single hub case, in addition to Assumption 2.1, that*

1. *the revenue functions* $r_{ij}(d)$ *are twice continuously differentiable and* $r_{ij}(1) \ge 0$*;*
2. *the revenue functions* $r_{ij}(d)$ *are strongly concave with some parameter* $u_{ij} > 0$ *and have a Lipschitz continuous gradient with some parameter* $U_{ij} > 0$*, i.e.,* $u_{ij} \le -r''_{ij}(d) \le U_{ij}$ *for all* $d \in [0,1]$*. Moreover,* $u_{ij} \ge \bar{u}$ *and* $U_{ij} \le \bar{U}$ *for some positive constants* $\bar{u}$ *and* $\bar{U}$*;*
3. *the arrival rates* $q_{i0}, q_{0i} \in \big[\frac{\underline{q}}{n}, \frac{\bar{q}}{n}\big]$ *for some positive constants* $\underline{q}$ *and* $\bar{q}$ *and all spokes* $i \in [n]$*.*

Then Assumption 4.1 holds.

We prove Lemma 4.8 in Appendix A.11. Typical private value distributions that lead to revenue functions satisfying Assumption 2.1 and Lemma 4.8 include uniform distributions, truncated exponential distributions and truncated normal distributions with support $[a, b]$ and $0 \le a \le b < \infty$, and truncated log-normal distributions with support $[a, b]$ and $0 < a \le b < \infty$.

5 Comparing to Static Pricing Policies

In this section we compare the Lagrangian policy to static pricing policies. A static pricing policy specifies a demand level d_{ij} for each origin-destination pair (i, j) and selects d_{ij} whenever the request type is (i, j) and the number of resources at location i is positive. From Banerjee et al. (2016) we can solve a fluid relaxation problem to get a set of static prices. For an arbitrary network with n locations, the fluid relaxation is (13)

$$
\begin{aligned}
V^{\mathrm{F}} = \max_{d_{ij} \in [0,1]} \quad & \sum_{i=1}^{n} \sum_{j=1}^{n} q_{ij} \cdot r_{ij}(d_{ij}) \\
\text{s.t.} \quad & \sum_{j=1}^{n} q_{ji} \cdot d_{ji} = \sum_{j=1}^{n} q_{ij} \cdot d_{ij}, \ \forall\, i \in [n].
\end{aligned}
\tag{13}
$$

The optimal value of (13) provides an upper bound on V^{OPT} and we denote the static policy by π^{F} and its performance by $V(\pi^{\mathrm{F}})$. We have $V(\pi^{\mathrm{F}}) = \frac{m}{m+n-1} V^{\mathrm{F}}$ from Whitt (1984) (especially Equation (13) therein). Banerjee et al. (2016) also show that the fluid policy can be quite different from the optimal static policy (Appendix E therein): $V(\pi^{\mathrm{F}})$ can be arbitrarily close to $\frac{m}{m+n-1}$ of the performance of the optimal static policy.

The reason that fluid-based static policies may perform worse than other static policies is that the flow balance constraint in (13) does not incorporate the probability that a location is empty. For example, for a hub-and-spoke network, the "exact" flow balance constraint for a general static pricing policy is:

$$
\mathbb{P}[\text{location } i \text{ not empty}] \cdot q_{i0} \cdot d_{i0} = \mathbb{P}[\text{location } 0 \text{ not empty}] \cdot q_{0i} \cdot d_{0i},
$$

whereas the flow balance equation in (13) is simply $q_{i0} \cdot d_{i0} = q_{0i} \cdot d_{0i}$ for each spoke i. In the large supply regime, since $\lim_{n\to\infty} \mathbb{P}[\text{location } i \text{ not empty}] = 1$ for any location $i \in [0 : n]$, this difference is inconsequential. However, in the large network regime we consider, the depletion probability at a node can be strictly positive, so it is essential to incorporate these probabilities.

In general, solving for an optimal static pricing policy for a general network seems to be difficult. For a hub-and-spoke network, however, we can use the same Lagrangian method to derive a performance bound for any static policy and characterize the optimal static policy in the large network limit. The construction and analysis are analogous to the previous sections, and we provide more details in Appendix D. Proposition 5.1 compares the performances of the optimal static policy and the fluid policy for a single hub network with symmetric spokes.

Proposition 5.1. *Consider a single hub network with n identical spokes and let $\hat{\gamma}(\beta) = \sum_{i \in [n]} \gamma_i(\beta)$ with $\gamma_i(\beta)$ as defined in (8). The performance of the fluid policy is $V(\pi^{\mathrm{F}}) = \frac{m}{m+n}\hat{\gamma}(1)$ and the performance V^{S} of the optimal static policy converges to $\hat{\gamma}\big(\frac{m}{m+n}\big)$ in the large network limit.*

We prove Proposition 5.1 in Appendix D.2. Since $\hat{\gamma}(\beta)$ is strictly concave by Lemma A.3 and $\hat{\gamma}(0) = 0$ by Lemma A.4, we have $V^{\mathrm{S}} > V(\pi^{\mathrm{F}})$ by Jensen's inequality. Finally, Example 5.1 shows that the optimal static policy is strictly suboptimal in the large network regime.

Example 5.1. Consider a single hub example with n identical spokes and $m = \frac{2}{3}n$ resources. All private values are uniformly distributed on $[0, 1]$ and the arrival rates are $q_{i0} = q_{0i} = \frac{1}{2n}$ for all $i \in [n]$. Since the request rates to and from a spoke are equal, the optimal solution to the fluid relaxation is $d_{i0} = d_{0i} = \frac{1}{2}$ for all $i \in [n]$, the fluid relaxation bound is the trivial bound $V^{\mathrm{F}} = \hat{\gamma}(1) = \frac{1}{4}$ that equals $\bar{r}$, and the performance of the fluid policy is $V(\pi^{\mathrm{F}}) = \frac{m}{m+n}V^{\mathrm{F}} = \frac{1}{10}$. On the other hand, $\hat{\gamma}(\beta) = \frac{1}{2} \cdot \frac{\beta}{1+\beta}$ and hence the performance of the optimal static policy will not exceed $\hat{\gamma}\big(\frac{m}{m+n}\big) = \frac{1}{7} \approx 0.143$ in the large network regime. The optimal static policy converges to $d_{i0} = \frac{m+n}{2m+n} = \frac{5}{7}$ and $d_{0i} = \frac{m}{2m+n} = \frac{2}{7}$ in the limit. The optimal static policy is strictly suboptimal: we can show a simple dynamic policy that limits each spoke to contain at most two resources achieves an asymptotic performance of approximately 0.152. Details for these calculations are provided in Appendix D.3.

6 Extensions

The policies and upper bounds developed in Section 4 can be extended to more general networks with multiple, interconnected hubs. In shared vehicle applications, we can imagine such network structures as capturing several nearby urban areas and other major traffic centers such as an airport (the hubs) that are surrounded by a large number of more distant suburbs (the spokes). Multiple hub systems are also used widely in other industries, such as airlines (e.g., Tran et al. 2017). In Section 6.1, we extend the Lagrangian relaxation to multiple-hubs networks in which resources can relocate between hubs and between a hub and a spoke but not between spokes. We provide performance bounds and prove asymptotic optimality only for the special case of "uniformly related" hubs. We further incorporate spoke-to-spoke connections in Section 6.2. Finally, we discuss how to incorporate nonzero relocation times in the bounds and policies in Section 6.3. While the theoretical analysis for general networks remains an open challenge, in Section 7 we study these approximations on several numerical examples and find that they lead to good performance.

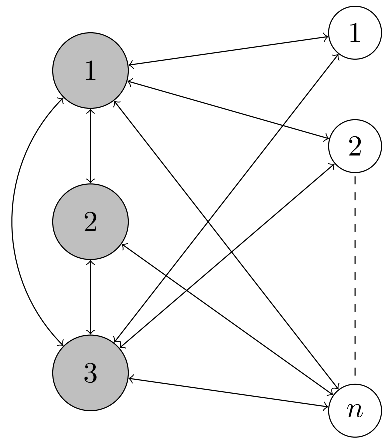

Figure 4: A hub-and-spoke network with 3 hubs (grey) and n spokes (white).

6.1 Multiple Hubs

In this section, we consider network structures with multiple, interconnected hubs, as illustrated in Figure 4. We assume resources can be relocated between hubs and between a hub and a spoke, but not between spokes. In the following, we let J denote the number of hubs and we use variable j to denote a hub and variable i to denote a spoke.

In the Lagrangian relaxation, in addition to dualizing the capacity constraint $\sum_{i\in[n]} x_i \leq m$ using a dual variable $\lambda \geq 0$ and dropping the constraint $d \leq x_j$ when some request (j,i) arrives, we dualize the flow balance constraint for each hub j, i.e., the constraint that, for each hub, the average inflow of resources is equal to the average outflow of resources. Note that flow balance is a "valid" equality in the sense that it is not an explicit constraint of the model, but every feasible policy must satisfy this constraint. In particular, denoting by $\mu_j \in \mathbb{R}$ the Lagrange multiplier for hub j, we introduce a reward μ_j for every resource that moves to hub j and a penalty $-\mu_j$ for every resource that leaves hub j.

Following this reasoning, for any $\lambda \geq 0$ and $\boldsymbol{\mu} = (\mu_j)_{j\in[J]} \in \mathbb{R}^J$, the Lagrangian relaxation provides an upper bound on the average revenue V^{OPT} of an optimal policy, which we denote by $\bar{V}^{\lambda,\boldsymbol{\mu}}$. Moreover, the Lagrangian relaxation decomposes over spokes with

$$\bar{V}^{\lambda,\boldsymbol{\mu}} = m\lambda + \sum_{i=1}^{n} h_i^{\lambda,\boldsymbol{\mu}} + \sum_{j,j'\in J} q_{jj'} \cdot g_{jj'}^{\boldsymbol{\mu}}, \tag{14}$$

where $g_{jj'}^{\boldsymbol{\mu}} \triangleq \max_{d\in[0,1]} \left\{ r_{jj'}(d) + d \cdot \left(\mu_{j'} - \mu_j\right) \right\}$ denotes the average revenue earned from a hub-to-hub request (j,j'), and $h_i^{\lambda,\boldsymbol{\mu}}$ denotes the average revenue of an optimal policy to each spoke i

problem. Let $d^{\boldsymbol{\mu}}_{jj'} \triangleq \operatorname{argmax}_{d\in[0,1]}\big\{r_{jj'}(d) + d\cdot\big(\mu_{j'} - \mu_j\big)\big\}$ denote the optimal demand value for the type (j,j') requests; since we relax the capacity constraints of each hub, $d^{\boldsymbol{\mu}}_{jj'}$ is independent of the resource levels of hubs j and j'. Finally, analogously to Proposition 3.2, $h_i^{\lambda,\boldsymbol{\mu}}$ is equal to the optimal value of the following optimization problem:

$$
\begin{aligned}
h_i^{\lambda,\boldsymbol{\mu}} = \max_{\substack{d_i(x,i,j)\in[0,1],\\ d_i(x,j,i)\in[0,1],\\ p_i(x)\geq 0}} \quad & \sum_{x=0}^{m} p_i(x)\cdot\sum_{j=1}^{J}\Big\{q_{ij}\cdot r_{ij}\Big(d_i(x,i,j)\Big) + q_{ji}\cdot r_{ji}\Big(d_i(x,j,i)\Big)\Big\} \\
& + \sum_{j=1}^{J}\mu_j\sum_{x=0}^{m}p_i(x)\cdot\Big\{q_{ij}\cdot d_i(x,i,j) - q_{ji}\cdot d_i(x,j,i)\Big\} - \lambda\cdot\sum_{x=0}^{m}x\cdot p_i(x) \\
\text{s.t.} \quad & \sum_{x=0}^{m}p_i(x) = 1, \qquad (15) \\
& p_i(x)\cdot\sum_{j=1}^{J}q_{ji}\cdot d_i(x,j,i) = p_i(x+1)\cdot\sum_{j=1}^{J}q_{ij}\cdot d_i(x+1,i,j),\ \forall\, x\in[0:m-1], \\
& d_i(0,i,j)=0,\ d_i(m,j,i)=0,\ \forall\, j\in[J].
\end{aligned}
$$

Compared to (6), the objective incorporates a penalty for the violation of the flow balance constraint of each hub. We can solve (15) in a similar manner to (6) using the algorithms developed in Lemma A.7. The Lagrangian dual problem

$$V^{\mathrm{R}} \triangleq \min_{\lambda\geq 0,\boldsymbol{\mu}\in\mathbb{R}^J} \bar{V}^{\lambda,\boldsymbol{\mu}}$$

provides the tightest possible Lagrangian relaxation upper bound and is equivalent to the problem of maximizing the average revenue subject to the constraints that the expected number of resources in the hubs is non-negative and all hubs are flow-balanced. The latter follows because the first-order condition with respect to μ_j implies that $\sum_{i=1}^{n}\sum_{x=0}^{m}p_i(x)\cdot\Big(q_{ij}\cdot d_i(x,i,j) - q_{ji}\cdot d_i(x,j,i)\Big) + \sum_{j'\in[J]}\big(q_{j'j}d_{j'j} - q_{jj'}d_{jj'}\big) = 0$, i.e., hub j is flow-balanced in expectation.

Finally, we can construct Lagrangian policies that attempt to leave $\delta = o(n)$ resources on average to the hubs, by solving a perturbed problem

$$V^{\mathrm{R}}(\delta) = \min_{\lambda\geq 0,\boldsymbol{\mu}\in\mathbb{R}^J} \bar{V}^{\lambda,\boldsymbol{\mu}} - \delta\lambda. \qquad (16)$$

Let $\lambda^*(\delta)$ and $\boldsymbol{\mu}^*(\delta)$ be an optimal solution to (16). We let our policy $\pi(\delta)$ be the adaptive control from solving (15) with the dual variables $\lambda^*(\delta)$ and $\boldsymbol{\mu}^*(\delta)$ for requests (i,j) or (j,i) between a hub

j and a spoke i, and the static control $d_{jj'}^{\boldsymbol{\mu}^*(\delta)}$ for requests of type (j, j') between hubs. We denote the performance of $\pi(\delta)$ in the original problem by $V^\pi(\delta)$. Analogous to Theorem 4.1, we can bound the performance gap by a term proportional to the depletion probabilities of the hubs in the original problem, plus a perturbation term, as shown in Theorem 6.1.

Theorem 6.1. *For any $\lambda^*(\delta)$ and $\boldsymbol{\mu}^*(\delta)$ optimal to (16), the corresponding Lagrangian policy $\pi(\delta)$ satisfies*

$$V^\pi(\delta) \le V^{\mathrm{OPT}} \le V^{\mathrm{R}} \le V^\pi(\delta) + \bar{r} \cdot \frac{\delta}{m-\delta} + (\bar{r} + \bar{\omega}) \cdot \sum_{j \in [J]} q_j \cdot \mathbb{P}\Big[X_j(\delta) = 0\Big],$$

where $\mathbb{P}\big[X_j(\delta) = 0\big]$ is the stationary probability that hub j runs out of resources in the original problem under the policy $\pi(\delta)$, and $q_j = \sum_{i \in [n]} q_{ji} + \sum_{j' \in [J]} q_{jj'}$ is the probability that hub j is the originating location of the request.

We prove Theorem 6.1 in Appendix A.12. Theorem 6.1 implies that the policy $\pi(\delta)$ only loses a small performance if the depletion probabilities of each hub are small. We conjecture that with proper choice of the parameter δ, the depletion probability of each hub decreases to zero and hence the Lagrangian policy is asymptotically optimal when the number of spokes n and resources m grow at the same rate, and the hubs and their interconnections remain fixed. The challenge in proving this conjecture is that the Lagrangian relaxation does not explicitly control the number of resources in *each* of the hubs. Corollary 4.5 and Proposition 4.6 imply that the probability that the sum of resources in the hubs being zero diminishes at an exponential rate; to show the depletion probability of each hub diminishes, we need to analyze the joint distribution of resources across hubs. In Appendix E, we fully characterize the joint distribution across the hubs for a special case that we refer to as *uniformly related hubs*, where the ratio of the arrival rate to a spoke from a hub to the arrival rate of the reverse trip is constant across hubs, and myopic pricing for hub-to-hub relocations yields balanced flow within hubs. We show that the stationary distribution across hubs is uniform with uniformly related hubs. This implies that the depletion probability for each hub is small and as a result, the policy $\pi(\delta)$ is asymptotically optimal when the number of spokes n and resources m grow at the same rate and the number of hubs is $o(n)$.

6.2 Incorporating Spoke-to-Spoke Connections

In this section we further extend the method to handle spoke-to-spoke connections. Although the approach applies to arbitrary networks, we expect this approach will work well when spoke-to-spoke requests constitute a small fraction of the total requests. Without loss of generality, we assume that no request relocates a resource within a spoke, i.e., $q_{ii} = 0$ for all $i \in [n]$, because the service provider can simply choose d^*_{ii} to maximize the immediate revenue when these requests arrive.

In the Lagrangian relaxation, as before, we relax the capacity constraint $\sum_{i=1}^n x_i \leq m$ with a dual variable $\lambda \geq 0$ and drop the constraint $d \leq x_j$ when some request (j, i) arrives, and relax the flow balance constraint for each hub j with a dual variable $\mu_j \in \mathbb{R}$ as before. We further relax the relocation constraint for each spoke-to-spoke request (i, i'), i.e., the constraint $x_{i',t+1} = x_{i't} + 1$ if a request (i, i') arrives at time t and is fulfilled, and $x_{i',t+1} = x_{i't}$ if not fulfilled, using the same dual variable $\nu_{ii'} \in \mathbb{R}$. This is equivalent to a model in which: (a) when a request (i, i') arrives and is fulfilled, the provider receives an additional reward $\nu_{ii'}$ and one resource exits the system from spoke i; and (b) the provider has the option to add one resource at the destination at a price $\nu_{ii'}$. Since every feasible policy to the original problem is feasible to the Lagrangian relaxation and attains an objective value that is no smaller, the Lagrangian relaxation provides a valid upper bound for all dual variables $\lambda \geq 0$, $\boldsymbol{\mu} \in \mathbb{R}^J$ and $\boldsymbol{\nu} \in \mathbb{R}^{n\times n}$, which we denote by $\bar{V}^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}}$.

Similar to the previous section, this Lagrangian relaxation also decomposes over spokes and hubs with the form

$$\bar{V}^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}} = m\lambda + \sum_{i=1}^n h_i^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}} + \sum_{j,j' \in [J]} q_{jj'} \cdot g_{jj'}^{\boldsymbol{\mu}},$$

where $g_{jj'}^{\boldsymbol{\mu}} \triangleq \max_{d\in[0,1]} \left\{ r_{jj'}(d) + d \cdot (\mu_{j'} - \mu_j) \right\}$ denotes the average revenue earned from a hub-to-hub request (j, j') as in (14), and $h_i^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}}$ denotes the average revenue of an optimal policy to each spoke i problem; this optimal average revenue may be calculated by solving an optimization problem similar to (15) that includes additional decision variables $d_i(x, i, i')$ that denote the demand values to use when a spoke-to-spoke request (i, i') arrives and spoke i has x resources, and decision variables $d_i(x, i', i)$ that denote the probability that the provider will add one resource in location i when a request (i', i) arrives with x resources in spoke i. The objective function further incorporates a penalty for the violation of the relocation constraint of each spoke-to-spoke connection. We define this problem formally and justify the above decomposition in Proposition B.12.

We can obtain the tightest possible Lagrangian relaxation upper bound by solving the Lagrangian dual problem

$$V^{\mathrm{R}} \triangleq \min_{\lambda \geq 0, \boldsymbol{\mu} \in \mathbb{R}^J, \boldsymbol{\nu} \in \mathbb{R}^{n\times n}} \bar{V}^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}}.$$

This problem is equivalent to maximizing the average revenue subject to the constraints that the expected number of resources in the hubs is non-negative, the in-flow and out-flow of each hub j is balanced in expectation, and for each spoke-to-spoke connection (i,i'), the out-flow of spoke i via requests (i,i') is equal to the in-flow of spoke i' via requests (i,i') in expectation. Proposition 6.2 shows that the Lagrangian relaxation provides tighter bounds than the fluid relaxation bound.

Proposition 6.2. *The fluid relaxation bound V^{F} in (13) is weaker than the Lagrangian relaxation bound $\min_{\boldsymbol{\mu},\boldsymbol{\nu}} \bar{V}^{\lambda=0,\boldsymbol{\mu},\boldsymbol{\nu}}$ where the dual variables $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ are optimal when $\lambda = 0$ is fixed, i.e.,*

$$V^{\mathrm{R}} \le \min_{\boldsymbol{\mu},\boldsymbol{\nu}} \bar{V}^{\lambda=0,\boldsymbol{\mu},\boldsymbol{\nu}} \le V^{\mathrm{F}}.$$

Thus, for a general network, regardless of which nodes we take to be hubs, this approach is guaranteed to provide tighter bounds than the fluid relaxation. Note also that the fluid relaxation corresponds to the case when all nodes are hubs. We prove Proposition 6.2 in Appendix B. Finally, we can construct Lagrangian policies that attempt to leave $\delta = o(n)$ resources on average in the hubs, by solving a perturbed problem

$$V^{\mathrm{R}}(\delta) = \min_{\lambda \ge 0, \boldsymbol{\mu} \in \mathbb{R}^J, \boldsymbol{\nu} \in \mathbb{R}^{n \times n}} \bar{V}^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}} - \delta\lambda. \tag{17}$$

Let $\lambda^*(\delta)$, $\boldsymbol{\mu}^*(\delta)$ and $\boldsymbol{\nu}^*(\delta)$ be an optimal solution to (17). We let our policy $\pi(\delta)$ be the adaptive control $d_i(x,i,j)$, $d_i(x,j,i)$ and $d_i(x,i,i')$ from solving the optimization problem for each spoke (given in (67) in Appendix B) with dual variables $\lambda^*(\delta)$, $\boldsymbol{\mu}^*(\delta)$ and $\boldsymbol{\nu}^*(\delta)$, for requests (i,j) or (j,i) between a hub j and a spoke i and requests (i,i') between spokes, and the static control $d_{jj'}^{\boldsymbol{\mu}^*(\delta)}$ for requests of type (j,j') between hubs.

6.3 Incorporating Relocation Times

In this section we further incorporate nonzero relocation times. We assume relocation times of resources on (i,j) are *i.i.d.* and we denote these relocation times by a positive random variable Γ_{ij} and denote their mean values by $\tau_{ij} \triangleq \mathbb{E}[\Gamma_{ij}]$. We again consider the embedded discrete time model where each period corresponds to a new request. With general relocation times, we need to track the number of resources in each location as well as the resources in transit and the relocation times they have already spent relocating.

We consider the same relaxations as in Section 6.2. First, we relax the capacity constraint that the number of resources in the hubs is non-negative with a dual variable $\lambda \ge 0$ and drop the constraint $d \le x_j$ when some request leaves the hub j: this involves that at the beginning of each period, we receive a reward $m\lambda$ but penalize every resource that is either in a spoke or in transit with a cost λ. We also relax the flow balance constraint of each hub j with a dual variable $\mu_j \in \mathbb{R}$, thus every resource that enters the hub j receives a reward μ_j and every resource that leaves the hub j incurs a penalty $-\mu_j$. Since we consider an infinite horizon setting, we assume without loss that all rewards are collected at the beginning of the relocation. Finally, we relax the relocation constraint for each spoke-to-spoke request (i, i') that involves two different spokes with a dual variable $\nu_{ii'} \in \mathbb{R}$. In effect, when a request (i, i') arrives and is fulfilled, we receive a reward $\nu_{ii'}$ and one resource exits the system from spoke i. We can add one resource at the destination at a cost $\nu_{ii'}$ and a waiting time $\Gamma_{ii'}$ regardless of the fulfillment. Since every feasible policy to the original problem is feasible to the Lagrangian problem and attains an objective value that is no smaller, the Lagrangian relaxation provides a valid upper bound $\bar{V}^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}}$ for any dual variables $\lambda \geq 0$, $\mu_j \in \mathbb{R}$ and $\nu_{ii'} \in \mathbb{R}$. Moreover, analogous to Proposition B.12, the relaxation still decomposes over spokes.

In Appendix F, we characterize the spoke problem when there is one hub and all relocation times follow exponential distributions. In this case, the resulting spoke problem has two state variables, which leads to a tractable optimization problem.

With general relocation times, the spoke problem is difficult to solve, because we need to track the relocation times of all the resources in transit. We further relax each spoke problem to provide a tractable upper bound by considering a relaxation that enables resources that are moving towards the spoke to be instantaneously available at the spoke. Since a resource incurs a penalty λ per period whether it is at the spoke or moving to it, it is always better to keep the resources at the spoke as this grants the decision maker flexibility to serve more requests. Thus, after the relaxation, we only need to track the number of resources in the spoke because resources that are moving out of the spoke do not need to be tracked[6]. We formalize this relaxation in Proposition B.13. We note that this relaxation may not provide asymptotically tight upper bounds, but this approach can provide high-quality bounds and policies in practice (see the examples in Section 7.2).

7 Numerical Examples

In this section, we examine the performance of the Lagrangian policy and the Lagrangian relaxation upper bound on some numerical examples. We consider two examples: a synthetic example with a single hub and a more realistic example based on data from RideAustin. In Appendix G, we show

[6] As in Banerjee et al. (2016) and Braverman et al. (2016), we can use Little's law for the expected number of resources moving out of a spoke to obtain that every resource leaving the spoke to a hub j incurs a penalty $\lambda\Lambda\tau_{ij}$ in expectation because relocation (i, j) takes $\Lambda \cdot \tau_{ij}$ periods on average.

the stationary distributions of the resources in the hub under various policies for the single hub example, and we provide another synthetic example with multiple hubs.

7.1 Single Hub Examples

We first consider examples with one hub. The number of spokes n increases linearly from 100 to 1000 with step size 100 and we assume that all spokes are identical, the arrival rates are $q_{i0} = q_{0i} = \frac{1}{2n}$ for all spokes $i \in [n]$, the number of resources is $m = 2n$, and the private values for all request types are uniformly distributed in $[0, 1]$.

For each fixed n, we calculate: (a) the performance of the policy $V^{\pi}(\delta)$ with $\delta = \sqrt{n \ln n}$; (b) the Lagrangian relaxation upper bound V^{R}; (c) the fluid relaxation bound V^{F}; and (d) the performance of the fluid-based static policy $V(\pi^{\mathrm{F}})$. We additionally consider a static pricing policy π^{S} with $d_{0i} = \frac{\beta}{1+\beta}$ and $d_{i0} = \frac{1}{1+\beta}$, using $\beta = \frac{\rho}{1+\rho}$ with $\rho = \frac{m}{n} - \frac{\delta}{n} = 2 - \sqrt{\frac{\ln n}{n}}$; based on the analysis in Appendices D.1 and D.2, π^{S} converges to the optimal static pricing policy in the large network regime. For each fixed n, we additional calculate: (e) the performance $V(\pi^{\mathrm{S}})$ of the static policy $V(\pi^{\mathrm{S}})$. We estimate the values $V^{\pi}(\delta)$, $V(\pi^{\mathrm{F}})$, and $V(\pi^{\mathrm{S}})$ with 100 sample paths and for each sample path we approximate the average revenue with a time average of the total revenue of the first $4000n$ time periods; this led to very low standard errors in the results (see (b) of Figure 5 for a sense of this).

Figure 5 shows the simulation results for the single-hub case. From Figure 5, the fluid relaxation bound is quite weak and the fluid-based static policy does not appear to converge to optimality. In fact, since the request rates from and to a location are equal for all locations, the optimal solution to (13) is $d_{ij} = d^*_{ij} = \frac{1}{2}$ for all request types and the fluid relaxation upper bound is the trivial bound $V^{\mathrm{F}} = \bar{r} = \frac{1}{4}$. The Lagrangian policy, however, performs very well and the gap between the Lagrangian relaxation bound and the performance of the Lagrangian policy clearly diminishes as n grows.

The policy π^{S} is also strictly suboptimal. Analogous to Example 5.1, the performance of any static policy is no larger than $\frac{1}{2} \cdot \frac{m}{2m+n} = \frac{1}{5}$ in the large network regime, and the performance $V(\pi^{\mathrm{S}})$ converges to this value as n increases. Note that the optimal static pricing policy is quite different from the fluid policy not only in terms of performance, but also in terms of controls. Intuitively, the optimal static policy will pool resources at the hub, but a dynamic pricing policy does so more efficiently.

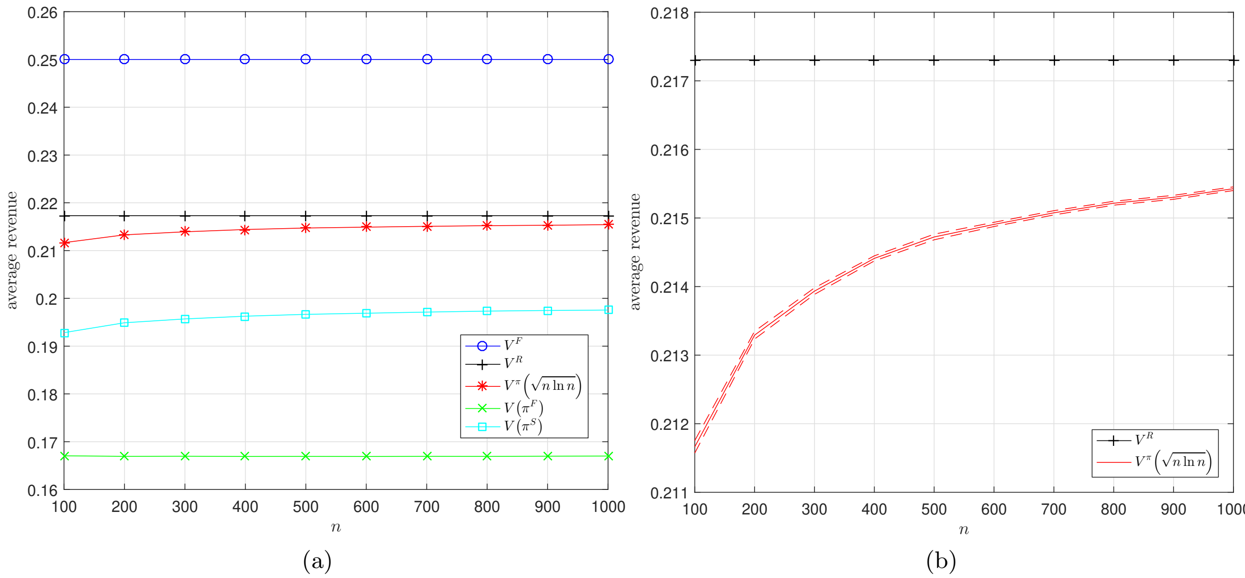

Figure 5: Simulation results of the single-hub case. (b) is simply a magnified version of (a) highlighting the performance of our policy and the Lagrangian relaxation upper bound. A 95% confidence interval around $V^\pi\left(\sqrt{n\ln n}\right)$ is plotted with dashed lines in (b).

7.2 RideAustin Example

In this section, we evaluate the policies and bounds using a model calibrated on empirical data from RideAustin (RideAustin 2017). The data set covers about 1.5 million transactions over a course of 10 months. Each transaction provides detailed information on the ride, including the longitude/latitude coordinates of the origin and destination points, the starting and end time, the unique driver ID, and the total fare. The objective of this analysis is to demonstrate the value that may be captured in practice by employing the Lagrangian method.

We first partition the city by clustering the origin and destination points into 100 clusters obtained from solving the k-center problem with the first few centers initialized with k-means clustering centers. The Voronoi diagram of the cluster centers results in a partition with $n = 100$ locations; we show the Voronoi diagram and the ride flow based on the partition in Figure 11 in Online Appendix G.3. To estimate model primitives, we assume a stationary demand arrival rate, and deterministic relocation times and log-normal private values for each route (i, j). Since the data only provides prices for requests that are fulfilled, we use prices as a rough approximation of the consumers' private values. We estimate the private value distribution parameters and the relocation probabilities q_{ij} using the full data set, and estimate the number of drivers m by restricting to data from January 2017 to February 2017, during which the system dynamics are quite stable. We approximate m with the number of active drivers in a 3-hour period and we choose $m = 400$, which is roughly the average number of drivers from 8am to 4pm.

To select a good set of hubs, we vary the total number of hubs J from 1 to 6, and for each value, we solve an integer linear program (18) to determine the specific locations to choose as hubs.

$$
\begin{aligned}
\min_{x_{ij}, y_i \in \{0,1\}} \quad & \sum_{i,j \in [n]: i \neq j} \left(1 - x_{ij}\right) \cdot q_{ij} \\
\text{s.t.} \quad & \sum_{i=1}^{n} y_i \leq J, \\
& x_{ij} \leq y_i + y_j, \ \forall\, i, j \in [n], \\
& \sum_{j: q_{ij} > 0} x_{ij} \geq 1, \ \sum_{j: q_{ji} > 0} x_{ji} \geq 1, \ \forall\, i \in [n].
\end{aligned}
\tag{18}
$$

In (18), $y_i = 1$ if location i is a hub, and $x_{ij} = 1$ if request (i, j) is between hubs or between a hub and a spoke. We minimize the sum of relocation probabilities of requests between two distinct spokes to ensure the hubs cover most of the relocations. The first constraint enforces that at most J hubs are selected. The second constraint imposes that $x_{ij} = 1$ only if either the origin or the destination is a hub. The last constraint enforces that each spoke is connected to a hub. Figure 12 in Online Appendix G.3 illustrates the locations of hubs obtained from solving (18) with different values of J.

For each fixed J, we calculate: (a) the Lagrangian relaxation upper bound V^{R}, (b) the performance (average revenue per request) of the Lagrangian policy $V^{\pi}(\delta)$ with the optimal choice of δ (which we denote by δ^*) and $\delta = 0$ respectively, (c) the performance of a static policy $\pi^{\mathrm{S}}(\delta)$ with the optimal choice of δ (which we denote by δ^*_{S}). Optimal choices of δ for the dynamic policies are roughly 160, 140, 140, 140, 140, and 160 for J from one to six; and of δ for the static policies are roughly 20, 20, 40, 40, 40, and 60 for J from one to six. In Figure 13 in Online Appendix G.3, we show how performance varies with δ for each value of J. The dynamic Lagrangian policy $\pi(\delta)$ incorporates multiple hubs, spoke-to-spoke transitions, and travel times as discussed in Section 6. The static policy $\pi^{\mathrm{S}}(\delta)$ is computed using a perturbed version of the Lagrangian relaxation described in Section 5. Namely, we solve the a perturbed Lagrangian relaxation with some δ like for dynamic pricing except that we enforce static controls in each spoke problem. Additionally, we compute the fluid relaxation bound V^{F}, and the performance of the fluid-based static policy $V(\pi^{\mathrm{F}})$. These quantities are independent of the number of hubs J. In the fluid relaxation (13), we additionally add the constraint that the number of resources in transit is no larger than m using

J	1	2	3	4	5	6
V^{R} (Lagrangian upper bound)	7.528	7.531	7.533	7.536	7.542	7.546
$V^{\pi}(\delta^*)$ (dynamic policy)	7.161	7.135	7.114	7.095	7.078	7.059
$V^{\pi}(0)$ (dynamic policy)	6.892	6.881	6.883	6.845	6.842	6.823
$V(\pi^{\mathrm{S}}(\delta^*_{\mathrm{S}}))$ (static policy)	6.915	6.896	6.885	6.892	6.881	6.873
$\frac{V^{\mathrm{R}}-V^{\pi}(\delta^*)}{V^{\pi}(\delta^*)}$ (dynamic gap)	5.13%	5.55%	5.88%	6.23%	6.55%	6.90%
$\frac{V^{\mathrm{R}}-V(\pi^{\mathrm{S}}(\delta^*_{\mathrm{S}}))}{V(\pi^{\mathrm{S}}(\delta^*_{\mathrm{S}}))}$ (static gap)	8.87%	9.21%	9.42%	9.35%	9.60%	9.80%

Fluid Relaxations	
V^{F} (fluid upper bound)	7.692
$V(\pi^{\mathrm{F}})$ (fluid policy)	6.046
$\frac{V^{\mathrm{F}}-V(\pi^{\mathrm{F}})}{V(\pi^{\mathrm{F}})}$ (fluid gap)	27.22%

Table 1: Performance bounds and mean policy performances for the RideAustin example.

Little's law, i.e., $\sum_{i=1}^{n}\sum_{j=1}^{n}\eta_{ij}d_{ij}\tau_{ij} \leq m$, as in Algorithm 5 of Banerjee et al. (2016).

Table 1 shows the simulation results. In general, there is a trade-off in choosing the optimal number of hubs: a small number of hubs retains the benefits of dynamic pricing at the spokes, while a large number of hubs guarantees that spoke-to-spoke flow is minimal. From the ride flow of the RideAustin (Figure 12(b) in Online Appendix), the requests at the central location in the Austin downtown dominates requests at other locations greatly, and this perhaps explains why one hub is sufficient to cover a great amount of flow and achieve the best performance. Retaining sufficient drivers in the hubs on average in the relaxation (i.e., $\delta = \delta^*$ for a fixed J) leads to a substantial performance improvement compared to only ensuring that the drivers in the hubs to be non-negative (i.e., $\delta = 0$) in the relaxation. The best dynamic pricing policy leads to a gap of 5.13% compared to the gap of 27.22% for the fluid policy. The gap of the best Lagrangian-based static policy is 8.87%, which substantially improves upon the fluid policy, but still underperforms compared to the (dynamic) Lagrangian policy.

8 Conclusions

We have considered dynamic pricing of resources that relocate over a network, and developed an approximate pricing policy and performance bound based on Lagrangian relaxations. We prove an explicit bound on the suboptimality of this policy that shows the approach is asymptotically optimal for hub-and-spoke networks with one hub or uniformly related hubs and with many locations and resources. We have also shown how to extend the approach to more general networks with multiple hubs and spoke-to-spoke connections and how to incorporate nonzero relocation times. While we have observed strong performance of this approach on numerical examples, ongoing work involves developing theoretical performance guarantees for these more complex systems. Our model could be further refined to include more realistic problem features, such as resources that periodically arrive and depart the system and non-stationary arrival rates and revenue functions. Although these features would potentially complicate a complete performance analysis, we believe the methods developed in this paper would nonetheless work well with such variations of the problem. In general, for policies from Lagrangian relaxations to perform well, we need the policy to behave similarly in the original and relaxed problems and this requires that the constraint that is being relaxed to be satisfied with high probability. To achieve this, we use a perturbed Lagrangian problem to push the system toward the interior and away from the boundary of the constraint we relax. We are optimistic that the perturbed Lagrangian could be useful in some other problems as well.

We are hopeful that supply-constrained large networks provide a suitable model of resource relocation problems in practice. In principle, resource relocation problems can be studied in the large supply regime by adopting network topologies in which nodes cover large geographical areas. These coarse network topologies would guarantee a large supply of resources per location at the expense of treating potentially distant resources as identical and interchangeable. Adopting finer network topologies with many nodes allows us to capture spatial supply and demand more accurately as resources within each node would be geographically closer. This is also consistent with insights from the ride-sharing industry. For example, Uber[7] indicates that "Deriving information and insights from data in the Uber marketplace requires analyzing data across an entire city. Because cities are geographically diverse, this analysis needs to happen at a fine granularity. Analysis at the finest granularity, the exact location where an event happens, is very difficult and expensive. Analysis on areas, such as neighborhoods within a city, is much more practical."

Acknowledgements

We thank Omar Besbes, an anonymous associate editor, and three anonymous reviewers for providing a number of valuable comments that greatly improved the paper.

[7]Source: `https://eng.uber.com/h3/`

A Proofs

A.1 Proof of Proposition 2.1

Let $N(T)$ be the number of requests during time $[0, T]$. Using the notation of Section 2.2, the long-run average number of relocations per resource per unit time Φ^π is

$$\begin{aligned}
\Phi^\pi &= \frac{1}{m} \cdot \lim_{T\to\infty} \mathbb{E}\left\{ \frac{1}{T} \cdot \sum_{t=1}^{N(T)} \sum_{i,j\in[n]} y_{ij,t} \cdot \mathbb{1}\left[\xi_t \le d^\pi_{ij,t}\right] \right\} \\
&\overset{\text{(i)}}{=} \frac{1}{m} \cdot \mathbb{E}\left\{ \lim_{T\to\infty} \frac{1}{T} \cdot \sum_{t=1}^{N(T)} \sum_{i,j\in[n]} y_{ij,t} \cdot \mathbb{1}\left[\xi_t \le d^\pi_{ij,t}\right] \right\} \\
&\overset{\text{(ii)}}{=} \frac{\sum_{i,j\in[n]} \eta_{ij}}{m} \cdot \mathbb{E}\left\{ \lim_{N\to\infty} \frac{1}{N} \cdot \sum_{t=1}^{N} \sum_{i,j\in[n]} y_{ij,t} \cdot \mathbb{1}\left[\xi_t \le d^\pi_{ij,t}\right] \right\} \\
&\overset{\text{(iii)}}{=} \frac{\sum_{i,j\in[n]} \eta_{ij}}{m} \underbrace{\lim_{N\to\infty} \frac{1}{N} \cdot \mathbb{E}\left\{ \sum_{t=1}^{N} \sum_{i,j\in[n]} y_{ij,t} \cdot d^\pi_{ij,t} \right\}}_{\Theta^\pi},
\end{aligned}$$

where (i) follows from the generalized dominated convergence theorem (Theorem 19 in Royden and Fitzpatrick 2010) because $\sum_{i,j\in[n]} y_{ij,t} \cdot \mathbb{1}\left[\xi_t \le d^\pi_{ij,t}\right] \in [0, 1]$ and $\lim_{T\to\infty} \mathbb{E}\left[\frac{N(T)}{T}\right] = \mathbb{E}\left[\lim_{T\to\infty} \frac{N(T)}{T}\right]$, (ii) from $\lim_{T\to\infty} \frac{N(T)}{T} = \sum_{i,j\in[n]} \eta_{ij}$ and $\lim_{T\to\infty} N(T) = \infty$, and (iii) from the bounded convergence theorem. In the last equation, Θ^π represents the expected throughput per period. Using $\Theta^\pi \le 1$ and $\sum_{i,j\in[n]} \eta_{ij} \le \bar{\eta} n$, we have

$$\Phi^\pi \le \bar{\eta} \frac{n}{m}.$$

For the other direction, let V^π denote the long-run average revenue of policy π. We have

$$V^\pi = \lim_{N\to\infty} \frac{1}{N} \cdot \mathbb{E}\left\{ \sum_{t=1}^{N} \sum_{i,j\in[n]} y_{ij,t} \cdot r_{ij}\left(d^\pi_{ij,t}\right) \right\} \le \bar{\omega} \cdot \lim_{N\to\infty} \frac{1}{N} \cdot \mathbb{E}\left\{ \sum_{t=1}^{N} \sum_{i,j\in[n]} y_{ij,t} \cdot d^\pi_{ij,t} \right\} = \bar{\omega} \cdot \Theta^\pi,$$

where the inequality is due to the mean value theorem and the facts that $r_{ij}(0) = 0$ and $\bar{\omega} > 0$ is a uniform bound on the derivatives of the one-period revenue functions by Assumption 2.1. Using the above inequality on V^π and $\sum_{i,j\in[n]} \eta_{ij} \ge \underline{\eta} n$, we obtain

$$\Phi^\pi \ge \frac{\underline{\eta} V^\pi}{\bar{\omega}} \cdot \frac{n}{m},$$

which completes the proof. The same result also holds with relocation times because incorporating relocation times only affects the controls $d^\pi_{ij,t}$ (as resources are blocked while relocating), and we can express resource utilization in terms of the controls as before.

A.2 Proof of Proposition 2.2

Since the network topology of the hub-and-spoke structure is strongly connected and the one-period revenue functions $r_{i0}(d)$ and $r_{0i}(d)$ are uniformly bounded by Assumption 2.1, Assumption 5.6.1 of Bertsekas (2012) holds. According to Proposition 5.6.2 of Bertsekas (2012), the average revenue V^{OPT} of an optimal policy does not depend on the initial state of the system, and moreover, there exists a solution to the Bellman equation (2) if randomized controls are allowed. Since the one-period revenue functions are concave, randomization does not improve performance and there must be a solution to (2) with controls being deterministic. Thus (2) has a solution. Finally, according to Proposition 5.6.1 of Bertsekas (2012), if $d^*(\mathbf{x}, s)$ attains the maximum in (2) for each state $(\mathbf{x}, s)$, the stationary policy $d^*(\mathbf{x}, s)$ is optimal.

A.3 Proof of Proposition 3.1

It is easy to see that (3) decomposes over spokes with each spoke problem being

$$\begin{aligned} \max_{\pi \in \Pi} \quad & \lim_{T \to \infty} \frac{1}{T} \cdot \mathbb{E}\left\{ \sum_{t=1}^{T} \left(y_{i0,t} \cdot r_{i0}\big(d^{\pi}_{i0,t}\big) + y_{0i,t} \cdot r_{0i}\big(d^{\pi}_{0i,t}\big) - \lambda \cdot x^{\pi}_{i,t} \right) \right\} \\ \text{s.t.} \quad & x^{\pi}_{i,t+1} = x^{\pi}_{i,t} - y_{i0,t} \cdot \mathbb{1}\big[\xi_t \le d^{\pi}_{i0,t}\big] + y_{0i,t} \cdot \mathbb{1}\big[\xi_t \le d^{\pi}_{0i,t}\big], \ \forall\, t \ge 1, \\ & 0 \le x^{\pi}_{i,t} \le m, \ \forall\, t \ge 1. \end{aligned} \tag{19}$$

We can interpret (19) as an average revenue problem for spoke i where in each time period, one request arrives following the same request rates as in the original problem, and every resource in spoke i incurs a holding cost λ. It is easy to see that Assumption 5.6.1 of Bertsekas (2012) holds for the spoke problem. By the same argument as in the proof of Proposition 2.2, the optimal average revenue h^{λ}_i does not depend on the initial state of spoke i, and moreover, h^{λ}_i together with some differential value functions $v^{\lambda}_i(x, i, 0)$, $v^{\lambda}_i(x, 0, i)$, and $v^{\lambda}_i(x, \varnothing)$ satisfies the Bellman equation (5). Since (3) decomposes into (19), the optimal average revenue $\bar{V}^{\lambda}$ of the Lagrangian relaxation does not depend on the initial state as well, and

$$\bar{V}^{\lambda} = m\lambda + \sum_{i=1}^{n} h^{\lambda}_i.$$

Finally, since every feasible policy to the original problem is feasible to the Lagrangian relaxation and attains an objective value that is no smaller, we have $\bar{V}^{\lambda} \ge V^{\text{OPT}}$.

A.4 Proof of Proposition 3.2

We first show that (6) is equivalent to the dual problem of the LP formulation of (5) in Section A.4.1.

A.4.1 Equivalence Between (5) and (6)

First, note that the optimal average revenue h_i^λ in (5) can be solved by the following linear program (20) following Section 5.5 of Bertsekas (2012).

$$
\begin{aligned}
\min_{\substack{h_i^\lambda, v_i^\lambda(x,i,0),\\ v_i^\lambda(x,0,i), v_i^\lambda(x,\varnothing)}} \quad & h_i^\lambda \\
\text{s.t.} \quad & h_i^\lambda + v_i^\lambda(x,i,0) \geq r_{i0}(d) + d\cdot v_i^\lambda(x-1) + (1-d)\cdot v_i^\lambda(x) - \lambda\cdot x, \\
& \qquad \forall\, x\in[0:m], d\in[0, 1\wedge x], \\
& h_i^\lambda + v_i^\lambda(x,0,i) \geq r_{0i}(d) + d\cdot v_i^\lambda(x+1) + (1-d)\cdot v_i^\lambda(x) - \lambda\cdot x, \\
& \qquad \forall\, x\in[0:m], d\in[0, 1\wedge(m-x)], \\
& h_i^\lambda + v_i^\lambda(x,\varnothing) \geq v_i^\lambda(x) - \lambda\cdot x,\ \forall\, x\in[0:m].
\end{aligned}
\tag{20}
$$

From Proposition 5.1.6 in Bertsekas (2012), any solution to (5) is an optimal solution to (20). Problem (20) is a semi-infinite linear program (see Section 4 in Anderson and Nash 1987) with a finite number of decision variables and infinitely many constraints.

According to Section 4.4 in Anderson and Nash (1987), the dual problem of (20) can be written as

$$
\begin{aligned}
\max_{\substack{F_{i,x}^{(i,0)}(d)\geq 0,\\ F_{i,x}^{(0,i)}(d)\geq 0,\\ p_i(x,\varnothing), p_i(x)\geq 0}} \quad & \sum_{x=0}^{m}\left\{\int_0^1 r_{i0}(d)\cdot dF_{i,x}^{(i,0)}(d) + \int_0^1 r_{0i}(d)\cdot dF_{i,x}^{(0,i)}(d)\right\} - \lambda\cdot\sum_{x=0}^{m} x\cdot p_i(x) \\
\text{s.t.} \quad & \sum_{x=0}^{m} p_i(x) = 1, \\
& p_i(x)\cdot q_{i0} = \int_0^1 dF_{i,x}^{(i,0)}(d),\ \forall\, x\in[0:m], \\
& p_i(x)\cdot q_{0i} = \int_0^1 dF_{i,x}^{(0,i)}(d),\ \forall\, x\in[0:m], \\
& p_i(x)\cdot(1-q_i) = p_i(x,\varnothing),\ \forall\, x\in[0:m], \\
& p_i(x) = \mathbb{1}\left[x\leq m-1\right]\cdot\int_0^1 d\cdot dF_{i,x+1}^{(i,0)}(d) + \mathbb{1}\left[x\geq 1\right]\cdot\int_0^1 d\cdot dF_{i,x-1}^{(0,i)}(d) \\
& + \int_0^1 (1-d)\cdot dF_{i,x}^{(i,0)}(d) + \int_0^1 (1-d)\cdot dF_{i,x}^{(0,i)}(d) + p_i(x,\varnothing),\ \forall\, x\in[0:m], \\
& F_{i,0}^{(i,0)}(d) = p_i(0)\cdot q_{i0},\ \forall\, d\in(0,1], \\
& F_{i,m}^{(0,i)}(d) = p_i(m)\cdot q_{0i},\ \forall\, d\in(0,1], \\
& F_{i,x}^{(i,0)}(d), F_{i,x}^{(0,i)}(d)\in M[0,1],\ \forall\, x\in[0:m],
\end{aligned}
\tag{21}
$$

where $M[0,1]$ is the set of Lebesgue-Stieltjes measures on interval $[0,1]$ with every $g(d)\in M[0,1]$ an increasing and right-continuous function with $g(0-)=0$. We can interpret the variables $F_{i,x}^{(i,0)}(d)$ and $F_{i,x}^{(0,i)}(d)$ as the joint probability that x resources are in spoke i, a request $(i,0)$ or $(0,i)$ arrives, and the service provider selects a demand level no larger than d. $p_i(x,\varnothing)$ is the probability that x resources are in spoke i and the request is one of any other types, and $p_i(x)$ is the probability with

x resources in spoke i.

We now show (21) and (6) are equivalent. To see this, note that every feasible solution to (21) represents a randomized control to the spoke problem. Specifically, at every state x with $p_i(x) > 0$, the provider selects a demand level according to the cumulative distribution $F_{i,x}^{(i,0)}(d)/(q_{i0} \cdot p_i(x))$ if a request $(i, 0)$ arrives, and cumulative distribution $F_{i,x}^{(0,i)}(d)/(q_{0i} \cdot p_i(x))$ if a request $(0, i)$ arrives. Since the one-period revenue functions $r_{i0}(d)$ and $r_{0i}(d)$ are concave, selecting the mean values $d_i(x, i, 0) = \int_0^1 d \cdot dF_{i,x}^{(i,0)}(d)/(q_{i0} \cdot p_i(x))$ for a request $(i, 0)$ and $d_i(x, 0, i) = \int_0^1 d \cdot dF_{i,x}^{(0,i)}(d)/(q_{0i} \cdot p_i(x))$ for a request $(0, i)$ can only be better. This implies that we can simply focus on deterministic controls in (21), which corresponds to (6). Thus, (21) and (6) are equivalent.

Finally, if we let $h_i^\lambda > \bar{r}$ and all the differential values be zero, we get a feasible solution to (20) with all constraints in (20) satisfied with strict inequality; thus strong duality holds according to Theorem 1 of Section 8.6 in Luenberger (1997).

Finally we point out that, from the complementary slackness property elaborated in the same theorem in Luenberger (1997), for all $x \in [0 : m]$ with $p_i(x) > 0$, we have:

$$h_i^\lambda + v_i^\lambda(x, i, 0) = r_{i0}\Big(d_i(x, i, 0)\Big) + d_i(x, i, 0) \cdot v_i^\lambda(x - 1) + \Big(1 - d_i(x, i, 0)\Big) \cdot v_i^\lambda(x) - \lambda \cdot x, \quad (22)$$

$$h_i^\lambda + v_i^\lambda(x, 0, i) = r_{0i}\Big(d_i(x, 0, i)\Big) + d_i(x, 0, i) \cdot v_i^\lambda(x + 1) + \Big(1 - d_i(x, 0, i)\Big) \cdot v_i^\lambda(x) - \lambda \cdot x, \quad (23)$$

and

$$h_i^\lambda + v_i^\lambda(x, \varnothing) = v_i^\lambda(x) - \lambda \cdot x, \quad (24)$$

where h_i^λ, $v_i^\lambda(x, i, 0)$, $v_i^\lambda(x, 0, i)$ and $v_i^\lambda(x, \varnothing)$ is an optimal solution to (20) and $p_i(x)$, $d_i(x, i, 0)$ and $d_i(x, 0, i)$ is an optimal solution to (6).

A.4.2 Equivalence Between (6) and (7)

From Lemma B.1, the support of the optimal probability distribution in (6) is $I_i = [0 : H_i]$ for some integer $H_i \in \mathbb{N}_+$. Introduce new variables $\beta_i(x) \geq 0$ for $x \in [0 : m - 1]$ such that $p_i(x + 1) = \beta_i(x) \cdot p_i(x)$. We can write (6) as

$$\begin{aligned}
h_i^\lambda = \max_{\substack{d_i(x,i,0)\in[0,1],\\ d_i(x,0,i)\in[0,1],\\ p_i(x),\beta_i(x)\geq 0}} \quad & \sum_{x=0}^{m} p_i(x)\Big[q_{i0} \cdot r_{i0}\Big(d_i(x, i, 0)\Big) + q_{0i} \cdot r_{0i}\Big(d_i(x, 0, i)\Big)\Big] - \lambda \cdot \sum_{x=0}^{m} x \cdot p_i(x) \\
\text{s.t.} \quad & \sum_{x=0}^{m} p_i(x) = 1, \\
& p_i(x) \cdot \beta_i(x) = p_i(x + 1), \ \forall\, x \in [0 : m - 1], \\
& q_{0i} \cdot d_i(x, 0, i) = \beta_i(x) \cdot q_{i0} \cdot d_i(x + 1, i, 0), \ \forall\, x \in [0 : m - 1], \\
& d_i(0, i, 0) = 0, \\
& d_i(m, 0, i) = 0.
\end{aligned} \quad (25)$$

The first part of the objective equals

$$
\begin{aligned}
&\sum_{x=0}^{m} p_i(x)\Big[q_{i0}\cdot r_{i0}\Big(d_i(x,i,0)\Big)+q_{0i}\cdot r_{0i}\Big(d_i(x,0,i)\Big)\Big]\\
&\overset{\text{(i)}}{=}\sum_{x=0}^{m-1}\Big\{p_i(x)\cdot q_{0i}\cdot r_{0i}\Big(d_i(x,0,i)\Big)+p_i(x+1)\cdot q_{i0}\cdot r_{i0}\Big(d_i(x+1,i,0)\Big)\Big\}\\
&=\sum_{x=0}^{m-1} p_i(x)\Big[q_{0i}\cdot r_{0i}\Big(d_i(x,0,i)\Big)+\beta_i(x)\cdot q_{i0}\cdot r_{i0}\Big(d_i(x+1,i,0)\Big)\Big],
\end{aligned}
\tag{26}
$$

where (i) is due to the constraints $d_i(0,i,0)=0$ and $d_i(m,0,i)=0$ and the fact that $r_{ij}(0)=0$. According to (26) and the constraints of (25), given $\beta_i(x)$, it is easy to solve $d_i(x,0,i)$ and $d_i(x+1,i,0)$ from the concave problem $\gamma_i(\beta)$ in (8), which is

$$
\begin{aligned}
\gamma_i(\beta)=\max_{d_{i0},d_{0i}\in[0,1]}\quad & q_{0i}\cdot r_{0i}(d_{0i})+\beta\cdot q_{i0}\cdot r_{i0}(d_{i0})\\
\text{s.t.}\quad & q_{0i}\cdot d_{0i}=\beta\cdot q_{i0}\cdot d_{i0}.
\end{aligned}
$$

Thus, (25) is equivalent to (27)

$$
\begin{aligned}
h_i^{\lambda}=\max_{p_i(x),\beta_i(x)\geq 0}\quad & \sum_{x=0}^{m-1} p_i(x)\cdot\gamma_i\Big(\beta_i(x)\Big)-\lambda\cdot\sum_{x=0}^{m} x\cdot p_i(x)\\
\text{s.t.}\quad & \sum_{x=0}^{m} p_i(x)=1,\\
& p_i(x)\cdot\beta_i(x)=p_i(x+1),\ \forall\, x\in[0:m-1].
\end{aligned}
\tag{27}
$$

Eliminating $\beta_i(x)$ from (27) yields (7), which is

$$
\begin{aligned}
h_i^{\lambda}=\max_{p_i(x)\geq 0}\quad & \sum_{x=0}^{m-1} p_i(x)\cdot\gamma_i\left(\frac{p_i(x+1)}{p_i(x)}\right)-\lambda\cdot\sum_{x=0}^{m} x\cdot p_i(x)\\
\text{s.t.}\quad & \sum_{x=0}^{m} p_i(x)=1,
\end{aligned}
$$

where we set $x\cdot\gamma_i\Big(\frac{y}{x}\Big)=0$ if $x=0$. Since the support of an optimal probability distribution in (7) is a sequence of consecutive interprets starting from zero, an optimal solution to (7) can be converted into a feasible solution to (27) and vice versa, thus the equivalence between (27) and (7). Lemma A.3 shows that the function $\gamma_i(\beta)$ is concave in β; this implies that (7) is a convex optimization problem, as we show in Lemma A.1.

Lemma A.1. *(7) is a convex optimization problem.*

Proof. It suffices to show the objective of (7) is concave in $p_i(x)$. Since $\gamma_i(\beta)$ is concave in β by Lemma A.3 and $x\cdot\gamma_i\big(\frac{y}{x}\big)$ is the perspective of $\gamma_i(\beta)$, $x\cdot\gamma_i\big(\frac{y}{x}\big)$ is jointly concave in (x,y) from Section 3.2.6 of Boyd and Vandenberghe (2004). This implies that the objective of (7) is concave in $p_i(x)$. $\square$

We can solve (7) efficiently with multiple methods. In Section A.4.3, we provide an specialized algorithm for solving (7) through its first-order optimality conditions. We provide some useful properties for $\gamma_i(\beta)$ as a preparation.

Lemma A.2. *$\gamma_i(\beta)$ is increasing in $\beta \in \mathbb{R}_+$.*

Proof. We show that for any $0 \le \beta_1 < \beta_2$, we have $\gamma_i(\beta_1) \le \gamma_i(\beta_2)$. Let d_{i0}^1 and d_{0i}^1 be an optimal solution to $\gamma_i(\beta_1)$. We have $q_{0i} \cdot d_{0i}^1 = \beta_1 \cdot q_{i0} \cdot d_{i0}^1$. It is easy to see that d_{0i}^1 and $\frac{\beta_1}{\beta_2} d_{i0}^1$ is feasible to $\gamma_i(\beta_2)$. Thus,

$$\begin{aligned}
\gamma_i(\beta_2) &\ge q_{0i} \cdot r_{0i}(d_{0i}^1) + \beta_2 \cdot q_{i0} \cdot r_{i0}\Big(\frac{\beta_1}{\beta_2} d_{i0}^1\Big) \\
&\overset{\text{(i)}}{\ge} q_{0i} \cdot r_{0i}(d_{0i}^1) + \beta_2 \cdot q_{i0} \cdot \frac{\beta_1}{\beta_2} \cdot r_{i0}\Big(d_{i0}^1\Big) \\
&= \gamma_i(\beta_1),
\end{aligned}$$

where (i) is due to $r_{i0}\Big(\frac{\beta_1}{\beta_2} d_{i0}^1\Big) \ge \frac{\beta_1}{\beta_2} \cdot r_{i0}\Big(d_{i0}^1\Big)$ because $r_{i0}(d)$ is concave and $r_{i0}(0) = 0$. □

Lemma A.3. *$\gamma_i(\beta)$ is strictly concave in $\beta \in \mathbb{R}_+$.*

Proof. For any $0 \le \beta_1 < \beta_2$ and $\alpha_1, \alpha_2 \in (0, 1)$ with $\alpha_1 + \alpha_2 = 1$, we let $\beta = \alpha_1 \cdot \beta_1 + \alpha_2 \cdot \beta_2$ and show that $\gamma_i(\beta) > \alpha_1 \cdot \gamma_i(\beta_1) + \alpha_2 \cdot \gamma_i(\beta_2)$.

Let d_{i0}^1 and d_{0i}^1 be an optimal solution to $\gamma_i(\beta_1)$, and d_{i0}^2 and d_{0i}^2 an optimal solution to $\gamma_i(\beta_2)$. Since $q_{0i} \cdot d_{0i}^1 = \beta_1 \cdot q_{i0} \cdot d_{i0}^1$ and $q_{0i} \cdot d_{0i}^2 = \beta_2 \cdot q_{i0} \cdot d_{i0}^2$, it is easy to see that $d_{0i} = \alpha_1 \cdot d_{0i}^1 + \alpha_2 \cdot d_{0i}^2$ and $d_{i0} = (\alpha_1 \cdot \beta_1 \cdot d_{i0}^1 + \alpha_2 \cdot \beta_2 \cdot d_{i0}^2)/(\alpha_1 \cdot \beta_1 + \alpha_2 \cdot \beta_2)$ is feasible to $\gamma_i(\beta)$. Thus,

$$\begin{aligned}
\gamma_i(\beta) &\ge q_{0i} \cdot r_{0i}\Big(\alpha_1 \cdot d_{0i}^1 + \alpha_2 \cdot d_{0i}^2\Big) + \beta \cdot q_{i0} \cdot r_{i0}\left(\frac{\alpha_1 \cdot \beta_1 \cdot d_{i0}^1 + \alpha_2 \cdot \beta_2 \cdot d_{i0}^2}{\alpha_1 \cdot \beta_1 + \alpha_2 \cdot \beta_2}\right) \\
&> q_{0i} \cdot \Big(\alpha_1 \cdot r_{0i}(d_{0i}^1) + \alpha_2 \cdot r_{0i}(d_{0i}^2)\Big) + q_{i0} \cdot \Big(\alpha_1 \cdot \beta_1 \cdot r_{i0}(d_{i0}^1) + \alpha_2 \cdot \beta_2 \cdot r_{i0}(d_{i0}^2)\Big) \\
&= \alpha_1 \cdot \gamma_i(\beta_1) + \alpha_2 \cdot \gamma_i(\beta_2)
\end{aligned}$$

where the second inequality is due to the strict concavity of the revenue functions and Jensen's inequality. □

For ease of exposition, in the following, we assume that $\gamma_i(\beta)$ is differentiable in β; otherwise, we can simply replace the derivatives of $\gamma_i(\beta)$ with its sub-gradients in the analysis.

Lemma A.4. *$\gamma_i(\beta)$ and its derivatives are bounded from above: $0 = \gamma_i(0) \le \gamma_i(\beta) \le q_{0i} \cdot (\bar{r} + \bar{\omega})$, and $0 \le \gamma_i'(\beta) \le \gamma_i'(0) \le q_{i0} \cdot (\bar{r} + \bar{\omega})$.*

Proof. It is easy to see from (8) that $\gamma_i(0) = 0$. Moreover, the objective of (8) satisfies that

$$\begin{aligned}
& q_{0i} \cdot r_{0i}(d_{0i}) + \beta \cdot q_{i0} \cdot r_{i0}(d_{i0}) \\
\le & q_{0i} \cdot \bar{r} + \beta \cdot q_{i0} \cdot d_{i0} \cdot \bar{\omega} \\
\le & q_{0i} \cdot (\bar{r} + \bar{\omega}),
\end{aligned}$$

where the first inequality is due to the mean value theorem and the facts that $r_{i0}(0) = 0$ and that $\bar{\omega}$ is the uniform bound on the derivatives of $r_{ij}(d)$ by Assumption 2.1. Thus $\gamma_i(\beta) \le q_{0i} \cdot (\bar{r} + \bar{\omega})$.

Finally, note that

$$\begin{aligned}\gamma_i(\beta) &\overset{\text{(i)}}{\leq} q_{0i} \cdot d_{0i} \cdot \bar{\omega} + \beta \cdot q_{i0} \cdot \bar{r} \\ &\overset{\text{(ii)}}{\leq} \beta \cdot q_{i0} \cdot (\bar{r} + \bar{\omega}),\end{aligned}$$

where (i) is from $r_{0i}(d_{0i}) \leq d_{0i} \cdot \bar{\omega}$ and (ii) from $q_{0i} \cdot d_{0i} = \beta \cdot q_{i0} \cdot d_{i0} \leq \beta \cdot q_{i0}$. We have

$$\gamma_i'(0+) = \lim_{\beta \to 0} \frac{\gamma_i(\beta) - \gamma_i(0)}{\beta} = \lim_{\beta \to 0} \frac{\gamma_i(\beta)}{\beta} \leq q_{i0} \cdot (\bar{r} + \bar{\omega}). \qquad \square$$

The remaining of Lemma A.4 is directly from Lemmas A.2 and A.3.

Lemma A.5. *Let $z(\beta) = \beta \cdot \gamma_i'(\beta) - \gamma_i(\beta)$ be a function of $\beta \in \mathbb{R}_+$. $z(\beta)$ is strictly decreasing in β and $z(0) = 0$.*

Proof. $z(0) = 0$ because $\gamma_i(0) = 0$ and $\gamma_i'(0)$ is bounded from Lemma A.4. To see that $z(\beta)$ is strictly decreasing in β, for any $0 \leq \beta_1 < \beta_2$ we have

$$\begin{aligned} z(\beta_1) - z(\beta_2) &= \gamma_i(\beta_2) - \gamma_i(\beta_1) + \beta_1 \cdot \gamma_i'(\beta_1) - \beta_2 \cdot \gamma_i'(\beta_2) \\ &= \Big\{\gamma_i(\beta_2) - \gamma_i(\beta_1) - \gamma_i'(\beta_2) \cdot (\beta_2 - \beta_1)\Big\} + \beta_1 \cdot \Big(\gamma_i'(\beta_1) - \gamma_i'(\beta_2)\Big) \\ &\geq \gamma_i(\beta_2) - \gamma_i(\beta_1) - \gamma_i'(\beta_2) \cdot (\beta_2 - \beta_1) \\ &> 0, \end{aligned}$$

where the first inequality is because $\gamma_i'(\beta_1) - \gamma_i'(\beta_2) \geq 0$ by the concavity of $\gamma_i(\beta)$, and the second inequality is due to the first-order condition of the strictly concave function $\gamma_i(\beta)$. $\square$

A.4.3 A Specialized Algorithm to Solve (7)

Since (7) is a convex program and all constraints are linear, strong duality holds. Let $f(\mathbf{p}) = \sum_{x=0}^{m-1} p_x \cdot \gamma_i\Big(\frac{p_{x+1}}{p_x}\Big) - \lambda \cdot \sum_{x=0}^{m} x \cdot p_x$ with $\mathbf{p} = (p_x)_{x \in [0:m]} \in \mathbb{R}_+^{m+1}$ denote the objective of (7). Relax the equality constraint $\sum_{x=0}^{m} p_i(x) = 1$ with a dual variable $r \in \mathbb{R}$ and let $L(\mathbf{p}, r) = f(\mathbf{p}) + r \cdot (1 - \sum_{x=0}^{m} p_x)$ denote the corresponding Lagrangian function and r^* denote the Lagrange multiplier. Proposition A.6 provides the optimality condition for (7).

Proposition A.6. *The following hold for (7).*

1. *The derivative of f is $\partial f / \partial p_x = -\lambda \cdot x + \gamma_i'\Big(\frac{p_x}{p_{x-1}}\Big) \cdot \mathbb{1}\Big[x \geq 1\Big] - z\Big(\frac{p_{x+1}}{p_x}\Big) \cdot \mathbb{1}\Big[x \leq m-1\Big]$;*
2. *$\mathbf{p} = (p_x)_{x \in [0:m]} \in \mathbb{R}_+^{m+1}$ is optimal to (7) and $r \in \mathbb{R}$ is a Lagrange multiplier if and only if (a) $\mathbf{p}$ is feasible to (7), and (b) $\frac{\partial f}{\partial p_x} = r$ for all $p_x > 0$ and $\frac{\partial f}{\partial p_x} \leq r$ for all $p_x = 0$;*
3. *The Lagrange multiplier $r^* = h_i^\lambda$ equals the optimal value of (7).*

Proof. Part 1 can be verified directly. Part 2 is from Proposition 6.2.5 in Bertsekas et al. (2003).

For part 3, suppose $\mathbf{p} = (p_x)_{x\in[0:m]}$ is an optimal solution to (7) with support $[0:H]$. We have

$$
\begin{aligned}
r^* &= \sum_{x\in[0:H]} p_x \cdot r^* \overset{\text{(i)}}{=} \sum_{x\in[0:H]} p_x \cdot \left(\frac{\partial f}{\partial p_x}\right) \\
&\overset{\text{(ii)}}{=} -\lambda \sum_{x\in[0:H]} p_x \cdot x + \sum_{x\in[1:H]} p_x \cdot \gamma_i'\left(\frac{p_x}{p_{x-1}}\right) - \sum_{x=0}^{H\bigwedge(m-1)} p_{x+1} \cdot \gamma_i'\left(\frac{p_{x+1}}{p_x}\right) + \sum_{x=0}^{H\bigwedge(m-1)} p_x \cdot \gamma_i\left(\frac{p_{x+1}}{p_x}\right) \\
&\overset{\text{(iii)}}{=} -\lambda \sum_{x\in[0:H]} p_x \cdot x + \sum_{x=0}^{H\bigwedge(m-1)} p_x \cdot \gamma_i\left(\frac{p_{x+1}}{p_x}\right) \\
&\overset{\text{(iv)}}{=} f(\mathbf{p}) = h_i^\lambda,
\end{aligned}
$$

where (i) is from part 2, (ii) from part 1 and the definition of $z(\beta) = \beta \cdot \gamma_i'(\beta) - \gamma_i(\beta)$, (iii) from $p_{H+1} = 0$ and thus $\sum_{x=0}^{H\bigwedge(m-1)} p_{x+1} \cdot \gamma_i'\left(\frac{p_{x+1}}{p_x}\right) = \sum_{x=1}^{H} p_x \cdot \gamma_i'\left(\frac{p_x}{p_{x-1}}\right)$, and (iv) from the fact that $p_x = 0$ for all $x > H$. $\square$

Lemma A.7 provides a bisection method to solve the optimality condition in Proposition A.6 part 2 efficiently.

Lemma A.7. *For any $r \geq 0$, let*

$$
m^* = \begin{cases} 0 & \text{if } \gamma_i'(0) \leq \lambda + r, \\ m & \text{if } \gamma_i'(0) > \lambda m + r, \\ \lceil \nu - 1 \rceil & \text{otherwise,} \end{cases}
$$

with $\nu = \frac{\gamma_i'(0)-r}{\lambda}$ and $\lceil x \rceil$ denoting the minimum integer that is no smaller than x.[8] *Let $\beta_x = 0$ for all $x \geq m^*$. If $m^* \geq 1$, set β_{m^*-1} to be the value that satisfies*

$$
\gamma_i'\big(\beta_{m^*-1}\big) = \lambda m^* + r, \tag{28}
$$

and set β_x for $x \leq m^* - 2$ recursively in the backward manner with[9]

$$
\gamma_i'\Big(\beta_x\Big) = z\Big(\beta_{x+1}\Big) + r + \lambda(x+1), \tag{29}
$$

where $z(\beta)$ is defined in Lemma A.5. We have

1. *β_x is decreasing in x: $\beta_0 \geq \cdots \geq \beta_{m^*-1} > 0 = \beta_{m^*} = \cdots = \beta_{m-1}$; and*
2. *if $r + z(\beta_0) = 0$, $r = r^*$ and the probabilities $p_i(x)$ that satisfy $p_i(x+1) = \beta_x \cdot p_i(x)$ for all $x \in [0:m-1]$ are optimal to (7); and*
3. *$r > r^*$ if $r + z(\beta_0) > 0$ and $r < r^*$ if $r + z(\beta_0) < 0$.*

From Lemma A.7 parts 2 and 3, we can solve the Lagrange multiplier $r^* = h_i^\lambda$ using a bisection method. Moreover, letting β_x^* be the values from (28) and (29) with $r = r^*$, the probabilities $p_i^*(x)$

[8] Equivalently, in case 3, m^* is the unique integer satisfying $\lambda m^* + r < \gamma_i'(0) \leq \lambda(m^*+1) + r$.

[9] If the right-hand side value of (29) is non-positive, return $r < r^*$.

that satisfy $p_i^*(x+1) = \beta_x^* \cdot p_i^*(x)$ for all $x \in [0:m-1]$ are optimal to (7). Since

$$p_i^*(x) = \left(\prod_{y=0}^{x-1} \beta_y^*\right) \cdot p_i^*(0),\ \forall\, x \in [m], \tag{30}$$

and these probabilities sum up to one, we have

$$p_i^*(0) = \left(1 + \sum_{x=1}^{m} \prod_{y=0}^{x-1} \beta_y^*\right)^{-1}. \tag{31}$$

From (30) and (31), we can compute $p_i^*(x)$ for all $x \le m$. Finally, since the ratios β_x^* are decreasing in $x \in [0:m-1]$ from Lemma A.7 part 1, the probability distribution $p_i^*(x)$ is discrete log-concave as defined in Definition B.1; this is the unique solution to (6) according to Proposition B.8.

Proof of Lemma A.7. Part 1: we prove by induction. As a base case, $\beta_{m^*-1} \ge \beta_{m^*} = 0$ and β_{m^*-1} satisfies (29) because $\beta_{m^*} = 0$, $z(0) = 0$ by Lemma A.5, and (28). Now for any $x \le m^* - 2$, we have

$$\begin{aligned}
\gamma_i'\Big(\beta_x\Big) &\overset{\text{(a)}}{=} z\Big(\beta_{x+1}\Big) + r + \lambda(x+1) \\
&= z\Big(\beta_{x+1}\Big) + r + \lambda(x+2) - \lambda \\
&\overset{\text{(b)}}{=} z\Big(\beta_{x+1}\Big) - z\Big(\beta_{x+2}\Big) + \gamma_i'\Big(\beta_{x+1}\Big) - \lambda \\
&\overset{\text{(c)}}{\le} \gamma_i'\Big(\beta_{x+1}\Big),
\end{aligned}$$

where (a) and (b) are because (29) holds at x and $x+1$, and (c) is from the facts that $\lambda \ge 0$, $z(\beta)$ is decreasing in β by Lemma A.5, and $\beta_{x+1} \ge \beta_{x+2}$ by assumption. Since $\gamma_i(\beta)$ is concave by Lemma A.3, we have $\beta_x \ge \beta_{x+1}$.

Part 2: this is essentially verifying the optimality condition as in Proposition A.6 part 2. Note that from Proposition A.6 part 1, the derivatives of the objective f only depend on the ratios $\beta_x = \frac{p_{x+1}}{p_x}$. If $r + z(\beta_0) = 0$, it is easy to check that the probabilities $p_i(x)$ with $p_i(x+1) = \beta_x \cdot p_i(x)$ for all $x \le m-1$ and the value r satisfy the optimality condition in Proposition A.6 part 2, with $\frac{\partial f}{\partial p_x} = r$ for all $x \le m^*$ and $\frac{\partial f}{\partial p_x} \le r$ for all $x \ge m^*+1$. Thus, $p_i(x)$ are optimal to (7) and $r = r^*$ equals the Lagrange multiplier.

Part 3: let $\beta_x(r)$ for all $0 \le x \le m-1$ and $m^*(r)$ denote the values of β_x and m^* with a specific r. Since $z(\beta)$ is decreasing in β from Lemma A.5, it suffices to show $\beta_x(r)$ is decreasing in r for all $0 \le x \le m-1$, which we will prove by induction. Clearly, $m^*(r)$ decreases in r. Thus, for any $r_1 > r_2$, $\beta_x(r_1) \le \beta_x(r_2)$ for all $x \ge m^*(r_1)$. Now, for any $x \le m^*(r_1) - 1$, suppose that $\beta_{x+1}(r_1) \le \beta_{x+1}(r_2)$. We have

$$\begin{aligned}
\gamma_i'\Big(\beta_x(r_1)\Big) &\overset{\text{(i)}}{=} z\Big(\beta_{x+1}(r_1)\Big) + r_1 + \lambda(x+1) \\
&\overset{\text{(ii)}}{\ge} z\Big(\beta_{x+1}(r_2)\Big) + r_2 + \lambda(x+1) \\
&\overset{\text{(iii)}}{=} \gamma_i'\Big(\beta_x(r_2)\Big),
\end{aligned}$$

where (i) and (iii) are from (29) and (ii) is from the facts that $z(\beta)$ is decreasing in β by Lemma

A.5, $\beta_{x+1}(r_1) \le \beta_{x+1}(r_2)$, and $r_1 > r_2$. Thus, $\beta_x(r_1) \le \beta_x(r_2)$ because $\gamma_i(\beta)$ is concave by Lemma A.3.

□

A.5 Proof of Proposition 3.3

Let h_i^λ, $v_i^\lambda(x,i,0)$, $v_i^\lambda(x,0,i)$ and $v_i^\lambda(x,\varnothing)$ be an optimal solution to (20). By the complementary slackness properties in (22)-(24) and the fact that I_i is closed under the Lagrangian policy by Proposition B.2, for all resource levels $x \in I_i = [0:H_i]$ we have

$$\begin{aligned} h_i^\lambda + v_i^\lambda(x,i,0) &= \max_{d\in[0,1\wedge x]} \left\{ r_{i0}(d) + d\cdot \Big(v_i^\lambda(x-1) - v_i^\lambda(x) \Big) \right\} + v_i^\lambda(x) - \lambda\cdot x, \\ h_i^\lambda + v_i^\lambda(x,0,i) &= \max_{d\in[0,1\wedge (H_i-x)]} \left\{ r_{0i}(d) + d\cdot \Big(v_i^\lambda(x+1) - v_i^\lambda(x) \Big) \right\} + v_i^\lambda(x) - \lambda\cdot x, \\ h_i^\lambda + v_i^\lambda(x,\varnothing) &= v_i^\lambda(x) - \lambda\cdot x, \end{aligned} \tag{32}$$

and the controls $d_i(x,i,0)$ and $d_i(x,0,i)$ of the Lagrangian policy attain the maximum in (32). We can interpret the Bellman equation (32) as an average revenue problem with states restricted to be in set I_i. Lemma A.8 shows the differential value functions in (32) are concave in x.

Lemma A.8. *The differential value functions $v_i^\lambda(x,i,0)$, $v_i^\lambda(x,0,i)$ and $v_i^\lambda(x,\varnothing)$ in (32) are concave in x for $x \in I_i$.*

We defer proof of Lemma A.8 to the end of this section. Let $\Delta v_i^\lambda(x) = v_i^\lambda(x) - v_i^\lambda(x-1)$ be the difference of the average differential values of two adjacent states. Lemma A.8 implies that $\Delta v_i^\lambda(x)$ is decreasing in x for $x \le H_i$. Since the one-period revenue functions $r_{i0}(d)$ and $r_{0i}(d)$ are strictly concave, the demand levels $d_i(x,i,0)$ and $d_i(x,0,i)$ that attain the maximum in (32) are unique. Moreover, since $d_i(x,i,0) = \operatorname{argmax}_{d\in[0,1\wedge x]}\{r_{i0}(d) - d\cdot \Delta v_i^\lambda(x)\}$ and the objective has increasing differences in d and $-\Delta v_i^\lambda(x)$, by the theory of monotone comparative statics (e.g., Topkis 1978, Milgrom and Shannon 1994, Topkis 2011), the unique optimal solution is increasing in $x \in I_i$ because $-\Delta v_i^\lambda(x)$ is increasing in x for $x \le H_i$ and $d_i(0,i,0) = 0$. Similar analysis implies the demand level $d_i(x,0,i) = \operatorname{argmax}_{d\in[0,1\wedge (H_i-x)]}\{r_{0i}(d) + d\cdot \Delta v_i^\lambda(x+1)\}$ is decreasing in $x \in I_i$.

Proof of Lemma A.8. Since the Lagrangian policy is optimal to (32) and is a unichain policy by Lemma B.1, Proposition 5.2.4 of Bertsekas (2012) implies that the differential value functions in (32) are unique up to a constant.

We now show the differential value functions are concave using a value iteration argument. Let $v_i^\lambda = \Big\{ v_i^\lambda(x,s) : x \in I_i, s \in \big\{(1,0),(0,1),\varnothing\big\} \Big\}$ be a set of value functions for states in the problem (32), and let Tv_i^λ be the one-step iteration with v_i^λ being the terminal values, i.e.,

$$\begin{aligned} Tv_i^\lambda(x,i,0) &= \max_{d\in[0,1\wedge x]} \left\{ r_{i0}(d) + d\cdot \Big(v_i^\lambda(x-1) - v_i^\lambda(x) \Big) \right\} + v_i^\lambda(x) - \lambda\cdot x, \\ Tv_i^\lambda(x,0,i) &= \max_{d\in[0,1\wedge (H_i-x)]} \left\{ r_{0i}(d) + d\cdot \Big(v_i^\lambda(x+1) - v_i^\lambda(x) \Big) \right\} + v_i^\lambda(x) - \lambda\cdot x, \\ Tv_i^\lambda(x,\varnothing) &= v_i^\lambda(x) - \lambda\cdot x, \end{aligned} \tag{33}$$

for all $x \in I_i$, with $v_i^\lambda(x) = q_{i0}\cdot v_i^\lambda(x,i,0) + q_{0i}\cdot v_i^\lambda(x,0,i) + (1-q_i)\cdot v_i^\lambda(x,\varnothing)$ being the average terminal values over request types. Lemma A.9 shows that the map T preserves concavity.

Lemma A.9. *If the value function* $v_i^\lambda = \Big\{v_i^\lambda(x,s) : x \in I_i, s \in \big\{(1,0),(0,1),\varnothing\big\}\Big\}$ *is concave in* x*, the one-step iteration* Tv_i^λ *is concave in* x *as well.*

Proof. The proof is standard (e.g., Proposition 5.2 of Talluri and Van Ryzin 2006) but we include it for completeness. Let $\Delta v_i^\lambda(x) = v_i^\lambda(x) - v_i^\lambda(x-1)$ for $x \le H_i$ be the difference of average values of two adjacent states. By assumption $\Delta v_i^\lambda(x)$ is decreasing in x. Let

$$\clubsuit := \Big\{Tv_i^\lambda(x+2,i,0) - Tv_i^\lambda(x+1,i,0)\Big\} - \Big\{Tv_i^\lambda(x+1,i,0) - Tv_i^\lambda(x,i,0)\Big\}.$$

We need to show that $\clubsuit \le 0$. Let $d_i(x,i,0)$ and $d_i(x,0,i)$ be the demand levels that attain the maximum in (33). From (33), for any $x \ge 0$ we have

$$\begin{aligned}
\clubsuit = {} & \Delta v_i^\lambda(x+2) - \Delta v_i^\lambda(x+1) \\
& + \Big\{r_{i0}\Big(d_i(x+2,i,0)\Big) - d_i(x+2,i,0)\cdot\Delta v_i^\lambda(x+2)\Big\} \\
& - \Big\{r_{i0}\Big(d_i(x+1,i,0)\Big) - d_i(x+1,i,0)\cdot\Delta v_i^\lambda(x+1)\Big\} \\
& - \Big\{r_{i0}\Big(d_i(x+1,i,0)\Big) - d_i(x+1,i,0)\cdot\Delta v_i^\lambda(x+1)\Big\} \\
& + \Big\{r_{i0}\Big(d_i(x,i,0)\Big) - d_i(x,i,0)\cdot\Delta v_i^\lambda(x)\Big\}.
\end{aligned}$$

Since $d_i(x+1,i,0)$ attains the maximum in (33),

$$r_{i0}\Big(d_i(x+1,i,0)\Big) - d_i(x+1,i,0)\cdot\Delta v_i^\lambda(x+1) \ge r_{i0}\Big(d_i(x,i,0)\Big) - d_i(x,i,0)\cdot\Delta v_i^\lambda(x+1)$$

and

$$r_{i0}\Big(d_i(x+1,i,0)\Big) - d_i(x+1,i,0)\cdot\Delta v_i^\lambda(x+1) \ge r_{i0}\Big(d_i(x+2,i,0)\Big) - d_i(x+2,i,0)\cdot\Delta v_i^\lambda(x+1).$$

Thus,

$$\begin{aligned}
\clubsuit \le {} & \Delta v_i^\lambda(x+2) - \Delta v_i^\lambda(x+1) \\
& + \Big\{r_{i0}\Big(d_i(x+2,i,0)\Big) - d_i(x+2,i,0)\cdot\Delta v_i^\lambda(x+2)\Big\} \\
& - \Big\{r_{i0}\Big(d_i(x+2,i,0)\Big) - d_i(x+2,i,0)\cdot\Delta v_i^\lambda(x+1)\Big\} \\
& - \Big\{r_{i0}\Big(d_i(x,i,0)\Big) - d_i(x,i,0)\cdot\Delta v_i^\lambda(x+1)\Big\} \\
& + \Big\{r_{i0}\Big(d_i(x,i,0)\Big) - d_i(x,i,0)\cdot\Delta v_i^\lambda(x)\Big\} \\
= {} & \Big(1 - d_i(x+2,i,0)\Big)\Big(\Delta v_i^\lambda(x+2) - \Delta v_i^\lambda(x+1)\Big) + d_i(x,i,0)\Big(\Delta v_i^\lambda(x+1) - \Delta v_i^\lambda(x)\Big) \\
\le {} & 0,
\end{aligned}$$

where the last inequality follows because $\Delta v_i^\lambda(x)$ is decreasing in x and demand levels are between zero and one. This implies $Tv_i^\lambda(x, i, 0)$ is concave in x. The same analysis on $Tv_i^\lambda(x, 0, i)$ implies $Tv_i^\lambda(x, 0, i)$ is concave in x as well. Finally, $Tv_i^\lambda(x, \varnothing)$ is concave in x because $Tv_i^\lambda(x, \varnothing) = v_i^\lambda(x) - \lambda \cdot x$ and $v_i^\lambda(x)$ is concave in x. □

Lemma A.9 and the convergence of value iteration in the proof of Proposition 5.6.2 in Bertsekas (2012) imply that the differential value functions in (32) are concave in x. □

A.6 Proof of Proposition 3.4

According to Danskin's Theorem (Proposition 4.5.1 in Bertsekas et al. 2003), the fact that the optimal probability distribution to (6) is unique (see Proposition B.8), and Proposition 4.2.4 in Bertsekas et al. (2003), the sub-differential of $\bar{V}^\lambda$ at any $\lambda \geq 0$ is a singleton

$$\partial \bar{V}^\lambda = \Big\{ m - \sum_{i=1}^{n} \sum_{x=0}^{m} x \cdot p_i(x) : p_i(x) \text{ is optimal to (6) with } \lambda \Big\}. \tag{34}$$

Thus by standard optimality conditions for convex optimization (Proposition 4.7.2 in Bertsekas et al. 2003), the dual variable λ^* is an optimal solution to (9) if and only if

$$\begin{aligned}
&\sum_{i=1}^{n} \sum_{x=0}^{m} x \cdot p_i(x) \leq m, \\
&\lambda^* \geq 0, \\
&\lambda^* \cdot \left(m - \sum_{i=1}^{n} \sum_{x=0}^{m} x \cdot p_i(x) \right) = 0, \\
&p_i(x) \text{ is an optimal solution to (6) with } \lambda^*,
\end{aligned}$$

which we can equivalently write as (10).

A.7 Proof of Lemma 4.2

First, we show that $\lambda^*(\delta) \leq \bar{r}/(m - \delta)$ if $\delta < m$. Since $\lambda^*(\delta)$ is an optimal solution to (11), $V^{\mathrm{R}}(\delta) = (m - \delta) \cdot \lambda^*(\delta) + \sum_{i=1}^{n} h_i^{\lambda^*(\delta)}$. It then suffices to show $V^{\mathrm{R}}(\delta) \leq \bar{r}$ and $h_i^\lambda \geq 0$ for all spokes $i \in [n]$ and dual variables $\lambda \geq 0$.

First, the optimality condition of (11) implies that (11) is equivalent to the problem of maximizing the average revenue subject to the constraint that the hub has at least δ resources in expectation. Since $\bar{r}$ is the uniform bound on the one-period revenue functions, $V^{\mathrm{R}}(\delta) \leq \bar{r}$. Second, h_i^λ is equal to the optimal value of (6) by Proposition 3.2. Let $p_i(0) = 1$, $p_i(x) = 0$ for all $x \geq 1$, $d_i(x, 0, i) = 0$ for all x, and $d_i(x, i, 0) = 1$ for all $x \geq 1$. This provides a feasible solution to (6) with an objective value of zero, thus $h_i^\lambda \geq 0$.

Combining the fact that $\lambda^*(\delta) \leq \bar{r}/(m - \delta)$ with (12) leads to the result.

A.8 Proof of Lemma 4.3

We prove the result by first showing that the value function of the Lagrangian policy in the relaxed system approximately solves the Bellman equation of the original system along the path induced by the Lagrangian policy in the original system. We then use a verification theorem to bound the total loss between these two systems. Finally, we extend to infinite horizon settings using a value iteration argument.

Step 1 (approximate Bellman equation). Let $v_{i,t}(x,i,0)$, $v_{i,t}(x,0,i)$ and $v_{i,t}(x,\varnothing)$ denote the value functions of the Lagrangian policy in each spoke i problem, with x resources and t time periods ahead, and the request type being $(i,0)$, $(0,i)$, or one of any other types, respectively. Let $v_{i,t}(x) = q_{i0}\cdot v_{i,t}(x,i,0) + q_{0i}\cdot v_{i,t}(x,0,i) + (1-q_i)\cdot v_{i,t}(x,\varnothing)$ be the average value functions over request types. The Bellman equation for each spoke problem with the Lagrangian policy is

$$
\begin{aligned}
v_{i,t}(x,i,0) &= r_{i0}\big(d_i(x,i,0)\big) + d_i(x,i,0)\cdot\Big(v_{i,t-1}(x-1) - v_{i,t-1}(x)\Big) + v_{i,t-1}(x),\\
v_{i,t}(x,0,i) &= r_{0i}\big(d_i(x,0,i)\big) + d_i(x,0,i)\cdot\Big(v_{i,t-1}(x+1) - v_{i,t-1}(x)\Big) + v_{i,t-1}(x),\\
v_{i,t}(x,\varnothing) &= v_{i,t-1}(x),
\end{aligned}
\tag{35}
$$

for all $x\in[0:m]$, where $d_i(x,i,0)$ and $d_i(x,0,i)$ are the demand values of the Lagrangian policy. Let $\Delta v_{i,t}(x) = v_{i,t}(x) - v_{i,t}(x-1)$ for $x\in[m]$ be the difference of the continuation values of two adjacent states. Lemma A.10 shows that $\Delta v_{i,t}(x)\le\bar{\omega}$ are uniformly bounded from above by the derivative bound $\bar{\omega}$ as defined in Assumption 2.1.

Lemma A.10. *The difference of the continuation values $\Delta v_{i,t}(x)$ satisfies $\Delta v_{i,t}(x)\le\bar{\omega}$ for all spokes i, time periods t, and resource levels $x\in[m]$, where $\bar{\omega}$ is the uniform bound on the derivatives of the one-period revenue functions as defined in Assumption 2.1.*

We prove Lemma A.10 at the end of this section. Let $V_t^{\mathrm{R}}(\mathbf{x},s)$ and $V_t(\mathbf{x},s)$ denote the continuation values of the Lagrangian policy in the relaxed and the original systems, with $\mathbf{x}=(x_i)_{i\in[n]}$ being the state of resources and s being the request type. Let $V_t^{\mathrm{R}}(\mathbf{x}) = \mathbb{E}_s\big[V_t^{\mathrm{R}}(\mathbf{x},s)\big]$ and $V_t(\mathbf{x}) = \mathbb{E}_s\big[V_t(\mathbf{x},s)\big]$ denote the average values over request types. The boundary conditions of the two systems are $V_0^{\mathrm{R}}(\mathbf{x},s) = V_0(\mathbf{x},s) = 0$. Because we relax the capacity constraint of the hub in the relaxed system, $V_t^{\mathrm{R}}(\mathbf{x})$ decouples over spokes with

$$
V_t^{\mathrm{R}}(\mathbf{x}) = \sum_{i=1}^{n} v_{i,t}(x_i),
\tag{36}
$$

where $v_{i,t}(x)$ are the average value functions of spoke i as in (35).

The Lagrangian policy takes different actions in the two systems at the same state $(\mathbf{x},s)$ only when $x_0=0$, $s=(0,i)$, and $x_i\le m-1$ for some spoke $i\in[n]$; let

$$
\mathcal{A} = \Big\{(\mathbf{x},s) : x_0 = 0, s=(0,i), x_i\le m-1, i\in[n]\Big\}
$$

be the set of states in which the policy can take different actions. Let $R(\mathbf{x},s)$ be the expected one-period revenue in the original system at state $(\mathbf{x},s)\in\mathcal{X}\times\big\{(i,0),(0,i): i\in[n]\big\}$, i.e.,

$$
\begin{aligned}
R(\mathbf{x},i,0) &= r_{i0}\Big(d_i(x_i,i,0)\Big),\\
R(\mathbf{x},0,i) &= r_{0i}\Big(d_i(x_i,0,i)\Big)\cdot\Big(1-\mathbb{1}\big[(\mathbf{x},0,i)\in\mathcal{A}\big]\Big).
\end{aligned}
\tag{37}
$$

Moreover, let

$$
\bar{R}_t(\mathbf{x},s) = V_t^{\mathrm{R}}(\mathbf{x},s) - R(\mathbf{x},s) - \mathbb{E}\Big[V_{t-1}^{\mathrm{R}}(\tilde{\mathbf{x}},\tilde{s})\Big|\mathbf{x},s\Big],
\tag{38}
$$

where the expectation is taken with respect the random variable $(\tilde{\mathbf{x}},\tilde{s})$, which is the next state in the original system under the Lagrangian policy when the current state is $(\mathbf{x},s)$. This value can be interpreted as the ex-ante compensation that needs to be given to the provider in the relaxed system at state $(\mathbf{x}, s)$ in order for her to be willing to switch from the current action to the action in the original system.

Fix a state $(\mathbf{x}, s)$ of the original system. We can write the value function of the relaxed system as follows:

$$\begin{aligned}
V_t^{\mathrm{R}}(\mathbf{x}, s) &= \mathbb{1}\left[(\mathbf{x}, s) \notin \mathcal{A}\right] \cdot V_t^{\mathrm{R}}(\mathbf{x}, s) + \mathbb{1}\left[(\mathbf{x}, s) \in \mathcal{A}\right] \cdot V_t^{\mathrm{R}}(\mathbf{x}, s) \\
&= \mathbb{1}\left[(\mathbf{x}, s) \notin \mathcal{A}\right] \cdot \left(R(\mathbf{x}, s) + \mathbb{E}\left[V_{t-1}^{\mathrm{R}}(\tilde{\mathbf{x}}, \tilde{s}) \middle| \mathbf{x}, s\right]\right) \\
&\quad + \mathbb{1}\left[(\mathbf{x}, s) \in \mathcal{A}\right] \cdot \left(R(\mathbf{x}, s) + \mathbb{E}\left[V_{t-1}^{\mathrm{R}}(\tilde{\mathbf{x}}, \tilde{s}) \middle| \mathbf{x}, s\right] + \bar{R}_t(\mathbf{x}, s)\right) \\
&= R(\mathbf{x}, s) + \mathbb{E}\left[V_{t-1}^{\mathrm{R}}(\tilde{\mathbf{x}}, \tilde{s}) \middle| \mathbf{x}, s\right] + \underbrace{\mathbb{1}\left[(\mathbf{x}, s) \in \mathcal{A}\right] \cdot \bar{R}_t(\mathbf{x}, s)}_{\varepsilon_t},
\end{aligned}$$

where the second equation follows from the Bellman equation for the relaxed system together with the fact that the evolution in the relaxed and original system coincide for all states not in $\mathcal{A}$ and using (38). We proceed by bounding the terms ε_t. Let $z(s)$ be the index of the spoke involved in type s and let $n(\mathbf{x}, s)$ be the resource level of spoke $z(s)$ when the state of resources is $\mathbf{x}$. We have

$$\begin{aligned}
\varepsilon_t &= \mathbb{1}\left[(\mathbf{x}, s) \in \mathcal{A}\right] \cdot \bar{R}_t(\mathbf{x}, s) \\
&\overset{\text{(i)}}{=} \mathbb{1}\left[(\mathbf{x}, s) \in \mathcal{A}\right] \cdot \left(V_t^{\mathrm{R}}(\mathbf{x}, s) - \mathbb{E}\left[V_{t-1}^{\mathrm{R}}(\tilde{\mathbf{x}}, \tilde{s}) \middle| \mathbf{x}, s\right]\right) \\
&\overset{\text{(ii)}}{=} \mathbb{1}\left[(\mathbf{x}, s) \in \mathcal{A}\right] \cdot \left(V_t^{\mathrm{R}}(\mathbf{x}, s) - V_{t-1}^{\mathrm{R}}(\mathbf{x})\right) \\
&\overset{\text{(iii)}}{=} \mathbb{1}\left[(\mathbf{x}, s) \in \mathcal{A}\right] \cdot \left[r_s\Big(d_{z(s)}\Big(n(\mathbf{x}, s), s\Big)\Big) + d_{z(s)}\Big(n(\mathbf{x}, s), s\Big) \cdot \Delta v_{z(s), t-1}\Big(n(\mathbf{x}, s) + 1\Big)\right] \\
&\overset{\text{(iv)}}{\leq} \mathbb{1}\left[(\mathbf{x}, s) \in \mathcal{A}\right] \cdot \left(\bar{r} + \bar{\omega}\right)
\end{aligned}$$

where (i) follows from (38) and the fact that $R(\mathbf{x}, s) = 0$ when $(\mathbf{x}, s) \in \mathcal{A}$ by (37), (ii) from $\tilde{\mathbf{x}} = \mathbf{x}$ when $(\mathbf{x}, s) \in \mathcal{A}$, (iii) from the Bellman equation of $V_t^{\mathrm{R}}(\mathbf{x}, s)$, the fact that the value function decomposes over spokes by (36) and the transition only involves $s = (0, i)$, (iv) from the definition of $\bar{r}$ in Assumption 2.1 and Lemma A.10. Putting everything together, we obtain that the value function for the relaxed system satisfies the following approximate Bellman equation in the original system

$$V_t^{\mathrm{R}}(\mathbf{x}, s) \leq R(\mathbf{x}, s) + \mathbb{E}\left[V_{t-1}^{\mathrm{R}}(\tilde{\mathbf{x}}, \tilde{s}) \middle| \mathbf{x}, s\right] + \mathbb{1}\left[(\mathbf{x}, s) \in \mathcal{A}\right] \cdot \left(\bar{r} + \bar{\omega}\right). \tag{39}$$

Step 2 (verification). Let $\{\mathbf{x}_\tau, s_\tau\}_{\tau \leq t}$ be the path of states of the original system with the Lagrangian policy. By taking expectations over the states $\{\mathbf{x}_\tau, s_\tau\}_{\tau < t}$ and using the boundary condition on the value function we have

$$
\begin{aligned}
V_t^{\mathrm{R}}(\mathbf{x}_t, s_t) &= \mathbb{E}\left[\sum_{\tau=1}^{t} V_\tau^{\mathrm{R}}(\mathbf{x}_\tau, s_\tau) - V_{\tau-1}^{\mathrm{R}}(\mathbf{x}_{\tau-1}, s_{\tau-1})\right] \\
&\overset{\text{(i)}}{=} \mathbb{E}\left[\sum_{\tau=1}^{t} V_\tau^{\mathrm{R}}(\mathbf{x}_\tau, s_\tau) - \mathbb{E}\Big[V_{\tau-1}^{\mathrm{R}}(\tilde{\mathbf{x}}, \tilde{s})|\mathbf{x}_\tau, s_\tau\Big]\right] \\
&\overset{\text{(ii)}}{\leq} \mathbb{E}\left[\sum_{\tau=1}^{t} R(\mathbf{x}_\tau, s_\tau)\right] + \big(\bar{r} + \bar{\omega}\big)\sum_{\tau=1}^{t} \mathbb{P}\Big[(\mathbf{x}_\tau, s_\tau) \in \mathcal{A}\Big] \\
&\overset{\text{(iii)}}{=} V_t(\mathbf{x}_t, s_t) + \big(\bar{r} + \bar{\omega}\big)\sum_{\tau=1}^{t} \mathbb{P}\Big[(\mathbf{x}_\tau, s_\tau) \in \mathcal{A}\Big] \\
&\overset{\text{(iv)}}{\leq} V_t(\mathbf{x}_t, s_t) + \big(\bar{r} + \bar{\omega}\big)\sum_{\tau=1}^{t} \mathbb{P}\Big[x_{0,\tau} = 0\Big],
\end{aligned}
\tag{40}
$$

where the (i) follows by the tower rule for conditional expectations and using the fact that the dynamics are Markovian, (ii) follows from (39) over $\tau \in [t]$ together with linearity of expectations, (iii) because $R(\mathbf{x}_\tau, s_\tau)$ is the expected one-period revenue in the original system at state $(\mathbf{x}_\tau, s_\tau)$, and (iv) because $\mathcal{A} \subseteq \big\{(\mathbf{x}, s) : x_0 = 0\big\}$.

Step 3 (value iteration). Taking an average over t time periods and letting t go to infinity gives

$$
\begin{aligned}
V^{\mathrm{R}}(\delta) &\overset{\text{(a)}}{=} \lim_{t\to\infty} \frac{1}{t} V_t^{\mathrm{R}}(\mathbf{x}_t, s_t) \\
&\overset{\text{(b)}}{\leq} \lim_{t\to\infty} \frac{1}{t}\Big(V_t(\mathbf{x}_t, s_t) + (\bar{r} + \bar{\omega}) \cdot \sum_{\tau=1}^{t} \mathbb{P}[x_{0,\tau} = 0]\Big) \\
&\overset{\text{(c)}}{=} V^\pi(\delta) + (\bar{r} + \bar{\omega}) \cdot \mathbb{P}\big[X_0(\delta) = 0\big],
\end{aligned}
$$

where (a) is due to the fact that the long-run time average of the total revenue converges to the average revenue of the policy by a value iteration argument (see Proposition 5.3.1 in Bertsekas 2012), (b) is from (40), and (c) from the same value iteration argument and the fact that the time-average limiting distribution converges to the stationary distribution because the Markov chain has a single recurrent class by Corollary B.11.

Proof of Lemma A.10. According to Proposition 3.3, the demand values of the Lagrangian policy are monotone in the resource levels of the spokes: for each spoke i and for all $x \in [m]$, we have $0 \leq d_i(x, 0, i) \leq d_i(x-1, 0, i) \leq 1$ and $0 \leq d_i(x-1, i, 0) \leq d_i(x, i, 0) \leq 1$. By coupling the private value of the arriving request when the number of resources is $x-1$ and x respectively, we can write the difference of the continuation values $\Delta v_{i,t}(x)$ in following recursive way:

$$\begin{aligned}
\Delta v_{i,t}(x) = q_{0i}\Big\{ &\Big(1 - d_i(x-1,0,i)\Big) \cdot \Delta v_{i,t-1}(x) \\
&+ \Big(d_i(x-1,0,i) - d_i(x,0,i)\Big) \cdot \Big(-G_{0i}\big(d_i(x-1,0,i)\big)\Big) \\
&+ d_i(x,0,i) \cdot \Big(-G_{0i}\big(d_i(x-1,0,i)\big) + G_{0i}\big(d_i(x,0,i)\big) + \Delta v_{i,t-1}(x+1)\Big)\Big\} \\
+ q_{i0}\Big\{ &\Big(1 - d_i(x,i,0)\Big) \cdot \Delta v_{i,t-1}(x) \\
&+ \Big(d_i(x,i,0) - d_i(x-1,i,0)\Big) \cdot G_{i0}\big(d_i(x,i,0)\big) \\
&+ d_i(x-1,i,0) \cdot \Big(-G_{i0}\big(d_i(x-1,i,0)\big) + G_{i0}\big(d_i(x,i,0)\big) + \Delta v_{i,t-1}(x-1)\Big)\Big\} \\
+ \Big(1 - q_i\Big) &\cdot \Delta v_{i,t-1}(x) \\
= q_{0i}\Big(&r_{0i}\big(d_i(x,0,i)\big) - r_{0i}\big(d_i(x-1,0,i)\big) + d_i(x,0,i) \cdot \Delta v_{i,t-1}(x+1)\Big) \\
+ q_{i0}\Big(&r_{i0}\big(d_i(x,i,0)\big) - r_{i0}\big(d_i(x-1,i,0)\big) + d_i(x-1,i,0) \cdot \Delta v_{i,t-1}(x-1)\Big) \\
+ \Big(1 &- q_{0i} \cdot d_i(x-1,0,i) - q_{i0} \cdot d_i(x,i,0)\Big)\Delta v_{i,t-1}(x),
\end{aligned} \tag{41}$$

with boundary conditions $\Delta v_{i,0}(x) = 0$ for all x.

We prove by induction. Clearly this is true for $t = 0$ by the boundary conditions that $\Delta v_{i,0}(x) = 0$. Now suppose $\Delta v_{i,t-1}(x) \leq \bar{\omega}$ for all spokes i and resource levels x. We show that $\Delta v_{i,t}(x) \leq \bar{\omega}$. From (41) we have

$$\begin{aligned}
\Delta v_{i,t}(x) \leq q_{0i}&\Big(r_{0i}\big(d_i(x,0,i)\big) - r_{0i}\big(d_i(x-1,0,i)\big)\Big) + q_{0i} \cdot d_i(x,0,i) \cdot \bar{\omega} \\
&+ q_{i0}\Big(r_{i0}\big(d_i(x,i,0)\big) - r_{i0}\big(d_i(x-1,i,0)\big)\Big) + q_{i0} \cdot d_i(x-1,i,0) \cdot \bar{\omega} \\
&+ \Big(1 - q_{0i} \cdot d_i(x-1,0,i) - q_{i0} \cdot d_i(x,i,0)\Big) \cdot \bar{\omega}.
\end{aligned}$$

To show $\Delta v_{i,t}(x) \leq \bar{\omega}$, it suffices to show that

$$\begin{aligned}
&q_{0i}\Big(r_{0i}\big(d_i(x,0,i)\big) - r_{0i}\big(d_i(x-1,0,i)\big)\Big) + q_{i0}\Big(r_{i0}\big(d_i(x,i,0)\big) - r_{i0}\big(d_i(x-1,i,0)\big)\Big) \\
&\leq \Bigg\{q_{0i}\Big(d_i(x-1,0,i) - d_i(x,0,i)\Big) + q_{i0}\Big(d_i(x,i,0) - d_i(x-1,i,0)\Big)\Bigg\} \cdot \bar{\omega}.
\end{aligned}$$

This is true because the left-hand side satisfies that

$$\begin{aligned}
&q_{0i}\Big(r_{0i}\big(d_i(x,0,i)\big) - r_{0i}\big(d_i(x-1,0,i)\big)\Big) + q_{i0}\Big(r_{i0}\big(d_i(x,i,0)\big) - r_{i0}\big(d_i(x-1,i,0)\big)\Big) \\
&\leq q_{0i}\Big|r_{0i}\big(d_i(x,0,i)\big) - r_{0i}\big(d_i(x-1,0,i)\big)\Big| + q_{i0}\Big|r_{i0}\big(d_i(x,i,0)\big) - r_{i0}\big(d_i(x-1,i,0)\big)\Big| \\
&\leq \Bigg\{q_{0i}\Big(d_i(x-1,0,i) - d_i(x,0,i)\Big) + q_{i0}\Big(d_i(x,i,0) - d_i(x-1,i,0)\Big)\Bigg\} \cdot \bar{\omega},
\end{aligned}$$

where the last inequality is due to the mean value theorem, the monotonicity property – i.e., $d_i(x, 0, i) \leq d_i(x-1, 0, i)$ and $d_i(x-1, i, 0) \leq d_i(x, i, 0)$ – and the fact that $\bar{\omega}$ is the uniform bound on the derivatives of the one-period revenue functions as in Assumption 2.1. □

A.9 Proof of Lemma 4.4

Let $\tilde{x}_{i,t}$ and $x_{i,t}$ denote the number of resources in locations $i \in [0:n]$ at time t in the relaxed and original systems, respectively. Lemma A.11 shows that if the two systems start at the same state and have the same sequence of requests and private values, $\tilde{x}_{i,t} \geq x_{i,t}$ for all spokes $i \in [n]$ and time periods t.

Lemma A.11. *If the relaxed and original systems start at the same state and have the same sequence of requests and private values, for any time period t, $\tilde{x}_{i,t} \geq x_{i,t}$ for all spokes $i \in [n]$ and $\tilde{x}_{0,t} \leq x_{0,t}$.*

Proof. We prove by induction. Since the two systems start at the same state, $\tilde{x}_{i,0} = x_{i,0}$ for all $i \in [n]$. For each spoke i, first suppose $\tilde{x}_{i,t-1} = x_{i,t-1}$. If request $(i, 0)$ arrives at time t, the Lagrangian policy takes the same action in the two systems, hence $\tilde{x}_{i,t} = x_{i,t}$. If request $(0, i)$ arrives, the Lagrangian policy takes different actions in the two systems only when the hub of the original system runs out of resources, in which case we have $\tilde{x}_{i,t} \geq \tilde{x}_{i,t-1} = x_{i,t-1} = x_{i,t}$. Next suppose $\tilde{x}_{i,t-1} \geq x_{i,t-1} + 1$ for spoke i. If request $(i, 0)$ arrives at time t, we have $\tilde{x}_{i,t} \geq \tilde{x}_{i,t-1} - 1 \geq x_{i,t-1} \geq x_{i,t}$. If request $(0, i)$ arrives, $\tilde{x}_{i,t} \geq \tilde{x}_{i,t-1} \geq x_{i,t-1} + 1 \geq x_{i,t}$. Thus by induction, $\tilde{x}_{i,t} \geq x_{i,t}$ for all spokes $i \in [n]$ and time periods t. Finally, since $\sum_{i=0}^{n} x_{i,t} = m$ and $\sum_{i=0}^{n} \tilde{x}_{i,t} = m$, we have $\tilde{x}_{0,t} \leq x_{0,t}$ for all time periods t. □

Lemma A.11 implies that for any integer k,

$$\mathbb{P}\Big[\tilde{X}_i(\delta) \leq k\Big] = \lim_{t\to\infty} \mathbb{P}\Big[\tilde{x}_{i,t} \leq k\Big] \leq \lim_{t\to\infty} \mathbb{P}\Big[x_{i,t} \leq k\Big] = \mathbb{P}\Big[X_i(\delta) \leq k\Big],$$

and

$$\mathbb{P}\Big[X_0(\delta) \leq k\Big] = \lim_{t\to\infty} \mathbb{P}\Big[x_{0,t} \leq k\Big] \leq \lim_{t\to\infty} \mathbb{P}\Big[\tilde{x}_{0,t} \leq k\Big] = \mathbb{P}\Big[\tilde{X}_0(\delta) \leq k\Big],$$

where the equations are because in both systems, the limiting distribution of the Markov chain converges to the unique stationary distribution, independently of the initial state, due to Corollaries B.10 and B.11. The inequalities follow from Lemma A.11 when we start the two systems with the same state and couple the sequence of requests and their private values. Note that in the proof we only use the fact that the Lagrangian policy only depends on the state of resources through the resource level of the spoke involved in the request type.

A.10 Proof of Proposition 4.6

Proposition 4.6 comes from Lemma A.12, which provide a concentration inequality for a sequence of independent random variables with discrete log-concave distributions (defined in Definition B.1) and uniformly bounded means.

Lemma A.12. *Let $\big\{X_i\big\}_{i=1}^{n}$ be a sequence of independent discrete log-concave random variables each with mean value $\mu_i = \mathbb{E}\big[X_i\big]$. If $\mu_i \leq c$ for all $i \leq [n]$ are uniformly bounded from above by some constant $c > 0$, then for any $\lambda \geq 1$ and letting $X = \sum_{i=1}^{n} X_i$ and $\mu = \mathbb{E}\big[X\big] = \sum_{i=1}^{n} \mu_i$, we have*

$$\mathbb{P}\Big[X \geq \lambda\mu\Big] \leq \exp\left\{ -\frac{(\lambda-1)\mu}{1+c} - \frac{n+\mu}{1+c} \ln\left(1 - \frac{\lambda\mu - \mu}{\lambda\mu + n}\right)\right\}. \tag{42}$$

We prove Lemma A.12 at the end of this section. We can apply Lemma A.12 to $\tilde{X}_i(\delta)$ for $i \in [n]$ because each $\tilde{X}_i(\delta)$ is log-concave by Proposition B.9, and $\tilde{X}_i(\delta)$ are independent because the joint distribution is equal to the product of their marginal distributions by Corollary B.10.

Let $\mu = \mathbb{E}\Big[\sum_{i=1}^n \tilde{X}_i(\delta)\Big]$ be the expected number of resources in the spokes of the relaxed system. We have $0 < \mu \leq m - \delta$. Applying (42) with $\lambda = \frac{m}{\mu}$ and $b = \frac{1}{1+c}$ gives

$$\begin{aligned}
\mathbb{P}\Big[\tilde{X}_0(\delta) \leq 0\Big] &= \mathbb{P}\left[\sum_{i=1}^n \tilde{X}_i(\delta) \geq m\right] \\
&\leq \exp\left\{ -b \cdot \left(\lambda\mu - \mu + (n+\mu)\cdot \ln\left(1 - \frac{\lambda\mu - \mu}{\lambda\mu + n}\right)\right)\right\} \\
&= \exp\left\{ -b \cdot \left(m - \mu + (n+\mu)\cdot \ln\left(\frac{n+\mu}{m+n}\right)\right)\right\} \\
&= \exp\left\{ b \cdot \left(\underbrace{(n+\mu)\cdot \ln\left(\frac{m+n}{n+\mu}\right) - (m-\mu)}_{\spadesuit}\right)\right\}.
\end{aligned} \tag{43}$$

Since $\ln x \leq \frac{x-1}{\sqrt{x}}$ for $x \geq 1$, we have

$$\spadesuit \leq (n+\mu)\cdot\frac{m-\mu}{n+\mu}\cdot\sqrt{\frac{n+\mu}{m+n}} - (m-\mu) = (m-\mu)\cdot\left(\sqrt{1-\frac{m-\mu}{m+n}} - 1\right) \leq -\frac{(m-\mu)^2}{2\cdot(m+n)} \leq -\frac{\delta^2}{2\cdot(m+n)},$$

where the second-to-last inequality is due to $\sqrt{1-x} - 1 \leq -\frac{x}{2}$ for $x \leq 1$. Thus from (43) we have

$$\mathbb{P}\Big[\tilde{X}_0(\delta) \leq 0\Big] \leq \exp\left(-\frac{b}{2}\cdot\frac{\delta^2}{m+n}\right).$$

Proof of Lemma A.12. We first provide an upper bound on the probability generating function of a log-concave random variable in Lemma A.14. The proof is based on the classic inequality bounding the factorial moments of a log-concave random variable, as we state in Lemma A.13.

Definition A.1 (Factorial Moment). Let $\mathbf{p} = \{p_i\}_{i=0}^\infty$ be a discrete distribution with all its support on non-negative integers. The factorial moment of $\mathbf{p}$ of order $r \geq 1$ is

$$\mu_{[r]} = \sum_{i=0}^\infty p_i \cdot \Big\{i\cdot(i-1)\cdots(i-r+1)\Big\} = \sum_{i=r}^\infty p_i \cdot \Big\{i\cdot(i-1)\cdots(i-r+1)\Big\} = \sum_{i=r}^\infty p_i \cdot \frac{i!}{(i-r)!}.$$

We set $\mu_{[0]} = 1$ for convenience.

Lemma A.13 (Theorem 2 in Keilson 1972). *Let $\mathbf{p} = \{p_i\}_{i=0}^\infty$ be a discrete log-concave distribution and let $\mu_{[r]}$ denote its order-r factorial moment. For any $r \geq 1$ we have*

$$\left\{\frac{\mu_{[r+1]}}{(r+1)!}\right\}^{1/(r+1)} \leq \left\{\frac{\mu_{[r]}}{r!}\right\}^{1/r} \leq \cdots \leq \frac{\mu_{[1]}}{1!} = \mu, \tag{44}$$

where μ denotes the mean value of $\mathbf{p}$. All inequalities in (44) hold with equalities when $\mathbf{p}$ is a geometric distribution, i.e., when $p_i = \theta(1-\theta)^i$ for some $0 < \theta \leq 1$.

Lemma A.14 provides an upper bound on the probability generating function of a log-concave random variable based on Lemma A.13.

Lemma A.14. *Let X be a discrete log-concave random variable (defined in Definition B.1) and let $\mu = \mathbb{E}[X]$ denote its mean value. We have $\mathbb{E}\left[z^X\right] \leq \frac{1}{1-\mu(z-1)}$ for all $1 \leq z < 1 + \frac{1}{\mu}$.*

Proof. First, we have

$$\mathbb{E}\left[z^X\right] = \sum_{i=0}^{\infty} p_i \cdot z^i \overset{\text{(i)}}{=} \sum_{i=0}^{\infty} p_i \sum_{j=0}^{i} \binom{i}{j}(z-1)^j \overset{\text{(ii)}}{=} \sum_{j=0}^{\infty} \frac{1}{j!} \cdot (z-1)^j \sum_{i=j}^{\infty} p_i \cdot \frac{i!}{(i-j)!} = \sum_{j=0}^{\infty} (z-1)^j \cdot \frac{\mu_{[j]}}{j!},$$

where (i) is due to the binomial expansion that $z^i = \left[(z-1)+1\right]^i = \sum_{j=0}^{i} \binom{i}{j} \cdot (z-1)^j$ for all $i \geq 0$ and (ii) follows from switching the order of summations by Tonelli's theorem because all terms are non-negative. (44) then implies that

$$\mathbb{E}\left[z^X\right] \leq \sum_{j=0}^{\infty} \mu^j (z-1)^j = \frac{1}{1-\mu(z-1)}. \qquad \square$$

We now prove Lemma A.12. The proof follows Theorem 2.1 in Janson (2018), which provides a concentration inequality for summations of independent geometric random variables using Chernoff inequality. Their results can be easily extended to random variables with log-concave distributions, as we present it here.

From Lemma A.14, the moment generating function of each random variable X_i can be bounded from above by

$$\mathbb{E}\left[e^{tX_i}\right] \leq \frac{1}{1-\mu_i(e^t-1)} = \frac{e^{-t}}{(1+\mu_i)e^{-t}-\mu_i}, \; \forall\, 0 \leq t < \ln\left(1+\frac{1}{\mu_i}\right).$$

Since for all $0 \leq t < \frac{1}{1+\mu_i} \leq \ln\left(1+\frac{1}{\mu_i}\right)$, the denominator satisfies

$$(1+\mu_i)e^{-t} - \mu_i \geq (1+\mu_i)\cdot(1-t) - \mu_i = 1-(1+\mu_i)t > 0,$$

we have

$$\mathbb{E}\left[e^{tX_i}\right] \leq \frac{e^{-t}}{1-(1+\mu_i)t}, \; \forall\, 0 \leq t < \frac{1}{1+\mu_i}.$$

As a result, for all $0 \leq t < \frac{1}{1+c} \leq \min_i \frac{1}{1+\mu_i}$ we have

$$\mathbb{E}\left[e^{tX}\right] = \prod_{i=1}^{n} \mathbb{E}\left[e^{tX_i}\right] \leq e^{-nt} \prod_{i=1}^{n} \Big(1-(1+\mu_i)t\Big)^{-1}$$

because the random variables are independent. By the Chernoff inequality, for all $0 \le t < \frac{1}{1+c}$,

$$
\begin{aligned}
\mathbb{P}\Big[X \ge \lambda\mu\Big] &\le e^{-t\lambda\mu}\mathbb{E}\Big[e^{tX}\Big] \\
&\le \exp\bigg(-t\lambda\mu - tn - \sum_{i=1}^{n} \ln\Big(1-(1+\mu_i)t\Big)\bigg) \\
&\overset{\text{(a)}}{\le} \exp\bigg(-t\lambda\mu - tn - \sum_{i=1}^{n} \frac{1+\mu_i}{1+c}\ln\Big(1-(1+c)t\Big)\bigg) \\
&= \exp\bigg(-t\lambda\mu - tn - \frac{n+\mu}{1+c}\ln\Big(1-(1+c)t\Big)\bigg),
\end{aligned}
$$

where (a) is due to the fact that the function $-\ln(1-x)$ is convex on $(0,1)$ and is 0 at $x=0$, thus by Jensen's inequality,

$$
-\ln(1-x) \le -\frac{x}{y}\ln(1-y), \ \forall\, 0 \le x \le y < 1.
$$

By choosing $t = \frac{(\lambda-1)\mu}{(1+c)(\lambda\mu+n)}$ (which is optimal here), we obtain (42). $\square$

A.11 Proof of Lemma 4.8

We prove a more general result stated in Lemma A.15 that under some regularity conditions on the function $\gamma_i(\beta)$ as defined in (8), Assumption 4.1 holds. We then show Lemma 4.8 are sufficient for the assumptions on $\gamma_i(\beta)$ in Lemma A.15 to hold.

Lemma A.15. *Suppose that function $\gamma_i(\beta)$ is differentiable, and on $\beta \in [0,1]$, is strongly concave with parameter $\ell_i > 0$ and has Lipschitz continuous gradient with parameter $L_i > 0$. Further assume that $n\ell_i \ge \bar{\ell}$, $nL_i \le \bar{L}$, and $q_{i0}, q_{0i} \le \frac{\bar{q}}{n}$ for all $i \in [n]$ and some positive constants $\bar{\ell}$, $\bar{L}$ and $\bar{q}$. Then $\lambda^*(\delta) \ge \frac{\bar{\lambda}}{n}$ for all $\delta \ge 0$ and some constants $\bar{\lambda} > 0$ and Assumption 4.1 holds.*

We prove Lemma A.15 in Appendix A.11.1; an overview of the key steps of the proof is as follows. Letting $p_i(x)$ be the optimal probabilities to (6), we can lower bound the ratio $\beta(x) = p_i(x+1)/p_i(x)$ through the first-order optimality conditions of (7) and show that these ratios are close to one for a large number of states x if the dual variable λ is sufficiently small. This implies that the expected number of resources at every spoke grows unbounded as λ goes to zero. Since the total number of resources at the spokes with $\lambda = \lambda^*(\delta)$ is no larger than $m - \delta$, $\lambda^*(\delta)$ cannot be too small, which in turn implies that the spoke resources are uniformly bounded.

We now show Lemma 4.8 provides sufficient conditions for the assumptions on $\gamma_i(\beta)$ in Lemma A.15 to hold based on the primitives of the problem, thus finished the proof.

Let $d_{i0}(\beta)$ denote an optimal solution to $\gamma_i(\beta)$. We can set $d_{i0}(0)$ arbitrarily because it does not affect the objective when $\beta = 0$. When $\beta > 0$, we can express d_{0i} in terms of d_{i0} and rewrite (8) as

$$
\gamma_i(\beta) = \max_{d_{i0} \in [0, 1 \wedge \frac{q_{0i}}{q_{i0}} \cdot \frac{1}{\beta}]} q_{0i} \cdot r_{0i}\Big(\beta \cdot \frac{q_{i0}}{q_{0i}} \cdot d_{i0}\Big) + \beta \cdot q_{i0} \cdot r_{i0}(d_{i0}). \tag{45}
$$

Since the revenue functions $r_{i0}(d)$ and $r_{0i}(d)$ are strictly concave, $d_{i0}(\beta)$ is unique for any $\beta > 0$. Let

$$
f(\beta, d) = q_{0i} \cdot r_{0i}\Big(\beta \cdot \frac{q_{i0}}{q_{0i}} \cdot d\Big) + \beta \cdot q_{i0} \cdot r_{i0}(d)
$$

be the objective of (45). $f(\beta,d)$ is concave in d with partial derivative

$$\frac{\partial f(\beta,d)}{\partial d} = \beta\cdot q_{i0}\cdot\bigg(r'_{0i}\Big(\beta\cdot\frac{q_{i0}}{q_{0i}}\cdot d\Big)+r'_{i0}(d)\bigg)$$

decreasing in d. Since $\frac{\partial f(\beta,d)}{\partial d}\big|_{d=0}=\beta\cdot q_{i0}\cdot\big(r'_{0i}(0)+r'_{i0}(0)\big)>0$, the optimal solution satisfies $d_{i0}(\beta)>0$ when $\beta>0$. $d_{i0}(\beta)$ may equal to the right end-point $\min\big\{1,\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}\big\}$. In the following, we study $\gamma_i(\beta)$ depending on whether the optimal solution is interior or one of the upper boundaries is binding.

Case 1 We have $d_{i0}(\beta)=1$ if $1\le\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}$ and $\frac{\partial f(\beta,d)}{\partial d}\big|_{d=1}\ge 0$, which is equivalent to $\beta\le\frac{q_{0i}}{q_{i0}}$ and $r'_{0i}\Big(\beta\cdot\frac{q_{i0}}{q_{0i}}\Big)+r'_{i0}(1)\ge 0$. Since $r_{i0}(d)$ and $r_{0i}(d)$ are strictly concave, there exists some $\underline{\beta}>0$ such that $r'_{0i}\Big(\underline{\beta}\cdot\frac{q_{i0}}{q_{0i}}\Big)+r'_{i0}(1)=0$ if and only if $r'_{0i}(0)+r'_{i0}(1)>0$. If it is the case, since $r'_{0i}(1)+r'_{i0}(1)<0$, we must have $\underline{\beta}\le\frac{q_{0i}}{q_{i0}}$. Moreover, $r'_{0i}\Big(\beta\cdot\frac{q_{i0}}{q_{0i}}\Big)+r'_{i0}(1)\ge 0$ for all $\beta\le\underline{\beta}$, and thus $d_{i0}(\beta)=1$ when $0<\beta\le\underline{\beta}$. If $r'_{0i}(0)+r'_{i0}(1)\le 0$, we set $\underline{\beta}=0$. When $\beta\in[0,\underline{\beta}]$, since $d_{i0}(\beta)=1$, we have

$$\gamma_i(\beta)=f\big(\beta,1\big)=q_{0i}\cdot r_{0i}\bigg(\beta\cdot\frac{q_{i0}}{q_{0i}}\bigg)+\beta\cdot q_{i0}\cdot r_{i0}(1),$$

which is concave and differentiable in β. The derivative

$$\gamma'_i(\beta)=q_{i0}\cdot\bigg(r'_{0i}\Big(\beta\cdot\frac{q_{i0}}{q_{0i}}\Big)+r_{i0}(1)\bigg)$$

is continuous on $[0,\underline{\beta}]$ because $r_{i0}(d)$ and $r_{0i}(d)$ are twice differentiable. This implies that if $\underline{\beta}>0$,

$$\gamma'_i(\underline{\beta}-)=q_{i0}\cdot\bigg(r'_{0i}\Big(\underline{\beta}\cdot\frac{q_{i0}}{q_{0i}}\Big)+r_{i0}(1)\bigg). \tag{46}$$

The second-order derivative satisfies

$$-\gamma''_i(\beta)=-\frac{q_{i0}^2}{q_{0i}}r''_{0i}(\beta)\in\bigg[\frac{\underline{q}^2}{\bar{q}}\frac{\bar{u}}{n},\frac{\bar{q}^2}{\underline{q}}\frac{\bar{U}}{n}\bigg],\ \forall\,\beta\in[0,\underline{\beta}]. \tag{47}$$

Case 2 $d_{i0}(\beta)=\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}$ if $\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}\le 1$ and $\frac{\partial f(\beta,d)}{\partial d}\big|_{d=\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}}\ge 0$, which is equivalent to $\beta\ge\frac{q_{0i}}{q_{i0}}$ and $r'_{0i}(1)+r'_{i0}\Big(\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}\Big)\ge 0$. Since $r_{i0}(d)$ and $r_{0i}(d)$ are strictly concave, there exists some $\bar{\beta}>0$ such that $r'_{0i}(1)+r'_{i0}\Big(\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}\Big)=0$ if and only if $r'_{0i}(1)+r'_{i0}(0)>0$. In this case, $\bar{\beta}$ has to be larger than $\frac{q_{0i}}{q_{i0}}$ and $r'_{0i}(1)+r'_{i0}\Big(\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}\Big)\ge 0$ for all $\beta\ge\bar{\beta}$. Thus, $d_{i0}(\beta)=\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}$ when $\beta\ge\bar{\beta}$. If $r'_{0i}(1)+r'_{i0}(0)\le 0$, we set $\bar{\beta}=\infty$. When $\beta\in[\bar{\beta},\infty)$, since $d_{i0}(\beta)=\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}$, we have

$$\gamma_i(\beta)=f\Big(\beta,d_{i0}(\beta)\Big)=q_{0i}\cdot r_{0i}(1)+\beta\cdot q_{i0}\cdot r_{i0}\Big(\frac{q_{0i}}{q_{i0}}\cdot\frac{1}{\beta}\Big).$$

Hence,

$$\gamma_i'(\beta) = q_{i0} \cdot r_{i0}\Big(\frac{q_{0i}}{q_{i0}} \cdot \frac{1}{\beta}\Big) - \frac{1}{\beta} \cdot q_{0i} \cdot r_{i0}'\Big(\frac{q_{0i}}{q_{i0}} \cdot \frac{1}{\beta}\Big),$$

and

$$\gamma_i''(\beta) = \frac{1}{\beta^3} \cdot \frac{q_{0i}^2}{q_{i0}} r_{0i}''\Big(\frac{q_{0i}}{q_{i0}} \cdot \frac{1}{\beta}\Big) < 0.$$

The first-order derivative is continuous on $[\bar{\beta}, \infty)$ because $r_{i0}(d)$ and $r_{0i}(d)$ are twice differentiable. Thus if $\bar{\beta} < \infty$,

$$\gamma_i'(\bar{\beta}+) = q_{i0} \cdot r_{i0}\Big(\frac{q_{0i}}{q_{i0}} \cdot \frac{1}{\bar{\beta}}\Big) - \frac{1}{\bar{\beta}} \cdot q_{0i} \cdot r_{i0}'\Big(\frac{q_{0i}}{q_{i0}} \cdot \frac{1}{\bar{\beta}}\Big) = q_{i0} \cdot \Bigg(r_{i0}\Big(d_{i0}(\bar{\beta})\Big) - d_{i0}(\bar{\beta}) \cdot r_{i0}'\Big(d_{i0}(\bar{\beta})\Big)\Bigg). \tag{48}$$

If $\bar{\beta} < 1$, since $\bar{\beta} \geq \frac{q_{0i}}{q_{i0}}$, the second-order derivative satisfies

$$-\gamma_i''(\beta) = -\frac{1}{\beta^3} \cdot \frac{q_{0i}^2}{q_{i0}} r_{0i}''\Big(\frac{q_{0i}}{q_{i0}} \cdot \frac{1}{\beta}\Big) \in \left[\frac{\underline{q}^2}{\bar{q}} \frac{\bar{u}}{n}, \frac{\bar{q}^2}{\underline{q}} \frac{\bar{U}}{n}\right], \ \forall\, \beta \in [\bar{\beta}, 1]. \tag{49}$$

Case 3 When $\beta \in [\underline{\beta}, \bar{\beta}]$, $d_{i0}(\beta)$ satisfies

$$\frac{\partial f(\beta, d)}{\partial d}\Big|_{d=d_{i0}(\beta)} = \beta \cdot q_{i0} \cdot \Bigg(r_{0i}'\Big(\beta \cdot \frac{q_{i0}}{q_{0i}} \cdot d_{i0}(\beta)\Big) + r_{i0}'\Big(d_{i0}(\beta)\Big)\Bigg) = 0.$$

Thus,

$$r_{0i}'\Big(\beta \cdot \frac{q_{i0}}{q_{0i}} \cdot d_{i0}(\beta)\Big) + r_{i0}'\Big(d_{i0}(\beta)\Big) = 0. \tag{50}$$

From (50) and the implicit function theorem, we have

$$\frac{d_{i0}(\beta)}{d\beta} = -\frac{q_{i0} \cdot d_{i0}(\beta) \cdot r_{0i}''\big(\beta \cdot \frac{q_{i0}}{q_{0i}} \cdot d_{i0}(\beta)\big)}{q_{i0} \cdot \beta \cdot r_{0i}''\big(\beta \cdot \frac{q_{i0}}{q_{0i}} \cdot d_{i0}(\beta)\big) + q_{0i} \cdot r_{i0}''\big(d_{i0}(\beta)\big)} \leq 0, \tag{51}$$

thus $d_{i0}(\beta)$ is decreasing in β.

Let us first assume $\bar{\beta} < \infty$ and consider the function $f(\beta, d)$ over $[\underline{\beta}, \bar{\beta}] \times [0, 1]$; since d in $f(\beta, d)$ is restricted to be in $\big[0, 1 \wedge \frac{q_{0i}}{q_{i0}} \cdot \frac{1}{\beta}\big]$, we extend $r_{0i}(d)$ smoothly over $[1, \infty]$ making sure $f(\beta, d)$ is well-defined on the support $[\underline{\beta}, \bar{\beta}] \times [0, 1]$. Note that in the extension we ensure that $r_{0i}(d)$ is strongly concave and has a Lipschitz continuous gradient with the same parameters. Since the partial derivative of $f(\beta, d)$ with respect to β is

$$\frac{\partial f(\beta, d)}{\partial \beta} = q_{i0} \cdot \Bigg(d \cdot r_{0i}'\Big(\beta \cdot \frac{q_{i0}}{q_{0i}} \cdot d\Big) + r_{i0}(d)\Bigg),$$

which is continuous in (β, d), by the envelope theorem, especially Corollary 4 in Milgrom and Segal

(2002), we have

$$\gamma_i'(\beta) = \frac{\partial f(\beta,d)}{\partial \beta}\Big|_{d=d_{i0}(\beta)} = q_{i0}\cdot\left(d_{i0}(\beta)\cdot r_{0i}'\Big(\beta\cdot\frac{q_{i0}}{q_{0i}}\cdot d_{i0}(\beta)\Big) + r_{i0}\Big(d_{i0}(\beta)\Big)\right)$$
$$\stackrel{\text{(i)}}{=} q_{i0}\cdot\left(r_{i0}\Big(d_{i0}(\beta)\Big) - d_{i0}(\beta)\cdot r_{i0}'\Big(d_{i0}(\beta)\Big)\right) \tag{52}$$

on $\beta\in(\underline{\beta},\bar{\beta})$, where (i) is due to the first-order condition (50). From (51) and (52), the second-order derivative is

$$\gamma_i''(\beta) = \frac{q_{i0}^2\cdot\big(d_{i0}(\beta)\big)^2\cdot r_{i0}''\big(d_{i0}(\beta)\big)\cdot r_{0i}''\big(\beta\cdot\frac{q_{i0}}{q_{0i}}\cdot d_{i0}(\beta)\big)}{q_{i0}\cdot\beta\cdot r_{0i}''\big(\beta\cdot\frac{q_{i0}}{q_{0i}}\cdot d_{i0}(\beta)\big) + q_{0i}\cdot r_{i0}''\big(d_{i0}(\beta)\big)} \le 0,\ \forall\ \underline{\beta}\le\beta\le\bar{\beta}, \tag{53}$$

thus $\gamma(\beta)$ is concave on $[\underline{\beta},\bar{\beta}]$. If $\underline{\beta}>0$, from Corollary 4 in Milgrom and Segal (2002), we have

$$\gamma_i'(\underline{\beta}+) = \frac{\partial f(\beta,d)}{\partial\beta}\Big|_{d=d_{i0}(\underline{\beta})=1} = q_{i0}\cdot\left(r_{0i}'\left(\underline{\beta}\cdot\frac{q_{i0}}{q_{0i}}\right) + r_{i0}(1)\right) \stackrel{\text{(a)}}{=} \gamma'(\underline{\beta}-),$$

where (a) is due to (46); hence $\gamma(\beta)$ is differentiable at $\beta=\underline{\beta}$. Analogously, if $\bar{\beta}<\infty$, again from Corollary 4 in Milgrom and Segal (2002), we have

$$\gamma_i'(\bar{\beta}-) = \frac{\partial f(\beta,d)}{\partial\beta}\Big|_{d=d_{i0}(\bar{\beta})} \stackrel{\text{(b)}}{=} q_{i0}\cdot\left(r_{i0}\Big(d_{i0}(\bar{\beta})\Big) - d_{i0}(\bar{\beta})\cdot r_{i0}'\Big(d_{i0}(\bar{\beta})\Big)\right) \stackrel{\text{(c)}}{=} \gamma'(\bar{\beta}+),$$

where (b) is from (52) and (c) is from (48); thus $\gamma(\beta)$ is differentiable at $\beta=\bar{\beta}$ as well. Combining three segments together, we know $\gamma(\beta)$ is differentiable everywhere. Note that we assume $\bar{\beta}<\infty$ in case 3. If this is not the case, following the same argument, we can show $\gamma_i(\beta)$ is differentiable in any bounded interval $[0,M]$ with $M>0$; thus again, $\gamma_i(\beta)$ is differentiable on $\mathbb{R}_+$.

Finally, from (50) and the mean value theorem we have

$$r_{0i}'(0) + r_{0i}''(\varepsilon_1)\cdot\beta\cdot\frac{q_{i0}}{q_{0i}}\cdot d_{i0}(\beta) + r_{i0}'(0) + r_{0i}''(\varepsilon_2)\cdot d_{i0}(\beta) = 0$$

for some $\varepsilon_1\in[0,\beta\cdot\frac{q_{i0}}{q_{0i}}\cdot d_{i0}(\beta)]$ and $\varepsilon_2\in[0,d_{i0}(\beta)]$. This implies that

$$d_{i0}(\beta) = \frac{-r_{0i}'(0) - r_{i0}'(0)}{r_{0i}''(\varepsilon_1)\cdot\beta\cdot\frac{q_{i0}}{q_{0i}} + r_{0i}''(\varepsilon_2)} \ge \frac{r_{0i}'(0)+r_{i0}'(0)}{\bar{U}}\cdot\frac{\underline{q}}{\beta\cdot\bar{q}+\underline{q}} \ge \frac{\bar{u}}{\bar{U}}\cdot\frac{\underline{q}}{\beta\cdot\bar{q}+\underline{q}}, \tag{54}$$

where in the last inequality we use the facts that $r_{ij}(1)\ge 0$, $r_{ij}(0)=0$, and $r_{ij}(1)\le r_{ij}(0)+r_{ij}'(0)\cdot(1-0)-\frac{(1-0)^2}{2}\bar{u}$ by strong concavity of $r_{ij}(d)$. Thus if $\underline{\beta}<1$, on the interval $[\underline{\beta},1\wedge\bar{\beta}]$, (53) and (54) implies that the second-order derivative satisfies

$$-\gamma_i''(\beta)\in\left[\frac{1}{2n}\frac{\bar{u}^3}{\bar{U}^2}\frac{\underline{q}^4}{\bar{q}(\bar{q}+\underline{q})^2},\frac{\bar{q}^2}{\underline{q}}\frac{\bar{U}}{n}\right],\ \forall\ \beta\in[\underline{\beta},1\wedge\bar{\beta}]. \tag{55}$$

From (47), (49), (55), the monotonicity of $\gamma_i'(\beta)$, and the mean value theorem on $\gamma_i'(\beta)$, $\gamma_i(\beta)$ is ℓ_i-strongly concave and has a L_i-Lipschitz continuous gradient on $\beta\in[0,1]$ with some constants ℓ_i

and L_i that satisfy $\frac{1}{2n}\frac{\bar{u}^3}{\bar{U}^2}\frac{\underline{q}^4}{\bar{q}(\bar{q}+\underline{q})^2} \le \ell_i \le L_i \le \frac{\bar{q}^2}{\underline{q}}\frac{\bar{U}}{n}$.

A.11.1 Proof of Lemma A.15

We let $\mathbb{E}\big[\tilde{X}_i^\lambda\big]$ be the expected number of resources in the spoke problem with $\lambda \ge 0$. Lemma A.16 shows that $\mathbb{E}\big[\tilde{X}_i^\lambda\big]$ can be arbitrarily large by choosing a sufficiently small λ.

Lemma A.16. *Suppose the assumptions on $\gamma_i(\beta)$ and q_{ij} in Lemma A.15 hold. Then for any $\rho > 0$, there exists a constant $c(\rho)$ such that $\mathbb{E}\big[\tilde{X}_i^\lambda\big] \ge \rho$ for $\lambda = \frac{c(\rho)}{n}$, all spokes $i \in [n]$ and m large enough.*

We prove Lemma A.16 in Appendix A.11.2. Since the perturbed problem (11) is a convex program in λ and the optimal probability distribution to each spoke problem with a given λ is unique by Proposition B.8, the objective $\bar{V}^\lambda - \delta\lambda$ is differentiable in λ with derivative $\frac{\partial(\bar{V}^\lambda - \delta\lambda)}{\partial\lambda} = (m-\delta) - \sum_{i\in[n]}\mathbb{E}\big[\tilde{X}_i^\lambda\big]$. From Lemma A.16, there exists a constant $\bar{\lambda}$ such that $\mathbb{E}\big[\tilde{X}_i^\lambda\big] \ge \frac{2m}{n}$ for $\lambda = \frac{\bar{\lambda}}{n}$ and all $i \in [n]$. Thus, the derivative at $\lambda = \frac{\bar{\lambda}}{n}$ is negative and as a result $\lambda^*(\delta) \ge \frac{\bar{\lambda}}{n}$. Finally, since $0 \le h_i^{\lambda^*(\delta)} \le (q_{i0}+q_{0i})\cdot\bar{r} - \lambda^*(\delta)\cdot\mathbb{E}\big[\tilde{X}_i(\delta)\big]$, we have $\mathbb{E}\big[\tilde{X}_i(\delta)\big] \le \bar{r}\cdot\frac{2\bar{q}}{\bar{\lambda}}$ for all spokes $i \in [n]$.

A.11.2 Proof of Lemma A.16

For ease of notation, in the proof, we consider a spoke i and drop the subscript i for $\gamma_i(\beta)$ by letting $\gamma(\beta) \triangleq \gamma_i(\beta)$; we drop the subscript i for ℓ_i and L_i as well. By assumption, $\gamma(\beta)$ is differentiable. Moreover, from Lemma A.3, $\gamma(\beta)$ is strictly concave; thus, the derivative $\gamma'(\beta)$ is strictly decreasing and the inverse $(\gamma')^{-1}(e)$ exists and strictly decreases as well. Finally, by assumption, on $\beta \in [0,1]$ $\gamma(\beta)$ is strongly concave with a constant $\ell > 0$, i.e.,

$$\gamma(\beta') \le \gamma(\beta) + \gamma'(\beta)\cdot\big(\beta'-\beta\big) - \frac{\ell}{2}\cdot\big(\beta'-\beta\big)^2,\ \forall\, \beta,\beta' \in [0,1], \tag{56}$$

and has a Lipschitz continuous gradient with a constant $L > 0$, i.e.,

$$|\gamma'(\beta) - \gamma'(\beta')| \le L\cdot|\beta-\beta'|,\ \forall\, \beta,\beta' \in [0,1]. \tag{57}$$

(56) and (57) imply that $\ell \le L$. We first construct lower bounds on the probability ratios $\beta_x = \frac{p_i(x+1)}{p_i(x)}$ in Lemma A.19 with $p_i(x)$ being the optimal probability distribution to the spoke problem (6). Lemmas A.17 and A.18 serve as preliminary results.

Lemma A.17. *Let $g(\beta,y) = (\gamma')^{-1}\Big(z(\beta)+\gamma'(0)-\lambda y\Big)$ with $\lambda > 0$ and $z(\beta) = \beta\cdot\gamma'(\beta)-\gamma(\beta)$ as in Lemma A.5. Let $\bar{y} = \frac{\gamma'(0)-\gamma(1)}{\lambda}$ and $\beta_y^* = \inf\big\{\beta \ge 0 : g(\beta,y) \le \beta\big\}$. Then,*

1. *$g(\beta,y)$ is strictly increasing in β and y;*
2. *$\beta_0^* = g(\beta_0^*,0) = 0$, $\bar{y} > 0$ and $\beta_{\bar{y}}^* = g(\beta_{\bar{y}}^*,\bar{y}) = 1$;*
3. *for $y \in [0,\bar{y}]$, $0 \le \beta_y^* \le 1$, $g(\beta_y^*,y) = \beta_y^*$, and β_y^* is increasing in y;*
4. *for $\beta \in [0,\beta_y^*]$, $g(\beta,y) \ge \alpha_y\beta + (1-\alpha_y)\beta_y^*$ with $\alpha_y = \big(1-\frac{\ell}{L}\big)+\frac{\ell}{L}\beta_y^*$.*

Proof. Part 1 is because both $(\gamma')^{-1}(e)$ and $z(\beta)$ are strictly decreasing and $\lambda > 0$. For part 2, since $g(0,0) = 0$, we have $\beta_0^* = g(\beta_0^*,0) = 0$. Moreover, since $\gamma(0) = 0$ and $\gamma(1) < \gamma(0)+\gamma'(0)\cdot(1-0)$ by strict concavity of $\gamma(\beta)$, $\gamma'(0) > \gamma(1)$ and hence $\bar{y} > 0$. Finally, let $\tilde{g}(\beta) \triangleq g(\beta,\bar{y}) = (\gamma')^{-1}\big(\gamma(1)+$

$z(\beta)\big)$. We show $\tilde{g}(1) = 1$ and $\tilde{g}(\beta) > \beta$ for all $\beta \in \mathbb{R}_+$ and $\beta \neq 1$; thus $\beta^*_{\bar{y}} = g(\beta^*_{\bar{y}}, \bar{y}) = 1$. To see this, first, note that $\tilde{g}(1) = (\gamma')^{-1}\big(\gamma(1) + z(1)\big) = (\gamma')^{-1}\big(\gamma'(1)\big) = 1$. Second, since $\gamma(\beta)$ is strictly concave, for any $\beta \neq 1$ we have

$$\gamma(1) < \gamma(\beta) + \gamma'(\beta) \cdot (1 - \beta) = \gamma'(\beta) - z(\beta).$$

Thus, $\gamma'\big(\tilde{g}(\beta)\big) = \gamma(1) + z(\beta) < \gamma'(\beta)$. Since $\gamma'(\beta)$ is strictly decreasing, we have $\tilde{g}(\beta) > \beta$.

Part 3: we already proved this for $y = 0$ and $y = \bar{y}$ in part 2. We now show that for any $y \in (0, \bar{y})$, $0 \leq \beta^*_y \leq 1$ and $g(\beta^*_y, y) = \beta^*_y$. First, $g(0, y) > g(0, 0) = 0$ and $g(1, y) < g(1, \bar{y}) = 1$ from part 1. Second, note that $g(\beta, y)$ is jointly continuous in (β, y) for $\beta \in [0, 1]$ and $y \in [0, \bar{y}]$. To see this, since $\gamma(\beta)$ has Lipschitz continuous gradient on $\beta \in [0, 1]$, $\gamma'(\beta)$ is continuous on $[0, 1]$ and $(\gamma')^{-1}(e)$ is continuous on $[\gamma'(1), \gamma'(0)]$. The continuity of $g(\beta, y)$ then follows from the definition of $z(\beta)$ and the fact that $g(\beta, y) \in [0, 1]$ when $\beta \in [0, 1]$ and $y \in [0, \bar{y}]$. Now let $h(\beta) = g(\beta, y) - \beta$. From above we know $h(0) > 0$, $h(1) < 0$, and $h(\beta)$ is continuous on $\beta \in [0, 1]$. The intermediate value theorem then implies the existence of a point $\beta^*_y \in [0, 1]$ that satisfies $g(\beta^*_y, \bar{y}) = \beta^*_y$. Finally, the monotonicity of β^*_y follows from the fact that $g(\beta, y)$ is increasing in y.

Part 4: from $\beta^*_y = g(\beta^*_y, y)$, we have $\gamma'(\beta^*_y) = z(\beta^*_y) + \gamma'(0) - \lambda y$. Thus,

$$g(\beta, y) = (\gamma')^{-1}\Big(z(\beta) + \gamma'(0) - \lambda y\Big) = (\gamma')^{-1}\Big(z(\beta) - z(\beta^*_y) + \gamma'(\beta^*_y)\Big).$$

Since $\gamma'(\beta)$ is strictly decreasing in β, $g(\beta, y) \geq \alpha_y \beta + (1 - \alpha_y)\beta^*_y$ if and only if

$$z(\beta) - z(\beta^*_y) + \gamma'(\beta^*_y) \leq \gamma'\Big(\alpha_y \beta + (1 - \alpha_y)\beta^*_y\Big).$$

Letting $h(\beta) = \gamma'(\beta) - z(\beta)$, above is equivalent to

$$\gamma'(\beta) - \gamma'\Big(\alpha_y \beta + (1 - \alpha_y)\beta^*_y\Big) \leq h(\beta) - h(\beta^*_y). \tag{58}$$

Since $\gamma(\beta)$ has Lipschitz continuous gradient on $[0, 1]$ with $L > 0$, the left-hand side is no larger than $L(1 - \alpha_y)(\beta^*_y - \beta)$. On the other hand, the right-hand side satisfies

$$\begin{aligned}
h(\beta) - h(\beta^*_y) &= (1 - \beta)\gamma'(\beta) - (1 - \beta^*_y)\gamma'(\beta^*_y) + \gamma(\beta) - \gamma(\beta^*_y) \\
&\geq (1 - \beta)\gamma'(\beta) - (1 - \beta^*_y)\gamma'(\beta^*_y) + \gamma'(\beta)(\beta - \beta^*_y) \\
&= (1 - \beta^*_y)(\gamma'(\beta) - \gamma'(\beta^*_y)) \\
&\geq \ell(1 - \beta^*_y)(\beta^*_y - \beta),
\end{aligned}$$

where the first inequality is from $\gamma(\beta) + \gamma'(\beta)(\beta^*_y - \beta) \geq \gamma(\beta^*_y)$ because $\gamma(\beta)$ is concave, and the second inequality is from the ℓ-strong concavity of $\gamma(\beta)$ on $[0, 1]$. Finally, (58) holds because $\ell(1 - \beta^*_y)(\beta^*_y - \beta) = L(1 - \alpha_y)(\beta^*_y - \beta)$ by the choice of α_y. □

In Lemma A.18, we provide lower and upper bounds on β^*_y for $y \in [0, \bar{y}]$.

Lemma A.18. *β^*_y defined in Lemma A.17 satisfies the following for $y \in [0, \bar{y}]$.*

1. $1 - \sqrt{\frac{2\lambda}{l}(\bar{y} - y)} \leq \beta^*_y \leq 1 - \sqrt{\frac{2\lambda}{L}(\bar{y} - y)}$;
2. $1 - \sqrt{1 - \frac{2\lambda}{L} y} \leq \beta^*_y \leq 1 - \sqrt{1 - \frac{2\lambda}{l} y}$, *where the second inequality holds for* $y \leq \frac{l}{2\lambda}$.

Proof. We first prove inequality (59) as a preparation.

$$\frac{\lambda}{L} \cdot \frac{1}{1-\beta_y^*} \leq \frac{d\beta_y^*}{dy} \leq \frac{\lambda}{\ell} \cdot \frac{1}{1-\beta_y^*},\ \forall\, 0 \leq \beta_y^* \leq 1. \tag{59}$$

To see this, note that $\gamma'(\beta_y^*) = z(\beta_y^*) + \gamma'(0) - \lambda y$ from $\beta_y^* = g(\beta_y^*, y)$ in Lemma A.17 part 2; thus letting $h(\beta) = \gamma'(\beta) - z(\beta)$, we have $h(\beta_y^*) = \gamma'(0) - \lambda y$. For any $\beta, \beta + \Delta\beta \in [0, 1]$,

$$\begin{aligned} h(\beta + \Delta\beta) - h(\beta) &= \gamma'(\beta + \Delta\beta) \cdot (1 - \beta - \Delta\beta) + \gamma(\beta + \Delta\beta) - \gamma'(\beta) \cdot (1 - \beta) - \gamma(\beta) \\ &= \gamma'(\beta + \Delta\beta) \cdot (1 - \beta - \Delta\beta) - \gamma'(\beta) \cdot (1 - \beta) + \gamma'(\beta + \varepsilon\Delta\beta) \cdot \Delta\beta \\ &= \Big(\gamma'(\beta + \Delta\beta) - \gamma'(\beta)\Big) \cdot (1 - \beta) + \Delta\beta \cdot \Big(\gamma'(\beta + \varepsilon\Delta\beta) - \gamma'(\beta + \Delta\beta)\Big), \end{aligned} \tag{60}$$

where the second equation is because $\gamma(\beta + \Delta\beta) = \gamma(\beta) + \gamma'(\beta + \varepsilon\Delta\beta) \cdot \Delta\beta$ for some $\varepsilon \in [0, 1]$ by the mean value theorem. For any $y, y + \Delta y \in [0, \bar{y}]$, let $\Delta\beta_y^* = \beta_{y+\Delta y}^* - \beta_y^*$; $\Delta\beta_y^* > 0$ if and only if $\Delta y > 0$ because β_y^* is increasing in y by Lemma A.17 part 3. Since $h(\beta_y^*) = \gamma'(0) - \lambda y$ for all $y \in [0, \bar{y}]$, we have

$$\lambda\Delta y = \lambda(y + \Delta y) - \lambda y = h(\beta_y^*) - h(\beta_{y+\Delta y}^*) = h(\beta_y^*) - h(\beta_y^* + \Delta\beta_y^*). \tag{61}$$

Combining (60) and (61) with the fact that $\gamma(\beta)$ is ℓ-strongly concave and has L-Lipschitz continuous gradient on $[0, 1]$ gives

$$\frac{\ell}{\lambda} \cdot \Big(1 - \beta_y^*\Big) + O\big(\Delta\beta_y^*\big) \leq \frac{\Delta y}{\Delta\beta_y^*} \leq \frac{L}{\lambda} \cdot \Big(1 - \beta_y^*\Big) + O\big(\Delta\beta_y^*\big).$$

Letting $\Delta\beta_y^*$ go to zero and rearranging gives (59).

Part 1: since $\beta_y^* \leq 1$, from (59) we have $(1 - \beta_y^*) \cdot d\beta_y^* \leq \frac{\lambda}{\ell} \cdot dy$. Since β_y^* is increasing in y, integrating both sides over $[y, \bar{y}]$ gives

$$\int_{\beta_y^*}^{\beta_{\bar{y}}^* = 1} (1 - \beta_y^*) \cdot d\beta_y^* = \frac{(1 - \beta_y^*)^2}{2} \leq \int_y^{\bar{y}} \frac{\lambda}{\ell} \cdot dy = \frac{\lambda}{\ell} \cdot (\bar{y} - y).$$

Rearranging gives the lower bound. A similar analysis applied to the first inequality in (59) yields the upper bound.

Part 2: proof is analogous to part 1 by integrating over $[0, y]$. □

Now we are ready to provide lower bounds on the ratios of the optimal probabilities to (6) in Lemma A.19.

Lemma A.19. *Let $\tilde{\beta}_1 = 0$ and consider $\tilde{\beta}_{y+1} = g(\tilde{\beta}_y, y)$. Let $\beta_x = \frac{p_i(x+1)}{p_i(x)}$ for $x \in [0 : m^* - 1]$ be the ratio of the optimal probabilities to the spoke problem (6), with m^* being the end point of the support of $p_i(x)$ as defined in Lemma A.7. Then,*

1. *$\tilde{\beta}_y$ is increasing in y;*
2. *$\tilde{\beta}_{y+1} \leq \beta_y^*$ for all integers $y \leq \bar{y}$;*
3. *$\tilde{\beta}_y \leq \beta_{m^*-y}$ for all $y \in [1 : m^*]$ when $m \geq \frac{\gamma'(0)}{\lambda}$;*
4. *$m^* \geq \bar{y} - 1$ when $m \geq \frac{\gamma'(0)}{\lambda}$.*

Proof. We prove parts 1 to 3 by induction. For part 1, as a base case, let $\tilde{\beta}_0 = 0$; it is easy to see that $\tilde{\beta}_1 = g(\tilde{\beta}_0, 0) = 0$ satisfies the iteration. By induction, $\tilde{\beta}_{y+1} = g(\tilde{\beta}_y, y) \geq g(\tilde{\beta}_{y-1}, y-1) = \tilde{\beta}_y$ because $g(\beta, y)$ is increasing in both β and y.

Part 2: as a base case, we have $\tilde{\beta}_1 = 0 \leq \beta_0^* = 0$. Now suppose $\tilde{\beta}_y \leq \beta_{y-1}^*$. Since β_y^* increases in y by Lemma A.17 part 3, $\tilde{\beta}_y \leq \beta_y^*$ and thus $\tilde{\beta}_{y+1} = g(\tilde{\beta}_y, y) \leq g(\beta_y^*, y) = \beta_y^*$ by monotonicity of $g(\beta, y)$.

Part 3: when $m \geq \frac{\gamma'(0)}{\lambda}$, $m^* < m$ and is the unique integer satisfying[10]

$$\lambda m^* + h_i^\lambda < \gamma'(0) \leq \lambda(m^* + 1) + h_i^\lambda. \tag{62}$$

From the first-order optimality conditions as in Lemma A.7, we have

$$\begin{aligned}
&\gamma'(\beta_{m^*-1}) = \lambda m^* + h_i^\lambda, \\
&\gamma'(\beta_{x-1}) = z(\beta_x) + \lambda x + h_i^\lambda, \ \forall\, x \leq m^* - 1.
\end{aligned}$$

As a base case, we have $\beta_{m^*-1} = (\gamma')^{-1}\big(\lambda m^* + h_i^\lambda\big) \geq (\gamma')^{-1}\big(\gamma'(0)\big) = 0 = \tilde{\beta}_1$, where the inequality follows from the facts that $\gamma'(\beta)$ decreases in β and $\lambda m^* + h_i^\lambda < \gamma'(0)$ by (62). In the induction step, suppose $\beta_{m^*-y} \geq \tilde{\beta}_y$. Then

$$\beta_{m^*-y-1} = (\gamma')^{-1}\Big(z(\beta_{m^*-y}) + h_i^\lambda + \lambda m^* - \lambda y\Big) \geq (\gamma')^{-1}\Big(z(\tilde{\beta}_y) + \gamma'(0) - \lambda y\Big) = \tilde{\beta}_{y+1},$$

where the inequality follows from the facts that $\gamma'(\beta)$ decreases in β, $\lambda m^* + h_i^\lambda < \gamma'(0)$ by (62), $z(\beta)$ decreases in β by Lemma A.5, and $\beta_{m^*-y} \geq \tilde{\beta}_y$ by assumption.

Part 4: first, note that the optimal value of the spoke problem h_i^λ satisfies $h_i^\lambda \leq h_i^{\lambda=0} - \lambda \mathbb{E}\big[\tilde{X}_i^\lambda\big] \leq \gamma(1)$, where the second inequality follows from the facts that $\mathbb{E}\big[\tilde{X}_i^\lambda\big] \geq 0$ and $\gamma(1)$ is the flow relaxation to $h_i^{\lambda=0}$. From (62) we have $m^* \geq \frac{\gamma'(0) - h_i^\lambda}{\lambda} - 1 \geq \frac{\gamma'(0) - \gamma(1)}{\lambda} - 1 = \bar{y} - 1$. $\square$

As a final preparation to the proof of Lemma A.16, we provide two more lemmas. Lemma A.20 bounds the gap $e_y = \beta_y^* - \tilde{\beta}_y$ from above.

Lemma A.20. *For any integer $y \leq \bar{y}$, the gap $e_y = \beta_y^* - \tilde{\beta}_y \geq 0$ satisfies*

1. $e_y \leq \alpha_{y-1} \cdot e_{y-1} + \beta_y^* - \beta_{y-1}^* \leq \alpha_{y-1} \cdot e_{y-1} + \frac{\lambda}{\ell} \cdot \frac{1}{1-\beta_y^*}$*;*
2. *for any $\tilde{\beta} \in [0,1]$, let $\tilde{y} = \inf\big\{y \geq 0 : \beta_y^* \geq \tilde{\beta}\big\}$ and $\tilde{\alpha} = \big(1 - \frac{\ell}{L}\big) + \frac{\ell}{L}\tilde{\beta}$; then for any integer $1 \leq y \leq \tilde{y}$, we have $e_y \leq \tilde{\alpha}^{y-1} \cdot e_1 + \frac{\lambda}{\ell} \cdot \frac{1}{1-\tilde{\beta}} \cdot \frac{1}{1-\tilde{\alpha}}$.*

Proof. Part 1: from Lemma A.17 part 4 we have

$$e_y = \beta_y^* - g(\tilde{\beta}_{y-1}, y-1) \leq \beta_y^* - \alpha_{y-1}\tilde{\beta}_{y-1} - (1 - \alpha_{y-1})\beta_{y-1}^* = \alpha_{y-1} \cdot e_{y-1} + \beta_y^* - \beta_{y-1}^*.$$

Moreover, from (59) we have

$$\beta_y^* - \beta_{y-1}^* = \int_{y-1}^{y} \left(\frac{d\beta_s^*}{ds}\right) ds \leq \int_{y-1}^{y} \frac{\lambda}{\ell} \cdot \frac{1}{1-\beta_s^*} ds \leq \frac{\lambda}{\ell} \cdot \frac{1}{1-\beta_y^*}. \tag{63}$$

[10] From (62), $m^* < \frac{\gamma'(0) - h_i^\lambda}{\lambda} \leq \frac{\gamma'(0)}{\lambda} \leq m$.

Part 2: since β_y^* is increasing in y by Lemma A.17 part 3, for any integer y with $1 \le y \le \tilde{y}$, $\beta_y^* \le \beta_{\tilde{y}}^* = \tilde{\beta}$ and hence $\alpha_y = (1-\frac{\ell}{L}) + \frac{\ell}{L}\beta_y^* \le \tilde{\alpha}$ and $e_y \le \alpha_{y-1} \cdot e_{y-1} + \frac{\lambda}{\ell} \cdot \frac{1}{1-\beta_y^*} \le \tilde{\alpha} \cdot e_{y-1} + \frac{\lambda}{\ell} \cdot \frac{1}{1-\tilde{\beta}}$. We now prove the inequality by induction. First, it is trivially true with $y=1$. Now suppose $y \le \tilde{y}$ and the inequality holds at $y-1$. We have

$$e_y \le \tilde{\alpha} \cdot e_{y-1} + \frac{\lambda}{\ell} \cdot \frac{1}{1-\tilde{\beta}} \le \tilde{\alpha}^{y-1} \cdot e_1 + \frac{\lambda}{\ell} \cdot \frac{1}{1-\tilde{\beta}} \cdot \Big(\frac{\tilde{\alpha}}{1-\tilde{\alpha}} + 1\Big) = \tilde{\alpha}^{y-1} \cdot e_1 + \frac{\lambda}{\ell} \cdot \frac{1}{1-\tilde{\beta}} \cdot \frac{1}{1-\tilde{\alpha}}. \ \square$$

Lemma A.21 shows that the expected value of a discrete random variable is increasing in the ratio of adjacent probabilities.

Lemma A.21. *Consider two integer-valued random variables X_i with $i \in \{1,2\}$, each with support $a_i \le x \le b_i$ and probability mass function $g_i(x)$. If $a_1 \ge a_2$, $b_1 \ge b_2$, and $\frac{g_1(x+1)}{g_1(x)} \ge \frac{g_2(x+1)}{g_2(x)}$ for all $a_1 \le x \le b_2 - 1$, then $\mathbb{E}[X_1] \ge \mathbb{E}[X_2]$.*

Proof. It is easy to check that $g_1(x')g_2(x) \ge g_2(x')g_1(x)$ for all $x' \ge x$. Thus by Section 1.C in Shaked and Shanthikumar (2007), X_1 dominates X_2 in the monotone likelihood ratio order, and this implies that X_1 first-order stochastically dominates X_2. $\square$

Proof of Lemma A.16: We show in Section D.1 that for any $\beta \in (0,1)$, the distribution of the resources in spoke i with $\beta_x = \frac{p_i(x+1)}{p_i(x)}$ being $\beta_x = \beta$ for $0 \le x \le k-1$ and $\beta_x = 0$ for $x \ge k$ has mean $B^k(\beta) = \frac{\beta^{k+2}k - (1+k)\beta^{k+1} + \beta}{(1-\beta)(1-\beta^{k+1})}$. Moreover, $\lim_{k\to\infty} B^k(\beta) = B^\infty(\beta) \triangleq \frac{\beta}{1-\beta}$. For any $\rho > 0$, pick $\bar{\beta}$ be such that $B^\infty(\bar{\beta}) = \frac{\bar{\beta}}{1-\bar{\beta}} = 2\rho$, i.e., $\bar{\beta} = \frac{2\rho}{1+2\rho}$. Since $\lim_{k\to\infty} B^k(\bar{\beta}) = B^\infty(\bar{\beta}) = 2\rho$, there exists an integer $N_1 \in \mathbb{N}$ such that $B^k(\bar{\beta}) \ge \rho$ for all $k \ge N_1$.

Select $\tilde{\beta} \in (\bar{\beta}, 1)$ and let $\Delta = \tilde{\beta} - \bar{\beta} > 0$. Let $\bar{y} = \frac{\gamma'(0) - \gamma(1)}{\lambda}$, $\tilde{y} = \inf\{y \ge 0 : \beta_y^* \ge \tilde{\beta}\}$, and $\tilde{y}^0$ the minimum integer that is at least $\tilde{y}$. Let $m \ge \frac{\gamma'(0)}{\lambda}$ (we will specify sufficient conditions for this later), hence $m^* \ge \bar{y} - 1$ from Lemma A.19 part 4. For any integer y satisfying $0 \le y \le \bar{y} - \tilde{y}^0 - 1 \le m^* - \tilde{y}^0$, the ratio $\beta_y = \frac{p_i(y+1)}{p_i(y)}$ with $p_i(y)$ optimal to (6) satisfies $\beta_y \ge \tilde{\beta}_{m^*-y} \ge \tilde{\beta}_{\tilde{y}^0-1}$, where the first inequality is from Lemma A.19 part 3 and the second one is because $\tilde{\beta}_y$ is increasing in y by Lemma A.19 part 1. Since (a): $\tilde{\beta}_{\tilde{y}^0-1} = \beta^*_{\tilde{y}^0-1} - e_{\tilde{y}^0-1}$, (b): $\beta^*_{\tilde{y}^0-1} \ge \tilde{\beta} - \frac{\lambda}{\ell} \cdot \frac{1}{1-\tilde{\beta}}$ analogous to (63) and noting that $\beta^*_{\tilde{y}} = \tilde{\beta}$, (c): $e_{\tilde{y}^0-1}$ can be bounded from above by Lemma A.20 in terms of e_1 with $\tilde{\alpha}$ defined therein satisfying $\tilde{\alpha} \le \big(1 - \frac{\bar{\ell}}{L}\big) + \frac{\bar{\ell}}{L}\tilde{\beta} < 1$, and (d): $e_1 = \beta_1^* - \tilde{\beta}_1 = \beta_1^* \le 1 - \sqrt{1 - \frac{2\lambda}{\ell}}$ by the upper bound in Lemma A.18 part 2, we have

$$\beta_y \ge \tilde{\beta}_{\tilde{y}^0-1} \ge \tilde{\beta} - \frac{\lambda}{\ell} \cdot \frac{1}{1-\tilde{\beta}} - \left(1 - \sqrt{1 - \frac{2\lambda}{\ell}}\right) - \frac{\lambda}{\ell} \cdot \frac{1}{1-\tilde{\beta}} \cdot \frac{1}{1-\tilde{\alpha}}, \ \forall\, 0 \le y \le \bar{y} - \tilde{y}^0 - 1.$$

Since the right-hand side of above converges to $\tilde{\beta}$ when λ diminishes and recall that $\ell \ge \frac{\bar{\ell}}{n}$, there exists a constant $c_1 > 0$ such that when $\lambda \le \frac{c_1}{n}$, $\beta_y \ge \tilde{\beta} - \Delta = \bar{\beta}$ for all $0 \le y \le \bar{y} - \tilde{y}^0 - 1$. Moreover, since $\bar{y} - \tilde{y} \ge \frac{\ell}{2\lambda} \cdot (1-\tilde{\beta})^2$ from Lemma A.18 part 1, there exists a constant $c_2 > 0$ such that $\bar{y} - \tilde{y}^0 \ge N_1$ when $\lambda \le \frac{c_2}{n}$.

Combining everything together, since the mean value $\mathbb{E}\big[\tilde{X}_i^\lambda\big]$ increases in $\beta_x = \frac{p_i(x+1)}{p_i(x)}$ by Lemma A.21, we have $\mathbb{E}\big[\tilde{X}_i^\lambda\big] \ge B^{N_1}(\bar{\beta}) \ge \rho$ when $\lambda = \frac{c}{n}$ with $c = \min\{c_1, c_2\}$. To ensure that $m \ge \frac{\gamma'(0)}{\lambda}$, note that since $\gamma'(0) \le q_{i0} \cdot (\bar{r} + \bar{\omega})$ by Lemma A.4 and $q_{i0} \le \frac{\bar{q}}{n}$, it suffices to set $m \ge \frac{\bar{q}\cdot(\bar{r}+\bar{\omega})}{c}$.

A.12 Proof of Theorem 6.1

The proof is analogous to the proof of Theorem 4.1. First, the same sensitivity analysis for $V^{\mathrm{R}}(\delta)$ in Section 4.1 implies that Lemma 4.2 still holds, i.e.,

$$V^{\mathrm{R}}(\delta) \leq V^{\mathrm{R}}(0) = V^{\mathrm{R}} \leq V^{\mathrm{R}}(\delta) + \bar{r} \cdot \frac{\delta}{m-\delta}. \tag{64}$$

Second, we can bound $V^{\mathrm{R}}(\delta) - V^{\pi}(\delta)$ from above in a similar manner as in Lemma 4.3. Specifically, the same argument in Lemma A.10 implies that the difference of the continuation values $\Delta v_{i,t}(x)$ for each spoke is still bounded from above by the derivative bound $\bar{\omega}$ in Assumption 2.1. This implies that every time the Lagrangian policy differs in the two systems, the difference in continuation values is at most $\bar{r} + \bar{\omega}$ if the request is from a hub to a spoke; if the request is between hubs, the difference in continuation values is at most $\bar{r}$.

Since the Lagrangian policy takes different actions in the relaxed and original systems at the same state $(\mathbf{x}, s)$ only when $x_j = 0$ and hub j is the originating location for some $j \in [J]$, following the same proof of Lemma 4.3, we have

$$V^{\mathrm{R}}(\delta) - V^{\pi}(\delta) \leq (\bar{r} + \bar{\omega}) \cdot \sum_{j \in [J]} q_j \cdot \mathbb{P}\Big[X_j(\delta) = 0\Big], \tag{65}$$

where $q_j = \sum_{i \in [n]} q_{ji} + \sum_{j' \in [J]} q_{jj'}$ is the probability that hub j is the originating location of the request. Combining (64) and (65) gives the desired result.

B Additional Results

Lemma B.1 (Lagrangian Policy in the Spoke Problem). *For each spoke problem and using the Lagrangian policy, we have*

1. *Set I_i is the single positive recurrent class and the chain is aperiodic;*
2. *$p_i(x)$ is the unique stationary distribution;*
3. *Set I_i takes the form of $I_i = [0 : H_i]$ for some non-negative integer $0 \leq H_i \leq m$;*
4. *The Lagrangian policy is optimal to each spoke problem.*

Proof. We prove Lemma B.1 through a sequence of properties. We say a set of states is *closed* if the state remains in the set when started at a state in the set. Proposition B.2 shows the set I_i is closed with the Lagrangian policy.

Proposition B.2. *Set I_i is closed with the Lagrangian policy.*

Proof. Suppose not. Without loss of generality we assume $x \in I_i$ and $x + 1 \in I_i^c$ but $d_i(x, 0, i) > 0$. From the balance constraint in (6) we have $p_i(x+1) \cdot q_{i0} \cdot d_i(x+1, i, 0) = p_i(x) \cdot q_{0i} \cdot d_i(x, 0, i) > 0$. This implies $p_i(x+1) > 0$ and a contradiction. □

Proposition B.3. *$p_i(x)$ is a stationary distribution with the Lagrangian policy.*

Proof. This is a direct result from the balance constraint in (6):

$$p_i(x) \cdot q_{0i} \cdot d_i(x, 0, i) = p_i(x+1) \cdot q_{i0} \cdot d_i(x+1, i, 0), \ \forall\, x \in [0 : m-1]. \qquad \square$$

Proposition B.3 implies all states in I_i are positive recurrent.

Corollary B.4. *All states $x \in I_i$ are positive recurrent.*

Proof. Since $p_i(x)$ is a stationary distribution from Proposition B.3 and $p_i(x) > 0$ for all states $x \in I_i$, states $x \in I_i$ are recurrent according to Theorem 6.5.4 in Durrett (2010). Moreover, each state $x \in I_i$ is positive recurrent because the set I_i is finite. □

Proposition B.5. *States $x \in I_i^c$ are transient.*

Proof. This is due to Proposition B.2 and the construction of the Lagrangian policy for states outside the set I_i. □

Proposition B.6. *Set I_i is irreducible.*

Proof. Suppose not. Since all states in set I_i are positive recurrent (Corollary B.4), I_i must contain at least two recurrent classes. Let $C \subsetneq I_i$ be a recurrent class and a strict subset of I_i. Without loss of generality, we assume state $x \in C$ whereas state $x-1 \in I_i \backslash C$ lies in another recurrent class. Since states $x-1$ and x do not reach each other, we must have $d_i(x-1,0,i) = d_i(x,i,0) = 0$. However, by the complementary slackness properties (22)-(24), we have $d_i(x-1,0,i) = \operatorname{argmax}_{d\in[0,1]}\big\{r_{0i}(d) + d \cdot \big(v_i^\lambda(x) - v_i^\lambda(x-1)\big)\big\}$ and $d_i(x,i,0) = \operatorname{argmax}_{d\in[0,1]}\big\{r_{i0}(d) + d \cdot \big(v_i^\lambda(x-1) - v_i^\lambda(x)\big)\big\}$ with $v_i^\lambda(x)$ being the average differential value functions in (20). Since the maximum points $d_{0i}^* = \operatorname{argmax}_{d\in[0,1]} r_{0i}(d)$ and $d_{i0}^* = \operatorname{argmax}_{d\in[0,1]} r_{i0}(d)$ are unique and strictly positive by Assumption 2.1, either $d_i(x-1,0,i) \geq d_{0i}^* > 0$ or $d_i(x,i,0) \geq d_{i0}^* > 0$; thus a contradiction. □

Proposition B.7. *The Markov chain is aperiodic.*

Proof. The chain stays at the current state in every time period when the request type is neither $(i,0)$ nor $(0,i)$, with probability $1 - q_i > 0$. □

From Corollary B.4 and Propositions B.5 and B.6, set I_i is the single positive recurrent class and all states outside I_i are transient. As a result, $p_i(x)$ is the unique stationary distribution. Since I_i is irreducible by Proposition B.6 and every transition can only increase or decrease the current state by one, I_i must take the form of $I_i = [L_i : H_i]$ which incorporates a sequence of consecutive integers. We now show $L_i = 0$. If $\lambda > 0$ but $L_i > 0$, shifting the probabilities and controls to the left by L_i with $\tilde{p}_i(x) = p_i(x+L_i)$, $\tilde{d}_i(x,i,0) = d_i(x+L_i,i,0)$ and $\tilde{d}_i(x,0,i) = d_i(x+L_i,0,i)$ yields a feasible solution to (6) with a strictly better objective value, thus a contradiction. If $\lambda = 0$, Lemma A.7 implies that $I_i = [0:m]$, i.e., the optimal distribution spans the whole range.

Finally, since the Markov chain has a single positive recurrent class and $p_i(x)$ is the unique stationary distribution, the average revenue of the Lagrangian policy does not depend on the initial state and can be expressed as the objective of (6). The strong duality in Proposition 3.2 implies that the average revenue of the Lagrangian policy is equal to the optimal average revenue h_i^λ, hence the Lagrangian policy is optimal to the spoke problem.

□

Proposition B.8. *The stationary distribution $p_i(x)$ that is optimal to (6) is unique.*

Proof. Suppose not. Let distribution $p_i^a(x)$ together with controls $d_i^a(x,i,0)$ and $d_i^a(x,0,i)$ and distribution $p_i^b(x)$ together with controls $d_i^b(x,i,0)$ and $d_i^b(x,0,i)$ be two optimal solutions to (6). $p_i^a(x)$ and $p_i^b(x)$ are not identical and we denote their supports by $[0:H_i^a]$ and $[0:H_i^b]$, respectively. Without loss of generality, let $d_i^a(x,i,0) = d_i^a(x,0,i) = 0$ for all states x with $p_i^a(x) = 0$, and $d_i^b(x,i,0) = d_i^b(x,0,i) = 0$ for all states x with $p_i^b(x) = 0$.

We first show that there exists a state x such that $p_i^a(x), p_i^b(x) > 0$ and either $d_i^a(x,i,0) \neq d_i^b(x,i,0)$ or $d_i^a(x,0,i) \neq d_i^b(x,0,i)$. To see this, note that if $H_i^a \neq H_i^b$, taking $x = \min\left\{H_i^a, H_i^b\right\}$, we have $p_i^a(x), p_i^b(x) > 0$ and $d_i^a(x,0,i) \neq d_i^b(x,0,i)$. Otherwise if $H_i^a = H_i^b$, since the distributions $p_i^a(x)$ and $p_i^b(x)$ are not identical and both of them sum up to one, there must exist a state $x \in [0 : H_i^a - 1]$ such that the ratios $\frac{p_i^a(x+1)}{p_i^a(x)}$ and $\frac{p_i^b(x+1)}{p_i^b(x)}$ are not equal. Since $p_i^a(x) \cdot q_{0i} \cdot d_i^a(x,0,i) = p_i^a(x+1) \cdot q_{i0} \cdot d_i^a(x+1,i,0)$ and $p_i^b(x) \cdot q_{0i} \cdot d_i^b(x,0,i) = p_i^b(x+1) \cdot q_{i0} \cdot d_i^b(x+1,i,0)$, we have either $d_i^a(x,0,i) \neq d_i^b(x,0,i)$ or $d_i^a(x+1,i,0) \neq d_i^b(x+1,i,0)$.

For any $\alpha_1, \alpha_2 > 0$ with $\alpha_1 + \alpha_2 = 1$, let $p_i(x) = \alpha_1 \cdot p_i^a(x) + \alpha_2 \cdot p_i^b(x)$ for all states x, $d_i(x,i,0) = \frac{\alpha_1 \cdot p_i^a(x)}{p_i(x)} \cdot d_i^a(x,i,0) + \frac{\alpha_2 \cdot p_i^b(x)}{p_i(x)} \cdot d_i^b(x,i,0)$ and $d_i(x,0,i) = \frac{\alpha_1 \cdot p_i^a(x)}{p_i(x)} \cdot d_i^a(x,0,i) + \frac{\alpha_2 \cdot p_i^b(x)}{p_i(x)} \cdot d_i^b(x,0,i)$ for all states x with $p_i(x) > 0$, and $d_i(x,i,0) = d_i(x,0,i) = 0$ for all states x with $p_i(x) = 0$. It is easy to see that $p_i(x)$, $d_i(x,i,0)$ and $d_i(x,0,i)$ are feasible to (6). Moreover, since the revenue functions $r_{i0}(d)$ and $r_{0i}(d)$ are strictly concave by Assumption 2.1, due to Jensen's inequality, the objective value with $p_i(x)$ and controls $d_i(x,i,0)$ and $d_i(x,0,i)$ is strictly larger than the objective values with the probability distributions $p_i^a(x)$ and $p_i^b(x)$. This violates the optimality of $p_i^a(x)$ and $p_i^b(x)$ and thus a contradiction. □

Proposition B.9 shows that the stationary distribution for each spoke is (discrete) log-concave.

Definition B.1 (Discrete Log-concavity, c.f., Keilson and Gerber 1971, Keilson 1972)**.** A discrete probability distribution $\mathbf{p} = \{p_i\}_{i=0}^{\infty}$ with all its support on non-negative integers is *discrete log-concave* (or simply *log-concave*) if (i) its support $I_{\mathbf{p}} = \{i \geq 0 : p_i > 0\}$ is a sequence of consecutive integers, i.e., for all $0 \leq n_1 \leq n \leq n_2$, if $n_1, n_2 \in I_{\mathbf{p}}$, then $n \in I_{\mathbf{p}}$; and (ii) $p_i^2 \geq p_{i-1} \cdot p_{i+1}$ for all $i \geq 1$.

Proposition B.9. *For each spoke $i \in [n]$, the stationary distribution $p_i(x)$ solved from (6) is discrete log-concave.*

Proof. From Lemma B.1, the support of $p_i(x)$ is $I_i = [0 : H_i]$ that is a sequence of consecutive integers. Secondly, from the flow balance constraint in (6)

$$p_i(x) \cdot q_{0i} \cdot d_i(x,0,i) = p_i(x+1) \cdot q_{i0} \cdot d_i(x+1,i,0), \ \forall\, x \in [0 : m-1],$$

we have

$$\left(p_i(x)\right)^2 \cdot d_i(x,i,0) \cdot d_i(x,0,i) = p_i(x-1) \cdot p_i(x+1) \cdot d_i(x+1,i,0) \cdot d_i(x-1,0,i)$$

for all $x \in [1 : m-1]$. Since the demand level $d_i(x,0,i)$ is decreasing in x and $d_i(x,i,0)$ is increasing in x for $x \in I_i$ by Proposition 3.3 and the support of $p_i(x)$ is a sequence of consecutive integers, we have $\left(p_i(x)\right)^2 \geq p_i(x-1) \cdot p_i(x+1)$ for all $x \in [1 : m-1]$. □

Corollary B.10 shows that the Lagrangian policy in the Lagrangian relaxation has a unique stationary distribution, which factors across spokes.

Corollary B.10 (Lagrangian Policy in the Relaxation)**.** *The Lagrangian policy is optimal to the Lagrangian relaxation. Moreover, let the system state be the resource levels $\mathbf{x} \in \bar{\mathcal{X}}$. Using the Lagrangian policy, the following hold:*

1. *The set $\prod_{i=1}^{n} I_i$ is a positive recurrent class and is aperiodic, and all states outside the set $\prod_{i=1}^{n} I_i$ are transient; and*

2. $q(\mathbf{x}) = \prod_{i=1}^n p_i(x_i)$ *is the unique stationary distribution.*

Proof. Since the Lagrangian relaxation decomposes over spokes and the Lagrangian policy is optimal to each spoke problem, the Lagrangian policy is optimal to the Lagrangian relaxation as well. From Lemma B.1, it is easy to see that the set $\prod_{i=1}^n I_i$ is positive recurrent and all states outside $\prod_{i=1}^n I_i$ are transient. To show the aperiodicity, let $\mathbf{x} \in \bar{\mathcal{X}}$ be a boundary state of the set $\prod_{i=1}^n I_i$. Without loss of generality we assume $\mathbf{x} \in \prod_{i=1}^n I_i$ whereas $\mathbf{x} + \boldsymbol{e}_{i'} \notin \prod_{i=1}^n I_i$ for some spoke $i' \in [n]$. State $\mathbf{x}$ stays unchanged when a request $(0, i')$ arrives, which occurs with probability $q_{0i'} > 0$. Thus the chain is aperiodic.

Since the Markov chain has a single positive recurrent class $\prod_{i=1}^n I_i$, the stationary distribution is unique. $q(\mathbf{x})$ is the stationary distribution if and only if $\sum_{\mathbf{x}\in\bar{\mathcal{X}}} q(\mathbf{x}) = 1$ and for all $\mathbf{x} \in \bar{\mathcal{X}}$,

$$\begin{aligned} q(\mathbf{x}) \cdot \sum_{i\in[n]} \Big[q_{i0} \cdot d_i\big(x_i, i, 0\big) + q_{0i} \cdot d_i\big(x_i, 0, i\big) \Big] = & \sum_{i\in[n]} \Big\{ \mathbb{1}\left[x_i \geq 1\right] \cdot q(\mathbf{x} - \boldsymbol{e}_i) \cdot q_{0i} \cdot d_i(x_i - 1, 0, i) \Big\} \\ & + \sum_{i\in[n]} \Big\{ \mathbb{1}\left[x_i \leq m-1\right] \cdot q(\mathbf{x} + \boldsymbol{e}_i) \cdot q_{i0} \cdot d_i(x_i + 1, i, 0) \Big\}. \end{aligned} \tag{66}$$

It is easy to see that $q(\mathbf{x}) = \prod_{i=1}^n p_i(x_i)$ satisfies (66) because of the flow balance constraint in (6). Thus, $q(\mathbf{x}) = \prod_{i=1}^n p_i(x_i)$ is the unique stationary distribution in the Lagrangian relaxation. ☐

Corollary B.11 shows that using the Lagrangian policy in the original problem also leads to a unichain policy; the proof is analogous to proof of Corollary B.10.

Corollary B.11 (Lagrangian Policy in the Original Problem). *Let the system state be the resource levels $\mathbf{x} \in \mathcal{X}$. Using the Lagrangian policy, the set $\prod_{i=1}^n I_i \cap \mathcal{X}$ is the single positive recurrent class and is aperiodic, and all states outside the set $\prod_{i=1}^n I_i \cap \mathcal{X}$ are transient.*

Proposition B.12 formalizes the decomposition across spokes and hubs for general networks.

Proposition B.12. *The Lagrangian relaxation bound $\bar{V}^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}}$ described in Section 6.2 decomposes over spokes with*

$$\bar{V}^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}} = m\lambda + \sum_{i=1}^n h_i^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}} + \sum_{j,j'\in[J]} q_{jj'} \cdot g_{jj'}^{\boldsymbol{\mu}},$$

where $g_{jj'}^{\boldsymbol{\mu}} \triangleq \max_{d\in[0,1]} \left\{ r_{jj'}(d) + d \cdot \left(\mu_{j'} - \mu_j\right)\right\}$ denotes the average revenue earned from a hub-to-hub request (j, j') as in (14), and $h_i^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}}$ denotes the average revenue of an optimal policy to each spoke i problem, which is equal to the optimal value of (67)

$$\begin{aligned} \max_{\substack{d_i(x,i,j)\in[0,1],\\ d_i(x,j,i)\in[0,1],\\ d_i(x,i,i')\in[0,1],\\ d_i(x,i',i)\in[0,1],\\ p_i(x)\geq 0}} \quad & \sum_{x=0}^m p_i(x) \cdot \left\{ \sum_{j=1}^J \Big[q_{ij} \cdot r_{ij}\Big(d_i(x,i,j)\Big) + q_{ji} \cdot r_{ji}\Big(d_i(x,j,i)\Big) \Big] + \sum_{i'=1}^n q_{ii'} \cdot r_{ii'}\Big(d_i(x,i,i')\Big) \right\} \\ & + \sum_{i'=1}^n \sum_{x=0}^m p_i(x) \cdot \Big(\nu_{ii'} \cdot q_{ii'} \cdot d_i(x,i,i') - \nu_{i'i} \cdot q_{i'i} \cdot d_i(x,i',i) \Big) \\ & + \sum_{j=1}^J \mu_j \sum_{x=0}^m p_i(x) \cdot \Big(q_{ij} \cdot d_i(x,i,j) - q_{ji} \cdot d_i(x,j,i) \Big) - \lambda \cdot \sum_{x=0}^m x \cdot p_i(x) \\ s.t. \quad & \sum_{x=0}^m p_i(x) = 1, \end{aligned} \tag{67}$$

$$
\begin{aligned}
&p_i(x) \cdot \left(\sum_{j=1}^{J} q_{ji} \cdot d_i(x,j,i) + \sum_{i'=1}^{n} q_{i'i} \cdot d_i(x,i',i) \right) \\
&= p_i(x+1) \cdot \left(\sum_{j=1}^{J} q_{ij} \cdot d_i(x+1,i,j) + \sum_{i'=1}^{n} q_{ii'} \cdot d_i(x+1,i,i') \right), \ \forall\, x \in [0:m-1], \\
&d_i(0,i,j) = 0, d_i(m,j,i) = 0, \ \forall\, j \in [J], \\
&d_i(0,i,i') = 0, d_i(m,i',i) = 0, \ \forall\, i' \in [n].
\end{aligned}
$$

Proof. To simplify the notation, in the proof we suppress the superscript $\lambda, \boldsymbol{\mu}, \boldsymbol{\nu}$ throughout that specifies the specific dual variables to use in the Lagrangian relaxation.

Since we assume the network topology is strongly connected, the Lagrangian relaxation bound $\bar{V}$ does not depend on the initial state of the system by the same argument as in Proposition 2.2. Moreover, since we relax the capacity constraint of each hub, we can express $\bar{V}$ as

$$
\bar{V} = m\lambda + h + \sum_{j,j' \in [J]} q_{jj'} \cdot g_{jj'},
$$

where $g_{jj'} = \max_{d \in [0,1]} \left\{ r_{jj'}(d) + d \cdot (\mu_{j'} - \mu_j) \right\}$ denotes the average revenue from a hub-to-hub request (j,j'), and h denotes the average revenue of an optimal control to requests that involve any of the spokes. Furthermore, h and some differential value functions $v(\mathbf{x}_\mathrm{S}, i, j)$, $v(\mathbf{x}_\mathrm{S}, j, i)$, $v(\mathbf{x}_\mathrm{S}, i, i')$ and $v(\mathbf{x}_\mathrm{S}, j, j')$ satisfy the following Bellman equation

$$
\begin{aligned}
h + v(\mathbf{x}_\mathrm{S}, i, j) &= \max_{d \in [0, 1 \wedge x_i]} \left\{ r_{ij}(d) + d \cdot \Big(v(\mathbf{x}_\mathrm{S} - \mathbf{e}_i) - v(\mathbf{x}_\mathrm{S}) + \mu_j \Big) \right\} + v(\mathbf{x}_\mathrm{S}) - \lambda \cdot \sum_{i \in [n]} x_i, \\
&\qquad \forall\, i \in [n], j \in [J], \\
h + v(\mathbf{x}_\mathrm{S}, j, i) &= \max_{d \in [0, 1 \wedge (m - x_i)]} \left\{ r_{ji}(d) + d \cdot \Big(v(\mathbf{x}_\mathrm{S} + \mathbf{e}_i) - v(\mathbf{x}_\mathrm{S}) - \mu_j \Big) \right\} + v(\mathbf{x}_\mathrm{S}) - \lambda \cdot \sum_{i \in [n]} x_i, \\
&\qquad \forall\, i \in [n], j \in [J], \\
h + v(\mathbf{x}_\mathrm{S}, i, i') &= \max_{\substack{d \in [0, 1 \wedge x_i] \\ a_1, a_2 \in \{0, 1 \wedge (m - x_{i'})\}}} r_{ii'}(d) + d \cdot \Big(v(\mathbf{x}_\mathrm{S} - e_i + a_1 \cdot e_{i'}) + \nu_{ii'} - a_1 \cdot \nu_{ii'} \Big) \\
&\quad + (1-d) \cdot \Big(v(\mathbf{x}_\mathrm{S} + a_2 \cdot e_{i'}) - a_2 \cdot \nu_{ii'} \Big) - \lambda \cdot \sum_{i \in [n]} x_i, \ \forall\, i \in [n], i' \in [n] \setminus \{i\}, \\
h + v(\mathbf{x}_\mathrm{S}, j, j') &= v(\mathbf{x}_\mathrm{S}) - \lambda \cdot \sum_{i \in [n]} x_i, \ \forall\, j, j' \in [J],
\end{aligned} \tag{68}
$$

for all $\mathbf{x}_\mathrm{S} \in [0:m]^n$, where $v(\mathbf{x}_\mathrm{S}) = \mathbb{E}_s[v(\mathbf{x}_\mathrm{S}, s)]$ denotes the average differential value function over request types, and the binary variables a_1 and a_2 in the third equation denote the decision of adding one resource in the destination when a spoke-to-spoke request (i,i') arrives, and the request is fulfilled or not, respectively.

It is easy to verify that the average revenue h and the differential value functions decompose over spokes with

$$
\begin{aligned}
&h = \sum_{i \in [n]} h_i, \\
&v(\mathbf{x}_{\mathrm{S}}, i, j) = v_i(x_i, i, j) + \sum_{k \neq i} v_k(x_k, \varnothing), \\
&v(\mathbf{x}_{\mathrm{S}}, j, i) = v_i(x_i, j, i) + \sum_{k \neq i} v_k(x_k, \varnothing), \\
&v(\mathbf{x}_{\mathrm{S}}, i, i') = v_i(x_i, i, i') + v_{i'}(x_{i'}, i, i') + \sum_{k \neq i, i'} v_k(x_k, \varnothing), \\
&v(\mathbf{x}_{\mathrm{S}}, j, j') = \sum_{k \in [n]} v_k(x_k, \varnothing),
\end{aligned}
\tag{69}
$$

where h_i denotes the average revenue of an optimal policy to each spoke i problem, and the differential value functions $v_i(x,i,j)$, $v_i(x,j,i)$, $v_i(x,i,i')$, $v_i(x,i',i)$ and $v_i(x,\varnothing)$ correspond to the state with x resources in spoke i and the request type being (i,j), (j,i), (i,i'), (i',i), or one of any other types, respectively. Moreover, h_i and the differential value functions satisfy the following Bellman equation

$$
\begin{aligned}
&h_i + v_i(x,i,j) = \max_{d \in [0, 1 \wedge x]} \Big\{ r_{ij}(d) + d \cdot \Big(v_i(x-1) - v_i(x) + \mu_j \Big) \Big\} + v_i(x) - \lambda \cdot x, \ \forall\, j \in [J], \\
&h_i + v_i(x,j,i) = \max_{d \in [0, 1 \wedge (m-x)]} \Big\{ r_{ji}(d) + d \cdot \Big(v_i(x+1) - v_i(x) - \mu_j \Big) \Big\} + v_i(x) - \lambda \cdot x, \ \forall\, j \in [J], \\
&h_i + v_i(x,i,i') = \max_{d \in [0, 1 \wedge x]} \Big\{ r_{ii'}(d) + d \cdot \Big(v_i(x-1) - v_i(x) + \nu_{ii'} \Big) \Big\} + v_i(x) - \lambda \cdot x, \ \forall\, i' \in [n] \setminus \{i\}, \\
&h_i + v_i(x,i',i) = \max_{a \in \{0, 1 \wedge (m-x)\}} \Big\{ v_i(x+a) - a \cdot \nu_{i'i} \Big\} - \lambda \cdot x, \ \forall\, i' \in [n] \setminus \{i\}, \\
&h_i + v_i(x,\varnothing) = v_i(x) - \lambda \cdot x,
\end{aligned}
\tag{70}
$$

for all $x \in [0:m]$ and each spoke i, with $v_i(x)$ being the average differential value function over request types. Here we only verify the decomposition of $v(\mathbf{x}_{\mathrm{S}}, i, i')$ that involves a spoke-to-spoke request (i,i'). Suppose the average revenues and differential value functions decompose over spokes by (69) and (70) holds. The right-hand side of the third equation in (68) is equal to

$$
\begin{aligned}
&\max_{\substack{d\in[0,1\wedge x_i]\\ a_1,a_2\in\{0,1\wedge(m-x_{i'})\}}} \quad r_{ii'}(d) + d\cdot\Big(v_i(x_i-1) + v_{i'}(x_{i'}+a_1) + \nu_{ii'} - a_1\cdot\nu_{ii'}\Big)\\
&\qquad\qquad + (1-d)\cdot\Big(v_i(x_i) + v_{i'}(x_{i'}+a_2) - a_2\cdot\nu_{ii'}\Big) + \sum_{k\neq i,i'} v_k(x_k) - \lambda\cdot\sum_{i\in[n]} x_i\\
&= \max_{d\in[0,1\wedge x_i]} r_{ii'}(d) + d\cdot\Big(v_i(x_i-1) + \nu_{ii'} + \max_{a_1\in\{0,1\wedge(m-x_{i'})\}} v_{i'}(x_{i'}+a_1) - a_1\cdot\nu_{ii'}\Big)\\
&\qquad + (1-d)\cdot\Big(v_i(x_i) + \max_{a_2\in\{0,1\wedge(m-x_{i'})\}} v_{i'}(x_{i'}+a_2) - a_2\cdot\nu_{ii'}\Big) + \sum_{k\neq i,i'} v_k(x_k) - \lambda\cdot\sum_{i\in[n]} x_i\\
&\overset{\text{(i)}}{=} \max_{a\in\{0,1\wedge(m-x_{i'})\}} v_{i'}(x_{i'}+a_1) - a\cdot\nu_{ii'}\\
&\qquad + \max_{d\in[0,1\wedge x_i]} r_{ii'}(d) + d\cdot\Big(v_i(x_i-1) + \nu_{ii'} - v_i(x_i)\Big) + v_i(x_i) + \sum_{k\neq i,i'} v_k(x_k) - \lambda\cdot\sum_{i\in[n]} x_i\\
&= \sum_{i\in[n]} h_i + v_i(x_i,i,i') + v_{i'}(x_{i'},i,i') + \sum_{k\neq i,i'} v_k(x_k,\varnothing)\\
&= h + v(\mathbf{x}_{\mathrm{S}},i,i').
\end{aligned}
$$

Note that (i) implies that although the provider can make the decision of adding one resource in the destination after knowing the outcome of the fulfillment, it loses nothing if she instead makes the decision before the outcome, by comparing the cost $\nu_{ii'}$ and the marginal value of having one more resource in the destination. Thus we can make the decisions at the origin and destination independently.

Finally, following the same argument as in Proposition 3.2, h_i is equal to the optimal value of (67), which is the dual formulation of the spoke problem. □

For an arbitrary network, we can divide the locations into hubs and spokes and consider the Lagrangian relaxation bound and policy in Section 6.2. Proposition 6.2 shows that the Lagrangian relaxation provides tighter bounds than the fluid relaxation bound V^{F}. We provide the proof here.

Proof of Proposition 6.2. The optimality condition of $\min_{\boldsymbol{\mu},\boldsymbol{\nu}} \bar{V}^{\lambda=0,\boldsymbol{\mu},\boldsymbol{\nu}}$ implies that it is equivalent to (71), the problem of maximizing the average revenue subject to the constraints that the in-flow and out-flow of each hub j is balanced in expectation, and the out-flow of spoke i through requests (i,i') is equal to the in-flow of spoke i' through requests (i,i') for each spoke-to-spoke connection (i,i').

$$
\begin{aligned}
\max_{\substack{d_i(x,i,j)\in[0,1],\\ d_i(x,j,i)\in[0,1],\\ d_i(x,i,i')\in[0,1],\\ d_i(x,i',i)\in[0,1],\\ d_{jj'}\in[0,1],\\ p_i(x)\geq 0}} \quad & \sum_{i=1}^{n}\sum_{x=0}^{m} p_i(x)\cdot\sum_{j=1}^{J}\left\{q_{ij}\cdot r_{ij}\Big(d_i(x,i,j)\Big)+q_{ji}\cdot r_{ji}\Big(d_i(x,j,i)\Big)\right\} \\
& \quad +\sum_{i=1}^{n}\sum_{x=0}^{m} p_i(x)\cdot\sum_{i'=1}^{n} q_{ii'}\cdot r_{ii'}\Big(d_i(x,i,i')\Big)+\sum_{j=1}^{J}\sum_{j'=1}^{J} q_{jj'}\cdot r_{jj'}(d_{jj'}) \\
\text{s.t.} \quad & \sum_{x=0}^{m} p_i(x)=1,\ \forall\, i\in[n], \\
& \sum_{i=1}^{n} q_{ij}\sum_{x=0}^{m} p_i(x)\cdot d_i(x,i,j)+\sum_{j'=1}^{J} q_{j'j}\cdot d_{j'j} \\
& \quad =\sum_{i=1}^{n} q_{ji}\sum_{x=0}^{m} p_i(x)\cdot d_i(x,j,i)+\sum_{j'=1}^{J} q_{jj'}\cdot d_{jj'},\ \forall\, j\in[J], \\
& p_i(x)\cdot\left(\sum_{j=1}^{J} q_{ji}\cdot d_i(x,j,i)+\sum_{i'=1}^{n} q_{i'i}\cdot d_i(x,i',i)\right)= \\
& \quad p_i(x+1)\cdot\left(\sum_{j=1}^{J} q_{ij}\cdot d_i(x+1,i,j)+\sum_{i'=1}^{n} q_{ii'}\cdot d_i(x+1,i,i')\right), \\
& \quad \forall\, x\in[0:m-1], i\in[n], \\
& q_{ii'}\cdot\sum_{x=0}^{m} p_i(x)\cdot d_i(x,i,i')=q_{ii'}\cdot\sum_{x=0}^{m} p_{i'}(x)\cdot d_{i'}(x,i,i'),\ \forall\, i,i'\in[n], \\
& d_i(0,i,j), d_i(0,i,i'), d_i(m,j,i), d_i(m,i',i)=0,\ \forall\, i\in[n], j\in[J], i'\in[n].
\end{aligned} \tag{71}
$$

For any optimal solution to (71), let $d_{ij}=\sum_{x=0}^{m} p_i(x)\cdot d_i(x,i,j)$, $d_{ji}=\sum_{x=0}^{m} p_i(x)\cdot d_i(x,j,i)$ and $d_{ii'}=\sum_{x=0}^{m} p_i(x)\cdot d_i(x,i,i')$ denote the average demand values. We show these demand values plus $d_{jj'}$ are feasible to the fluid relaxation. First, by the second constraint in (71) we have

$$
\sum_{i=1}^{n} q_{ij}d_{ij}+\sum_{j'=1}^{J} q_{j'j}\cdot d_{j'j}=\sum_{i=1}^{n} q_{ji}d_{ji}+\sum_{j'=1}^{J} q_{jj'}\cdot d_{jj'},\ \forall\, j\in[J], \tag{72}
$$

which implies that the flow at each hub is balanced. Second, summing both sides of the third constraint over $x\in[0:m-1]$ plus the fourth and last constraints gives

$$
\sum_{j=1}^{J} q_{ji}\cdot d_{ji}+\sum_{i'=1}^{n} q_{i'i}\cdot d_{i'i}=\sum_{j=1}^{J} q_{ij}d_{ij}+\sum_{i'=1}^{n} q_{ii'}\cdot d_{ii'},\ \forall\, i\in[n], \tag{73}
$$

which implies that the flow at each spoke is balanced as well. Thus from (72) and (73), the demand values d_{ij}, d_{ji}, $d_{ii'}$ and $d_{jj'}$ are feasible to the fluid relaxation. Finally, by Jensen's inequality

$$\begin{aligned} V^{\mathrm{F}} &\geq \sum_{i=1}^{n}\sum_{j=1}^{J}\Big(q_{ij}\cdot r_{ij}(d_{ij}) + q_{ji}\cdot r_{ji}(d_{ji})\Big) + \sum_{i=1}^{n}\sum_{i'=1}^{n} q_{ii'}\cdot r_{ii'}(d_{ii'}) + \sum_{j=1}^{J}\sum_{j'=1}^{J} q_{jj'}\cdot r_{jj'}(d_{jj'}) \\ &\geq \min_{\boldsymbol{\mu},\boldsymbol{\nu}} \bar{V}^{\lambda=0,\boldsymbol{\mu},\boldsymbol{\nu}}. \end{aligned}$$

□

In the case with general relocation times, Proposition B.13 further relax the spoke problem to provide a tractable upper bound.

Proposition B.13. *With general relocation times, the Lagrangian relaxation bound $\bar{V}^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}}$ decomposes over spokes as*

$$\bar{V}^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}} = m\lambda + \sum_{i=1}^{n} h_i^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}} + \sum_{j,j'\in[J]} q_{jj'}\cdot g_{jj'}^{\boldsymbol{\mu}},$$

where $g_{jj'}^{\boldsymbol{\mu}} \triangleq \max_{d\in[0,1]}\left\{r_{jj'}(d) + d\cdot\big(\mu_{j'} - \mu_j - \lambda\cdot\Lambda\tau_{jj'}\big)\right\}$ denotes the average revenue earned from a hub-to-hub request (j,j') and $h_i^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}}$ denotes the average revenue of an optimal policy to each spoke i problem. Moreover, $h_i^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}}$ is no larger than $\hat{h}_i$ which is the optimal value of (74).

$$\begin{aligned}
\max_{\substack{d_i(x,i,j)\in[0,1],\\ d_i(x,j,i)\in[0,1],\\ d_i(x,i,i')\in[0,1],\\ d_i(x,i',i)\in[0,1],\\ p_i(x)\geq 0}} \quad & \sum_{x=0}^{m}\sum_{j=1}^{J} p_i(x)\cdot q_{ij}\cdot\Big[r_{ij}\Big(d_i(x,i,j)\Big) + \Big(\mu_j - \lambda\Lambda\tau_{ij}\Big)\cdot d_i(x,i,j)\Big] \\
& + \sum_{x=0}^{m}\sum_{j=1}^{J} p_i(x)\cdot q_{ji}\cdot\Big[r_{ji}\Big(d_i(x,j,i)\Big) - \mu_j\cdot d_i(x,j,i)\Big] \\
& + \sum_{x=0}^{m}\sum_{i'\neq i} p_i(x)\cdot q_{ii'}\cdot\Big[r_{ii'}\Big(d_i(x,i,i')\Big) + \nu_{ii'}\cdot d_i(x,i,i')\Big] \\
& - \sum_{x=0}^{m}\sum_{i'\neq i} p_i(x)\cdot q_{i'i}\cdot\nu_{i'i}\cdot d_i(x,i',i) + \Big(\sum_{x=1}^{m} p_i(x)\Big)\cdot q_{ii}\cdot r_{ii}^{*} - \lambda\cdot\sum_{x=0}^{m} x\cdot p_i(x) \\
\text{s.t.} \quad & \sum_{x=0}^{m} p_i(x) = 1, \qquad (74) \\
& p_i(x)\cdot\left\{\sum_{j=1}^{J} q_{ji}\cdot d_i(x,j,i) + \sum_{i'\neq i} q_{i'i}\cdot d_i(x,i',i)\right\} \\
& = p_i(x+1)\cdot\left\{\sum_{j=1}^{J} q_{ij}\cdot d_i(x+1,i,j) + \sum_{i'\neq i} q_{ii'}\cdot d_i(x+1,i,i')\right\},\ \forall\, x\in[0:m-1], \\
& d_i(0,i,j) = 0, d_i(m,j,i) = 0,\ \forall\, j\in[J], \\
& d_i(0,i,i') = 0, d_i(m,i',i) = 0,\ \forall\, i'\in[n]\setminus\{i\}.
\end{aligned}$$

Proof. The decomposition is analogous to Proposition B.12. Note that compared to Proposition B.12, the extra term $\lambda\cdot\Lambda\tau_{jj'}$ in $g_{jj'}^{\boldsymbol{\mu}}$ comes from the fact that the relocation (j,j') takes $\Lambda\cdot\tau_{jj'}$ periods on average (this is because requests follow a Possion process of rate Λ that is independent of the relocation times) and each period incurs a penalty λ.

We now show the average revenue $h_i^{\lambda,\boldsymbol{\mu},\boldsymbol{\nu}}$ of the spoke problem is no larger than $\hat{h}_i$. To see this, we first set the binding condition of the spoke problem to be the sum of resources in the spoke and transiting to it is no larger than m; then the resources that are moving out of the spoke are irrelevant. We then allow that the resources that are moving to the spoke can be instantaneously available at the spoke. Because a resource incurs a penalty λ per period no matter it is in the spoke or moving to it, it is always better to keep the resources at the spoke as this increases the opportunity to serve the requests; this yields (74). We conjecture that the relaxation works well when incoming relocation times are not long.

The Lagrangian relaxation with the optimal dual variable corresponds to maximizing the average revenue subject to the sum of resources that are in the spokes and transiting to the hubs no larger than m in expectation. □

C More Discussions on the Lagrangian Dual Problem

Recall that the Lagrangian dual problem (9) is

$$V^{\mathrm{R}} = \min_{\lambda \geq 0} \bar{V}^{\lambda},$$

which is a convex optimization problem. According to (4) and (20), V^{R} is equal to the optimal value of (75).

$$\begin{aligned}
\min_{\substack{\lambda\geq 0, h_i^{\lambda},\\ v_i^{\lambda}(x,i,0),\\ v_i^{\lambda}(x,0,i),\\ v_i^{\lambda}(x,\varnothing)}} \quad & m\lambda + \sum_{i=1}^{n} h_i^{\lambda} \\
\text{s.t.} \quad & h_i^{\lambda} + v_i^{\lambda}(x,i,0) \geq \max_{d\in[0,1\wedge x]} \left\{ r_{i0}(d) + d\cdot\Big(v_i^{\lambda}(x-1) - v_i^{\lambda}(x)\Big)\right\} + v_i^{\lambda}(x) - \lambda\cdot x, \\
& \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall\, x\leq m, i\in[n], \\
& h_i^{\lambda} + v_i^{\lambda}(x,0,i) \geq \max_{d\in[0,1\wedge (m-x)]} \left\{ r_{0i}(d) + d\cdot\Big(v_i^{\lambda}(x+1) - v_i^{\lambda}(x)\Big)\right\} + v_i^{\lambda}(x) - \lambda\cdot x, \\
& \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \forall\, x\leq m, i\in[n], \\
& h_i^{\lambda} + v_i^{\lambda}(x,\varnothing) \geq v_i^{\lambda}(x) - \lambda\cdot x,\ \forall\, x\leq m, i\in[n].
\end{aligned} \tag{75}$$

Analogous to Proposition 3.2, V^{R} is the optimal value of (76) as well, which is maximizing the average revenue subject to the constraint that the expected number of resources in the hub is non-negative.

$$\begin{aligned}
\max_{\substack{d_i(x,i,0)\in[0,1],\\ d_i(x,0,i)\in[0,1],\\ p_i(x)\geq 0}} \quad & \sum_{i=0}^{n}\sum_{x=0}^{m} p_i(x)\Big[q_{i0}\cdot r_{i0}\Big(d_i(x,i,0)\Big) + q_{0i}\cdot r_{0i}\Big(d_i(x,0,i)\Big)\Big] \\
\text{s.t.} \quad & \sum_{i=1}^{n}\sum_{x=0}^{m} x\cdot p_i(x) \leq m, \\
& \sum_{x=0}^{m} p_i(x) = 1,\ \forall\, i\in[n],
\end{aligned} \tag{76}$$

$$
\begin{aligned}
&p_i(x) \cdot q_{0i} \cdot d_i(x, 0, i) = p_i(x+1) \cdot q_{i0} \cdot d_i(x+1, i, 0),\ \forall\, x \in [0 : m-1],\ \forall\, i \in [n],\\
&d_i(0, i, 0) = 0,\ \forall\, i \in [n],\\
&d_i(m, 0, i) = 0,\ \forall\, i \in [n].
\end{aligned}
$$

Finally, we can solve (9) efficiently using a cutting plane method (Section 8.3 of Bertsekas et al. 2003) with sub-gradients given in (34).

D Optimal Static Pricing in the Large Network Regime

In this section, we use the same Lagrangian method to derive a performance bound for any static policy and characterize the optimal static policy in the large network limit. When we focus on static pricing policies and relax the capacity constraint $\sum_{i=1}^n x_i \le m$ with a dual variable $\lambda \ge 0$, the problem again decomposes over spokes and for any $\lambda \ge 0$, we have an upper bound $\bar{V}^{S,\lambda}$ on the performance of any static policy. Let $h_i^{\mathrm{S}}(\lambda)$ denote the performance of the spoke problem. We have $\bar{V}^{S,\lambda} = m\lambda + \sum_{i=1}^n h_i^{\mathrm{S}}(\lambda)$, and

$$
\begin{aligned}
h_i^{\mathrm{S}}(\lambda) = \max_{\substack{d_{0i}\in[0,1],\\ d_{i0}\in[0,1],\\ p_i(x)\ge 0}} \quad & \sum_{x=0}^{m} p_i(x) \cdot \Big\{ q_{i0} \cdot r_{i0}\big(d_{i0}\big) \cdot \mathbb{1}[x>0] + q_{0i} \cdot r_{0i}\big(d_{0i}\big) \cdot \mathbb{1}[x<m] \Big\} - \lambda \cdot \sum_{x=0}^{m} x \cdot p_i(x)\\
\text{s.t.} \quad & \sum_{x=0}^{m} p_i(x) = 1,\\
& p_i(x) \cdot q_{0i} \cdot d_{0i} = p_i(x+1) \cdot q_{i0} \cdot d_{i0},\ \forall\, x \in [0 : m-1].
\end{aligned} \tag{77}
$$

Note that we restrict to static pricing in the spoke problem (77).

We can solve the best possible bound $V^{\mathrm{S,R}} = \min_{\lambda\ge 0} \bar{V}^{S,\lambda}$ and compute a static policy $\pi^{\mathrm{S}}(\delta)$ from a perturbed problem $V^{\mathrm{S,R}}(\delta) = \min_{\lambda \ge 0} \bar{V}^{S,\lambda} - \delta\lambda$. Let $V^{\mathrm{S}}(\delta)$ denote the performance of policy $\pi^{\mathrm{S}}(\delta)$ in the original problem. Analogous to Section 4, Theorem D.1 shows that $\pi^{\mathrm{S}}(\delta)$ converges to the optimal static policy in the large network regime with a proper choice of δ.

Theorem D.1. *The average revenue $V^{\mathrm{S}}(\delta)$ of the static pricing policy $\pi^{\mathrm{S}}(\delta)$ satisfies*

$$
0 \le V^{\mathrm{S,R}} - V^{\mathrm{S}}(\delta) \le (\bar{r} + \bar{\omega}) \cdot \mathbb{P}\big[X_0(\delta) = 0\big] + \bar{r} \cdot \frac{\delta}{m-\delta},
$$

where $\mathbb{P}\big[X_0(\delta)=0\big]$ is the stationary probability that the hub runs out of resources in the original problem under the policy $\pi^{\mathrm{S}}(\delta)$. Moreover, if there exist some constants $\bar{q} > 0$, $\underline{\beta} \in (\frac{m}{m+n}, 1)$ and $\varepsilon > 0$ such that $q_{i0}, q_{0i} \le \frac{\bar{q}}{n}$ and $\big(1-\underline{\beta}\big)^2 \cdot \min_{i\le n} \gamma_i'\big(\underline{\beta}\big) \ge \frac{\varepsilon}{n}$[11], then

$$
\mathbb{P}\Big[X_0(\delta) \le 0\Big] \le \exp\left(-\frac{b}{2} \cdot \frac{\delta^2}{m+n} \right)
$$

[11] Since the functions $\gamma_i(\beta)$ are concave, if $\gamma_i(\beta)$ are continuously differentiable, this is equivalent to requiring $\min_{i\le n} \gamma_i'\big(\frac{m}{m+n}\big) \ge \frac{\varepsilon}{n} \cdot \big(\frac{m+n}{n}\big)^2$ for some $\varepsilon > 0$.

for some constant $b > 0$ when m and n grow at the same rate. Thus,

$$V^{\mathrm{S,R}} - V^{\mathrm{S}}(\delta) \leq O\left(\sqrt{\frac{\ln n}{n}}\right)$$

if we set $\delta = \sqrt{\frac{1}{b} \cdot (m+n) \cdot \ln n}$.

D.1 Proof of Theorem D.1

We first rewrite the spoke problem (77). An optimal solution to (77) satisfies $d_{i0}, d_{0i} > 0$, and we let $\beta = \frac{q_{0i} \cdot d_{0i}}{q_{i0} \cdot d_{i0}}$. By the flow balance constraint in (77), we have $p_i(x+1) = \beta \cdot p_i(x)$, and hence

$$p_i(x) = \beta^x \cdot p_i(0),\ \forall\, x \in [m]. \tag{78}$$

Since these probabilities sum up to one, we have

$$p_i(0) = \left(1 + \sum_{x=1}^{m} \beta^x\right)^{-1} = \begin{cases} \frac{1}{m+1} & \text{if } \beta = 1, \\ \frac{1-\beta}{1-\beta^{m+1}} & \text{otherwise.} \end{cases} \tag{79}$$

The first part of the objective of $h_i^{\mathrm{S}}(\lambda)$ can be written as

$$\sum_{x=0}^{m-1} p_i(x) \cdot \Big[q_{0i} \cdot r_{0i}\big(d_{0i}\big) + \beta \cdot q_{i0} \cdot r_{i0}\big(d_{i0}\big)\Big] = \Big(1 - p_i(m)\Big) \cdot \Big[q_{0i} \cdot r_{0i}\big(d_{0i}\big) + \beta \cdot q_{i0} \cdot r_{i0}\big(d_{i0}\big)\Big].$$

As a result, we can rewrite $h_i^{\mathrm{S}}(\lambda)$ as

$$h_i^{\mathrm{S}}(\lambda) = \max_{\beta \geq 0} A^m(\beta) \cdot \gamma_i(\beta) - \lambda \cdot B^m(\beta), \tag{80}$$

where

$$A^m(\beta) = 1 - p_i(m) = 1 - \beta^m \cdot p_i(0) = \begin{cases} \frac{m}{m+1} & \text{if } \beta = 1, \\ \frac{1-\beta^m}{1-\beta^{m+1}} & \text{otherwise,} \end{cases}$$

and

$$B^m(\beta) = \sum_{x=0}^{m} x \cdot p_i(x) = p_i(0) \cdot \sum_{x=1}^{m} x\beta^x = \begin{cases} \frac{m}{2} & \text{if } \beta = 1, \\ \frac{\beta^{m+2}m - (1+m)\beta^{m+1} + \beta}{(1-\beta)(1-\beta^{m+1})} & \text{otherwise.} \end{cases}$$

We first provide some useful properties for $A^m(\beta)$ and $B^m(\beta)$.

Lemma D.2 (Monotonicity)**.** *$B^m(\beta)$ is strictly increasing in $\beta \geq 0$.*

Proof. For any $\beta_1 < \beta_2$ and $i \in \{1, 2\}$, let Z_i be a discrete random variable with support $[0\!:\!m]$ and density function $g_i(\cdot)$ specified by (78) and (79) using a parameter β_i. Since for any $0 \leq x < y \leq m$, $\frac{g_1(y)}{g_1(x)} = \beta_1^{y-x} < \beta_2^{y-x} = \frac{g_2(y)}{g_2(x)}$, Z_2 dominates Z_1 in the monotone likelihood ratio order (see Section 1.C of Shaked and Shanthikumar 2007). Hence, Z_2 first-order stochastically dominates Z_1 and as a result, $B^m(\beta_2) = \mathbb{E}[Z_2] > \mathbb{E}[Z_1] = B^m(\beta_1)$; the strict inequality is because Z_1 and Z_2 have distinct density functions (see Theorem 1.A.8 of Shaked and Shanthikumar 2007). $\square$

Lemma D.3 (Uniform convergence of $A^m(\beta)$ and $B^m(\beta)$)**.** *Let $A^\infty(\beta) = \begin{cases} 1 & \text{if } 0 \leq \beta \leq 1 \\ \frac{1}{\beta} & \text{if } \beta > 1 \end{cases}$ and*

$B^{\infty}(\beta) = \begin{cases} \frac{\beta}{1-\beta} & \text{if } 0 \leq \beta < 1 \\ \infty & \text{if } \beta \geq 1 \end{cases}$. *Then,*

1. $0 \leq A^{\infty}(\beta) - A^{m}(\beta) \leq \frac{1}{m+1}$ *for all* $\beta \geq 0$*;*
2. $0 \leq B^{\infty}(\beta) - B^{m}(\beta) \leq \frac{(m+1)\beta^{m+1}}{1-\beta}$ *for all* $\beta \in [0,1)$*, and* $\lim_{m\to\infty} B^{m}(\beta) = \infty$ *for all* $\beta \geq 1$*.*

Proof. To see the first part, note that if $\beta < 1$, we have

$$0 \leq A^{\infty}(\beta) - A^{m}(\beta) = \frac{\beta^m - \beta^{m+1}}{1-\beta^{m+1}} = \frac{\beta^m}{\sum_{i=0}^{m}\beta^i} \leq \frac{1}{m+1}.$$

If $\beta = 1$, $A^{\infty}(\beta) - A^{m}(\beta) = 1 - \frac{m}{m+1} = \frac{1}{m+1}$. If $\beta > 1$, we have

$$0 \leq A^{\infty}(\beta) - A^{m}(\beta) = \frac{\beta - 1}{\beta \cdot (\beta^{m+1} - 1)} = \frac{1}{\sum_{i=0}^{m}\beta^{i+1}} \leq \frac{1}{m+1}.$$

For the second part, if $\beta < 1$, we have

$$0 \leq B^{\infty}(\beta) - B^{m}(\beta) = \frac{(m+1)\cdot\beta^{m+1}}{1-\beta^{m+1}} \leq \frac{(m+1)\cdot\beta^{m+1}}{1-\beta}.$$

If $\beta \geq 1$, it is easy to see that $\lim_{m\to\infty} B^{m}(\beta) = \infty$. □

Lemma D.4 (Derivatives of $A^m(\beta)$ and $B^m(\beta)$)**.** *The derivatives of* $A^m(\beta)$ *and* $B^m(\beta)$ *satisfy*

1. $0 \geq \frac{d}{d\beta}A^m(\beta) \geq -m\beta^{m-1}$ *for all* $\beta \in [0,1)$*; and*
2. $0 \leq \frac{d}{d\beta}B^m(\beta) \leq \frac{d}{d\beta}B^{\infty}(\beta) = \frac{1}{(1-\beta)^2}$ *for all* $\beta \in [0,1)$*.*

Proof. For the first part, note that for any $\beta \in [0,1)$ we have

$$\frac{d}{d\beta}A^m(\beta) = \frac{\beta^{m-1}}{(1-\beta^{m+1})^2} \cdot \Big((m+1)\beta - \beta^{m+1} - m\Big).$$

Since $1 - \beta^{m+1} = (1-\beta)\cdot(1+\beta+\cdots+\beta^m)$ and

$$\begin{aligned} (m+1)\beta - \beta^{m+1} - m &= -\Big[(1-\beta)\cdot m + \beta^{m+1} - \beta\Big] = -(1-\beta)\cdot(m - \beta - \beta^2 - \cdots - \beta^m) \\ &= -(1-\beta)^2 \cdot \Big(m + (m-1)\beta + (m-2)\beta^2 + \cdots + \beta^{m-1}\Big), \end{aligned}$$

we have

$$0 \geq \frac{d}{d\beta}A^m(\beta) = -\beta^{m-1} \cdot \frac{\Big(m + (m-1)\beta + (m-2)\beta^2 + \cdots + \beta^{m-1}\Big)}{(1+\beta+\cdots+\beta^m)^2} \geq -m\beta^{m-1},$$

where the last inequality is because

$$m + (m-1)\beta + (m-2)\beta^2 + \cdots + \beta^{m-1} \leq m\cdot(1+\beta+\cdots+\beta^m) \leq m\cdot(1+\beta+\cdots+\beta^m)^2.$$

For the second part, since $B^\infty(\beta) - B^m(\beta) = \frac{(m+1)\cdot\beta^{m+1}}{1-\beta^{m+1}}$, we have

$$\frac{d}{d\beta}B^\infty(\beta) - \frac{d}{d\beta}B^m(\beta) = \frac{(m+1)^2\beta^m}{(1-\beta^{m+1})^2} \geq 0.$$

Thus,

$$0 \leq \frac{d}{d\beta}B^m(\beta) \leq \frac{d}{d\beta}B^\infty(\beta) = \frac{1}{(1-\beta)^2},$$

where the first inequality is because $B^m(\beta)$ is increasing in β by Lemma D.3. $\square$

D.1.1 Proof of Part One

Let $\lambda(\delta)$ denote an optimal solution to the perturbed problem $V^{\mathrm{S,R}}(\delta)$, d_{i0} and d_{0i} an optimal solution to the spoke problem $h_i^{\mathrm{S}}\big(\lambda(\delta)\big)$, and $\beta_i = \frac{q_{0i}d_{0i}}{q_{i0}d_{i0}}$. Let random variables $X_i(\delta)$ and $\tilde{X}_i(\delta)$ denote the number of resources in location $i \in [0:n]$ under the stationary distributions of policy $\pi^{\mathrm{S}}(\delta)$ in the original and the relaxed systems[12], respectively. Then,

$$\begin{aligned}
V^{\mathrm{S}}(\delta) &= \sum_{i=1}^{n}\Big\{q_{i0}\cdot\mathbb{P}\big[X_i(\delta)\geq 1\big]\cdot r_{i0}(d_{i0}) + q_{0i}\cdot\mathbb{P}\big[X_0(\delta)\geq 1\big]\cdot r_{0i}(d_{0i})\Big\} \\
&= \sum_{i=1}^{n}\Big\{q_{0i}\cdot\mathbb{P}\big[X_0(\delta)\geq 1\big]\cdot\frac{d_{0i}}{d_{i0}}\cdot r_{i0}(d_{i0}) + q_{0i}\cdot\mathbb{P}\big[X_0(\delta)\geq 1\big]\cdot r_{0i}(d_{0i})\Big\} \\
&= \sum_{i=1}^{n}\mathbb{P}\big[X_0(\delta)\geq 1\big]\cdot\Big\{q_{0i}\cdot r_{0i}(d_{0i}) + \beta_i\cdot q_{i0}\cdot r_{i0}(d_{i0})\Big\} \\
&= \mathbb{P}\big[X_0(\delta)\geq 1\big]\cdot\sum_{i=1}^{n}\gamma_i(\beta_i),
\end{aligned} \tag{81}$$

where the second equality is from the flow balance equation $q_{0i}\cdot\mathbb{P}[X_0(\delta)\geq 1]\cdot d_{0i} = q_{i0}\cdot\mathbb{P}[X_i(\delta)\geq 1]\cdot d_{i0}$, and the last equality is because d_{i0} and d_{0i} are optimal to $\gamma_i(\beta_i)$.

On the other hand, analogous to Section 4.2, $V^{\mathrm{S,R}}(\delta)$ is equal to the performance of $\pi^{\mathrm{S}}(\delta)$ in the relaxed system, thus

$$V^{\mathrm{S,R}}(\delta) = \sum_{i=1}^{n}A^m(\beta_i)\gamma_i(\beta_i) \leq \sum_{i=1}^{n}\gamma_i(\beta_i). \tag{82}$$

Combining (81) and (82) implies

$$V^{\mathrm{S,R}}(\delta) - V^{\mathrm{S}}(\delta) \leq \Big(1-\mathbb{P}\big[X_0(\delta)\geq 1\big]\Big)\cdot\sum_{i=1}^{n}\gamma_i(\beta_i) \leq \mathbb{P}\big[X_0(\delta)=0\big]\cdot(\bar{r}+\bar{\omega})\cdot\sum_{i=1}^{n}q_{0i}, \tag{83}$$

where the last inequality is due to Lemma A.4. Finally, analogous to Lemma 4.2, we have

$$V^{\mathrm{S,R}}(\delta) \leq V^{\mathrm{S,R}} \leq V^{\mathrm{S,R}}(\delta) + \bar{r}\cdot\frac{\delta}{m-\delta}. \tag{84}$$

[12] We consider the same relaxed system as in Section 4.2.

From (83) and (84) we have

$$V^{\mathrm{S,R}} - V^{\mathrm{S}}(\delta) = \Big(V^{\mathrm{S,R}} - V^{\mathrm{S,R}}(\delta)\Big) + \Big(V^{\mathrm{S,R}}(\delta) - V^{\mathrm{S}}(\delta)\Big) \le (\bar{r}+\bar{\omega})\cdot \mathbb{P}\big[X_0(\delta)=0\big] + \bar{r}\cdot \frac{\delta}{m-\delta},$$

which is analogous to Theorem 4.1.

D.1.2 Proof of Part Two

From Corollary 4.5 we have

$$\mathbb{P}\big[X_0(\delta)=0\big] \le \mathbb{P}\big[\tilde{X}_0(\delta)\le 0\big] = \mathbb{P}\Big[\sum_{i=1}^{n}\tilde{X}_i(\delta)\ge m\Big].$$

Since $\tilde{X}_i(\delta)$ are truncated geometric random variables with success probability $1-\beta_i$ and end point m, they are log-concave. If $\mathbb{E}\big[\tilde{X}_i(\delta)\big]\le c$ for some constant $c>0$ and all $i\in[n]$, from Proposition 4.6 and Corollary 4.7, we have

$$\mathbb{P}\Big[\tilde{X}_0(\delta)\le 0\Big] \le \exp\bigg(-\frac{b}{2}\cdot\frac{\delta^2}{m+n}\bigg)$$

with $b=\frac{1}{1+c}$. Moreover, if we choose $\delta=\sqrt{\frac{m+n}{b}\cdot\ln n}$, then $V^{\mathrm{S}}-V^{\mathrm{S}}(\delta)\le O\Big(\sqrt{\frac{\ln n}{n}}\Big)$ when m and n grow at the same rate. We now show that under the additional assumptions regarding q_{ij} and $\gamma_i(\beta)$, we can find such a constant c.

Let $\beta_i^m(\lambda)$ denote an optimal solution to the spoke problem $h_i^{\mathrm{S}}(\lambda)$ with some $\lambda\ge 0$. Lemma D.5 shows that $\beta_i^m(\lambda)$ shrinks towards zero when λ becomes large.

Lemma D.5. *For any $\lambda>0$ and $i\in[n]$, $\beta_i^m\big(\lambda\big)\le\frac{2q_{0i}\cdot(\bar{r}+\bar{\omega})}{\lambda+2q_{0i}\cdot(\bar{r}+\bar{\omega})}<1$ when m is large enough.*

Proof. Let $h_i^{\mathrm{S}}\big(\beta,\lambda\big)=A^m(\beta)\gamma_i(\beta)-\lambda B^m(\beta)$. It suffices to show that $h_i^{\mathrm{S}}\big(\beta,\lambda\big)\le h_i^{\mathrm{S}}\big(0,\lambda\big)=0$ for all $\beta\ge\bar{\beta}\triangleq\frac{2q_{0i}\cdot(\bar{r}+\bar{\omega})}{\lambda+2q_{0i}\cdot(\bar{r}+\bar{\omega})}$ and large enough m. Since $\lim_{m\to\infty}B^m\big(\bar{\beta}\big)=\frac{\bar{\beta}}{1-\bar{\beta}}$ by Lemma D.3, there exists some $\bar{m}\in\mathbb{N}_+$ such that for all $m\ge\bar{m}$, $B^m\big(\bar{\beta}\big)\ge\frac{1}{2}\cdot\frac{\bar{\beta}}{1-\bar{\beta}}$. Thus, for any $\beta\ge\bar{\beta}$ and $m\ge\bar{m}$,

$$h_i^{\mathrm{S}}\big(\beta,\lambda\big)=A^m(\beta)\gamma_i(\beta)-\lambda B^m(\beta)\le q_{0i}\cdot(\bar{r}+\bar{\omega})-\lambda B^m\big(\bar{\beta}\big)\le q_{0i}\cdot(\bar{r}+\bar{\omega})-\frac{\lambda}{2}\cdot\frac{\bar{\beta}}{1-\bar{\beta}}=0,$$

where the first inequality is due to the facts that $A^m(\beta)\le 1$, $\gamma_i(\beta)\le q_{0i}\cdot(\bar{r}+\bar{\omega})$ from Lemma A.4, and $B^m\big(\bar{\beta}\big)$ is increasing in β by Lemma D.2. □

Lemma D.6 shows that $\beta_i^m(\lambda)$ is large when λ is small.

Lemma D.6. *For any $\tilde{\beta}<1$, $\beta_i^m(\lambda)>\tilde{\beta}$ if $\lambda<(1-\tilde{\beta})^2\gamma_i'(\tilde{\beta})$ and m is large enough.*

Proof. Let $\varepsilon=(1-\tilde{\beta})^2\gamma_i'(\tilde{\beta})-\lambda>0$ and $\delta=\frac{1}{2}\cdot\frac{\varepsilon}{(\bar{r}+\bar{\omega})\cdot(q_{i0}+q_{0i})}$. Since by Lemmas D.3 and D.4, $A^m(\beta)$ and $\frac{d}{d\beta}A^m(\beta)$ converge to 1 and 0 uniformly on $[0,\tilde{\beta}]$ when m grows to infinity, there exists a constant $\underline{m}$ independent of spoke i such that for all $m\ge\underline{m}$ and $\beta\le\tilde{\beta}$, $A^m(\beta)\ge 1-\delta$ and

$\frac{d}{d\beta}A^m(\beta) \geq -\delta$. Thus, the derivative of the objective of the spoke problem satisfies

$$\begin{aligned}
&\frac{d}{d\beta}\Big(A^m(\beta)\gamma_i(\beta) - \lambda B^m(\beta)\Big) \\
=&\Big(\frac{d}{d\beta}A^m(\beta)\Big)\gamma_i(\beta) + A^m(\beta)\gamma_i'(\beta) - \lambda\frac{d}{d\beta}B^m(\beta) \\
\geq& - \delta \cdot \gamma_i(\beta) + (1-\delta)\cdot\gamma_i'(\beta) - \lambda\frac{d}{d\beta}B^m(\beta) \\
\overset{\text{(i)}}{\geq}& \gamma_i'(\beta) - \frac{\lambda}{(1-\beta)^2} - \delta\cdot(\bar{r}+\bar{\omega})\cdot(q_{i0}+q_{0i}) \\
\overset{\text{(ii)}}{\geq}& \frac{\varepsilon}{(1-\beta)^2} - \frac{\varepsilon}{2} \\
>& 0,
\end{aligned}$$

where (i) is because $\frac{d}{d\beta}B^m(\beta) \leq \frac{1}{(1-\beta)^2}$ by Lemma D.4 and $\gamma_i(\beta) \leq q_{0i}(\bar{r}+\bar{\omega})$ and $\gamma_i'(\beta) \leq q_{i0}(\bar{r}+\bar{\omega})$ by Lemma A.4, and (ii) is from the definitions of ε and δ. Therefore, the objective is strictly increasing in $[0, \tilde{\beta}]$ and as a result, $\beta_i^m(\lambda) > \tilde{\beta}$ when $m \geq \underline{m}$. $\square$

Let $\lambda(\delta)$ denote an optimal solution to the perturbed problem $V^{\mathrm{S,R}}(\delta)$. First, we have $\lambda(\delta) \geq \frac{\varepsilon}{n}$. If not, then from Lemma D.6, $\beta_i^m(\lambda(\delta)) > \underline{\beta}$ for all $i \in [n]$ when m is large enough. Since $B^m(\beta)$ is increasing in β by Lemma D.2, $\sum_{i=1}^n B^m\big(\beta_i^m(\lambda(\delta))\big) \geq nB^m\big(\underline{\beta}\big)$. Moreover, since $\lim_{m\to\infty} B^m(\beta) = \frac{\beta}{1-\beta}$ for all $\beta < 1$ and $\frac{\underline{\beta}}{1-\underline{\beta}} > \frac{\frac{m}{m+n}}{1-\frac{m}{m+n}} = \frac{m}{n}$, $\sum_{i=1}^n B^m\big(\beta_i^m(\lambda(\delta))\big) > m$ when m is large enough. But this contradicts with the fact that $\sum_{i=1}^n B^m\big(\beta_i^m(\lambda(\delta))\big) \leq m - \delta$ by the optimality condition of $\lambda(\delta)$. Hence, $\lambda(\delta) \geq \frac{\varepsilon}{n}$, and from Lemma D.5, we have

$$\beta_i^m\big(\lambda(\delta)\big) \leq \frac{2q_{0i}\cdot(\bar{r}+\bar{\omega})}{\lambda(\delta) + 2q_{0i}\cdot(\bar{r}+\bar{\omega})} \leq \bar{\beta} \triangleq \frac{2\bar{q}\cdot(\bar{r}+\bar{\omega})}{\varepsilon + 2\bar{q}\cdot(\bar{r}+\bar{\omega})} < 1.$$

Finally, letting $c = \frac{\bar{\beta}}{1-\bar{\beta}}$, we have $\mathbb{E}\big[\tilde{X}_i(\delta)\big] = B^m\big(\beta_i^m(\lambda(\delta))\big) \leq \frac{\beta_i^m(\lambda(\delta))}{1-\beta_i^m(\lambda(\delta))} \leq \frac{\bar{\beta}}{1-\bar{\beta}} = c$.

D.2 Proof of Proposition 5.1

Since spokes are identical, we drop the index i and let $\gamma(\beta) = \gamma_i(\beta)$ and $h^{\mathrm{S}}(\lambda) = h_i^{\mathrm{S}}(\lambda)$ for ease of notation. First, it is easy to see that $V^{\mathrm{F}} = \hat{\gamma}(1)$ and thus $V(\pi^{\mathrm{F}}) = \frac{m}{m+n}\hat{\gamma}(1)$.

Let $\rho = \frac{m}{n}$ and $\hat{\lambda} = n\lambda$. Then

$$\begin{aligned}
V^{\mathrm{S,R}} &= \min_{\lambda\geq 0}\Big\{m\lambda + n\cdot h^{\mathrm{S}}(\lambda)\Big\} \\
&= \min_{\lambda\geq 0}\Big\{m\lambda + n\cdot\max_{\beta\geq 0}\Big\{A^m(\beta)\cdot\gamma(\beta) - \lambda\cdot B^m(\beta)\Big\}\Big\} \\
&= \min_{\hat{\lambda}\geq 0}\Big\{\rho\hat{\lambda} + \max_{\beta\geq 0}\Big\{A^m(\beta)\cdot\hat{\gamma}(\beta) - \hat{\lambda}\cdot B^m(\beta)\Big\}\Big\}.
\end{aligned}$$

Lemma D.7 shows that solving the problem with $m = \infty$ provides an upper bound on $V^{\mathrm{S,R}}$ in the large network limit.

Lemma D.7 (Interchange of operations).

$$\lim_{n\to\infty} V^{\mathrm{S,R}} \le \min_{\hat{\lambda}\ge 0}\left\{\rho\hat{\lambda} + \max_{\beta\ge 0}\left\{A^{\infty}(\beta)\cdot\hat{\gamma}(\beta) - \hat{\lambda}\cdot B^{\infty}(\beta)\right\}\right\}. \tag{85}$$

We prove Lemma D.7 at the end of this section. In the following, we optimize the right-hand side of (85). Note that

$$\begin{aligned}
&\min_{\hat{\lambda}\ge 0}\left\{\rho\hat{\lambda} + \max_{\beta\ge 0}\left\{A^{\infty}(\beta)\cdot\hat{\gamma}(\beta) - \hat{\lambda}\cdot B^{\infty}(\beta)\right\}\right\}\\
=&\min_{\hat{\lambda}\ge 0}\max_{\beta\in[0,1]}\left\{\rho\hat{\lambda} + \hat{\gamma}(\beta) - \hat{\lambda}\cdot\frac{\beta}{1-\beta}\right\}\\
=&\max_{\beta\in[0,1]}\left\{\hat{\gamma}(\beta) \quad \text{s.t. } \frac{\beta}{1-\beta}\le\rho\right\},
\end{aligned} \tag{86}$$

where the first equation is because we can restrict the domain to $\beta\in[0,1]$ by Lemma D.9. Since $\hat{\gamma}(\beta)$ is concave by Lemma A.3, the second equation is due to the fact that the right-hand side is a convex program and strong duality holds. Since $\hat{\gamma}(\beta)$ together with $\frac{\beta}{1-\beta}$ is increasing in β by Lemma A.2, the optimal solution β^* satisfies $\frac{\beta^*}{1-\beta^*}=\rho$ and thus $\beta^*=\frac{\rho}{1+\rho}=\frac{m}{m+n}$. Combining (85) and (86), we have $\lim_{n\to\infty}V^{\mathrm{S,R}}\le\hat{\gamma}(\beta^*)=\hat{\gamma}\big(\frac{m}{m+n}\big)$, and this provides an upper bound on the performance of any static pricing policy in the large network regime.

On the other hand, for any n, let $\tilde{\rho}=\rho-O\big(\sqrt{\ln n/n}\big)$, $\tilde{\beta}=\frac{\tilde{\rho}}{1+\tilde{\rho}}$, and $\tilde{d}_{i0}$ and $\tilde{d}_{0i}$ an optimal solution to $\hat{\gamma}(\tilde{\beta})$. Based on the analysis in Appendix D.1 (especially equation (81) and Section D.1.2), the performance of the static policy using $\tilde{d}_{i0}$ and $\tilde{d}_{0i}$ converges to the upper bound $\hat{\gamma}\big(\frac{m}{m+n}\big)$ in the large network regime. Thus, the bound is tight and $V^{\mathrm{S}}=\hat{\gamma}\big(\frac{m}{m+n}\big)$ in the limit.

D.2.1 Proof of Lemma D.7

We let $h^m\big(\beta,\hat{\lambda}\big)\triangleq A^m(\beta)\hat{\gamma}(\beta)-\hat{\lambda}B^m(\beta)$, $h^m\big(\hat{\lambda}\big)=\max_{\beta\ge 0}h^m\big(\beta,\hat{\lambda}\big)$ and $\beta^m\big(\hat{\lambda}\big)$ denote an optimal solution to $h^m\big(\hat{\lambda}\big)$. Let $h^{\infty}\big(\beta,\hat{\lambda}\big)\triangleq A^{\infty}(\beta)\hat{\gamma}(\beta)-\hat{\lambda}B^{\infty}(\beta)$ and $h^{\infty}\big(\hat{\lambda}\big)$ and $\beta^{\infty}\big(\hat{\lambda}\big)$ denote the maximum value and point over β for a given $\hat{\lambda}$. From Lemma A.4, we have $\hat{\gamma}(\beta)\le c\triangleq n\cdot q_{0i}\cdot(\bar{r}+\bar{\omega})$ for all $\beta\ge 0$.

Lemma D.8. *For any* $\hat{\lambda}>0$, $\beta^m\big(\hat{\lambda}\big)\le\frac{2c}{\hat{\lambda}+2c}<1$ *when* m *is large enough.*

Proof. This is exactly Lemma D.5. □

Lemma D.9. $\beta^{\infty}\big(0\big)=1$ *and* $\beta^{\infty}\big(\hat{\lambda}\big)\le\frac{c}{\hat{\lambda}+c}<1$ *for any* $\hat{\lambda}>0$.

Proof. For any $\hat{\lambda}>0$ and $\beta\ge\bar{\beta}=\frac{c}{\hat{\lambda}+c}$, analogous to the proof of Lemma D.5, we have

$$h^{\infty}\big(\beta,\hat{\lambda}\big)=A^{\infty}(\beta)\hat{\gamma}(\beta)-\hat{\lambda}B^{\infty}(\beta)\le c-\hat{\lambda}\cdot\frac{\bar{\beta}}{1-\bar{\beta}}=0=h^{\infty}\big(0,\hat{\lambda}\big).$$

Thus, $\beta^{\infty}\big(\hat{\lambda}\big)\le\frac{c}{\hat{\lambda}+c}$. Suppose $\hat{\lambda}=0$. Then $h^{\infty}\big(\beta,\hat{\lambda}\big)=A^{\infty}(\beta)\hat{\gamma}(\beta)$. If $\beta\le 1$, $A^{\infty}(\beta)\hat{\gamma}(\beta)=\hat{\gamma}(\beta)$ is increasing by Lemma A.2. If $\beta\ge 1$, $A^{\infty}(\beta)\hat{\gamma}(\beta)=\frac{\hat{\gamma}(\beta)}{\beta}$ with derivative $\frac{d}{d\beta}\Big(\frac{\hat{\gamma}(\beta)}{\beta}\Big)=\frac{\beta\cdot\hat{\gamma}'(\beta)-\hat{\gamma}(\beta)}{\beta^2}\le 0$ because the numerator is non-positive by Lemma A.5. Thus, $\beta^{\infty}\big(0\big)=1$. □

Lemma D.10. *For any $\hat{\lambda} \geq 0$, $\lim_{m\to\infty} h^m(\hat{\lambda}) = h^\infty(\hat{\lambda})$.*

Proof. If $\hat{\lambda} > 0$, from Lemmas D.8 and D.9, we can restrict the domain to be $\beta \in \left[0, \frac{2c}{2c+\hat{\lambda}}\right]$. From Lemma D.3, $A^m(\beta)$ and $B^m(\beta)$ converge to $A^\infty(\beta)$ and $B^\infty(\beta)$ uniformly on $\beta \in \left[0, \frac{2c}{2c+\hat{\lambda}}\right]$ as m goes to infinity; this implies $\lim_{m\to\infty} h^m(\hat{\lambda}) = h^\infty(\hat{\lambda})$. If $\hat{\lambda} = 0$, the result follows analogously because $A^m(\beta)$ converges to $A^\infty(\beta)$ uniformly on $\beta \geq 0$ by Lemma D.3. $\square$

Now we are ready to prove Lemma D.7. By the definition of $V^{\mathrm{S,R}}$, for any $\hat{\lambda} \geq 0$ we have

$$V^{\mathrm{S,R}} \leq \rho\hat{\lambda} + \max_{\beta\geq 0}\left\{A^m(\beta)\cdot\hat{\gamma}(\beta) - \hat{\lambda}\cdot B^m(\beta)\right\}.$$

Taking limits on both sides and noting Lemma D.10, we have

$$\lim_{n\to\infty} V^{\mathrm{S,R}} \leq \rho\hat{\lambda} + \max_{\beta\geq 0}\left\{A^\infty(\beta)\cdot\hat{\gamma}(\beta) - \hat{\lambda}\cdot B^\infty(\beta)\right\}.$$

Finally, minimizing the right-hand side over $\hat{\lambda} \geq 0$ gives the desired result.

D.3 More Details on Example 5.1

Since $q_{i0} = q_{0i} = \frac{1}{2n}$ and all private values are uniformly distributed on $[0,1]$, $\gamma_i(\beta) = \frac{1}{2n}\cdot\frac{\beta}{1+\beta}$ and $d_{0i}(\beta) = \frac{\beta}{1+\beta}$ and $d_{i0}(\beta) = \frac{1}{1+\beta}$ are optimal to $\gamma_i(\beta)$. Since spokes are identical, we drop the index i and let $\gamma(\beta) = \gamma_i(\beta)$ for ease of notation. Let $\rho = \frac{m}{n} = \frac{2}{3}$, $\hat{\gamma}(\beta) = n\cdot\gamma(\beta) = \frac{1}{2}\cdot\frac{\beta}{1+\beta}$ and $\hat{\lambda} = n\lambda$.

Letting $\beta^* = \frac{\rho}{1+\rho} = \frac{m}{m+n} = \frac{2}{5}$, from Proposition 5.1 and the proof therein, $\hat{\gamma}(\beta^*) = \frac{1}{7}$ is a tight upper bound on the performance of the optimal static policy in the large network limit, and the asymptotically optimal static policy converges to $d_{i0}(\beta^*) = \frac{5}{7}$ and $d_{0i}(\beta^*) = \frac{2}{7}$.

We now provide a lower bound on the Lagrangian relaxation bound V^{R} and show that it is strictly larger than $V^{\mathrm{S,R}}$ in the large network regime. Since $\lim_{n\to\infty}\left\{V^{\mathrm{R}} - V^{\mathrm{OPT}}\right\} = 0$ by Corollary 4.7, this implies that no static pricing policy is asymptotically optimal in the regime. We can also get a simple dynamic pricing policy that is strictly better than the optimal static policy as a byproduct.

Since the number of resources per location is relatively small, it is beneficial to keep the number of resources in each spoke to be small, thus retaining some resources in the hub. Motivated by this, we consider a family of cutoff policies with some parameter k that keeps at most k resources in a spoke and uses static controls to manage these resources. Analogously, we can provide a bound on the performance of any cutoff policy with a parameter k by relaxing the constraint that the hub has non-negative resources with a dual variable $\lambda \geq 0$; the best performance bound is given by

$$V^{\mathrm{k,R}} = \min_{\hat{\lambda}\geq 0}\left\{\rho\hat{\lambda} + \max_{\beta\geq 0}\left\{A^k(\beta)\cdot\hat{\gamma}(\beta) - \hat{\lambda}\cdot B^k(\beta)\right\}\right\},$$

which is equivalent to maximizing the average revenue subject to the constraints that the expected number of resources in the hub is non-negative and we restrict to cut-off policies. The max-min inequality implies

$$\begin{aligned}
V^{\mathrm{k,R}} &\geq \max_{\beta\geq 0}\min_{\hat{\lambda}\geq 0}\left\{\rho\hat{\lambda} + A^k(\beta)\cdot\hat{\gamma}(\beta) - \hat{\lambda}\cdot B^k(\beta)\right\} \\
&= \max_{\beta\geq 0}\left\{A^k(\beta)\hat{\gamma}(\beta) \quad \text{s.t. } B^k(\beta)\leq\rho\right\} \\
&\geq A^k\big(\tilde{\beta}\big)\hat{\gamma}\big(\tilde{\beta}\big) \quad \text{where } \tilde{\beta} : B^k\big(\tilde{\beta}\big) = \rho.
\end{aligned}$$

Take $k = 2$. We can solve $\tilde{\beta} = \frac{\sqrt{33}-1}{8} \approx 0.593$ from $B^2\big(\tilde{\beta}\big) = \frac{2\tilde{\beta}^2+\tilde{\beta}}{\tilde{\beta}^2+\tilde{\beta}+1} = \rho = \frac{2}{3}$. Since $V^{\mathrm{R}} \geq V^{\mathrm{k,R}}$ for any $k \in \mathbb{N}$, we have

$$V^{\mathrm{R}} \geq V^{\mathrm{k=2,R}} \geq A^{k=2}\big(\tilde{\beta}\big)\hat{\gamma}\big(\tilde{\beta}\big) \approx 0.152 > \frac{1}{7} \geq V^{\mathrm{S,R}}.$$

Finally, analogous to the discussion at the end of Section D.2, we can construct cut-off policies with asymptotic performance equal to $A^{k=2}\big(\tilde{\beta}\big)\hat{\gamma}\big(\tilde{\beta}\big)$ by using a perturbed $\beta = \tilde{\beta} - O\big(\sqrt{\ln n/n}\big)$.

E Performance Analysis of Uniformly Related Hubs

In this section, we analyze the performance of policy $\pi(\delta)$ for the special case with *uniformly related hubs* as described in Definition E.1.

Definition E.1 (Uniformly Related Hubs). A network with uniformly related hubs is a hub-and-spoke network with multiple hubs where:

1. for each spoke $i \in [n]$, the revenue functions $r_{ij}(d)$ and $r_{ji}(d)$ are identical across hubs, i.e., $r_{ij}(d)$ are identical across $j \in [J]$ and so are $r_{ji}(d)$;
2. for each spoke $i \in [n]$, the request rates satisfy $q_{ij} = c_i \cdot q_{ji} \geq 0$ for all hubs and some constant $c_i > 0$;
3. for any two hubs j and j', the request rates and the maximum points of the revenue functions as defined in Assumption 2.1 satisfy $d^*_{jj'} \cdot q_{jj'} = d^*_{j'j} \cdot q_{j'j}$.

Assumption E.1 is a form of symmetry in which, for each spoke, the revenue functions and ratio of requests are identical across hubs. We allow, however, revenue functions and request rates to be different across spokes. Note that a sufficient condition for part 3 is that, for any two hubs j and j', the request rates satisfy $q_{jj'} = q_{j'j}$ and the revenue functions satisfy $r_{jj'}(d) = r_{j'j}(d)$. Proposition E.1 shows that with uniformly related hubs, $\mu_j = 0$ for all $j \in [J]$ constitutes an optimal solution to (16) (this is equivalent to the flow balance constraint in (16) being redundant) and the controls $d_i(x, i, j)$ and $d_i(x, j, i)$ in the Lagrangian policy derived from (16) are identical across hubs. We can show that when hubs are uniformly related, at optimality, the flow between each hub-spoke pair is balanced. Summing over each spoke we obtain that hubs are balanced. Moreover, using the Lagrangian policy in the original system, we obtain a unique stationary distribution with a special form.

Proposition E.1. *For a network with uniformly related hubs, the following properties hold:*

(a) The flow balance constraint of (16) is redundant;

(b) *For the Lagrangian policy $\pi(\delta)$, all controls $d_i(x,i,j)$ and $d_i(x,j,i)$ are identical across hubs j and $d_{jj'} = d^*_{jj'}$ for all hub-to-hub requests;*

(c) *In the original system, the Markov chain with the policy $\pi(\delta)$ has a single recurrent class $\mathcal{C}$ and is aperiodic; thus with this policy we obtain a unique stationary distribution over states, which we denote by $\mathbb{P}\big(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}\big)$, where $\mathbf{x}_{\mathrm{H}} = (x_j)_{j\in[J]} \in \mathbb{N}^J$ and $\mathbf{x}_{\mathrm{S}} = (x_i)_{i\in[n]} \in \mathbb{N}^n$ denote the number of resources in the hubs and the spokes, respectively;*

(d) *$\mathbb{P}(\cdot)$ is reversible in $\mathcal{C}$; and*

(e) *Conditioned on a state $\mathbf{x}_{\mathrm{S}} = (x_i)_{i\in[n]}$ of the spokes, $\mathbb{P}(\cdot)$ is uniform across the resources in the hubs, i.e., $\mathbb{P}\big(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}\big) = \mathbb{P}\big(\mathbf{x}'_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}\big)$ for any $\mathbf{x}_{\mathrm{H}} = (x_j)_{j\in[J]} \in \mathbb{N}^J$ and $\mathbf{x}'_{\mathrm{H}} = (x'_j)_{j\in[J]} \in \mathbb{N}^J$ with $\sum_{j\in[J]} x_j = \sum_{j\in[J]} x'_j = m - \sum_{i\in[n]} x_i$.*

We prove Proposition E.1 at the end of this section. Part (e) of this result implies that the number of resources in the spokes $\mathbf{x}_{\mathrm{S}}$ only provides information on the number of resources in the hubs $\mathbf{x}_{\mathrm{H}}$ through their summations $\sum_{i\in[n]} x_i$, and vice versa; this follows directly from the reversibility property in part (d). We will use this fact to bound the probability that any hub is depleted in the following analysis. In particular, we consider a high multiplicity model and we show with a particular choice of the parameter δ, the depletion probability $\mathbb{P}\big[X_j(\delta) = 0\big]$ of each hub j of the original system diminishes to zero as the number of spokes n increases, the ratio $\frac{m}{n}$ remains fixed, and the number of hubs J grows at rate $o(n)$.

In the high multiplicity model, we assume the n spokes can be divided into S distinct spoke types and the number of spokes of each type s is fixed to be a proportion $\alpha_s > 0$ of n. All spokes of a given type have the same revenue functions and have the same arrival rates into and out of each hub (these rates may vary across hubs).

From Corollary 4.5 and Proposition 4.6, we have that the probability that all hubs run out of resources in the original system goes to zero when n increases and the ratio $\frac{m}{n}$ remains fixed. Because the resources in the hubs are uniformly distributed according to Proposition E.1, part (e), the depletion probability of each hub j of the original system diminishes to zero as well, as we show in Proposition E.2.

Proposition E.2. *Let $\underline{\alpha} = \min_{s\in[S]} \alpha_s$ and let $b = \frac{1}{1+m/(\underline{\alpha}\cdot n)}$. In the high multiplicity model, the depletion probabilities of each hub j of the original system are equal and satisfy*

$$\mathbb{P}\Big[X_j(\delta) = 0\Big] \le \exp\left(-\frac{b}{8}\cdot\frac{\delta^2}{m+n-\frac{1}{2}\delta}\right) + \frac{J-1}{\frac{1}{2}\delta + J - 1}.$$

Putting Theorem 6.1 and Proposition E.2 together, we obtain the following result.

Corollary E.3. *Under the high multiplicity model, the Lagrangian policy $\pi(\delta)$ with $0 \le \delta < m$ satisfies*

$$V^{\pi}(\delta) \le V^{\mathrm{OPT}} \le V^{\pi}(\delta) + \bar{r}\cdot\frac{\delta}{m-\delta} + (\bar{r}+\bar{\omega})\cdot\left\{\exp\left(-\frac{b}{8}\cdot\frac{\delta^2}{m+n-\frac{1}{2}\delta}\right) + \frac{J-1}{\frac{1}{2}\delta+J-1}\right\},$$

where b is as in Proposition E.2. Moreover, if m and n grow at the same rate and J grows at a *sub-linear rate with $J = o(n)$, by choosing* $\delta = \max\left\{\sqrt{nJ}, \sqrt{\frac{4}{b}\cdot(m+n)\cdot\ln n}\right\}$, *we have*

$$V^{\mathrm{OPT}} - V^{\pi}(\delta) \le O\Bigg(\max\left\{\sqrt{\frac{J}{n}}, \sqrt{\frac{\ln n}{n}}\right\}\Bigg).$$

E.1 Proof of Proposition E.1

Parts (a) & (b): We first show in Lemma E.4 that the Lagrangian policies with dual variables $\mu_j = 0$ for all $j \in [J]$ and any $\lambda \ge 0$ have the desired properties as stated in Proposition E.1, parts (a) and (b).

Lemma E.4. *Let $d_i(x,i,j)$, $d_i(x,j,i)$, and $d_{jj'}$ be the controls of a Lagrangian policy with dual variables $\mu_j = 0$ for all $j \in [J]$ and some $\lambda \ge 0$. We have*

1. *$d_i(x,i,j)$ and $d_i(x,j,i)$ are identical across hubs; and*
2. *$d_{jj'} = d^*_{jj'}$; and*
3. *the in-flow and out-flow of each hub j is balanced in expectation, i.e., for each $j \in [J]$,*

$$\sum_{i=1}^{n} q_{ij}\sum_{x=0}^{m} p_i(x)\cdot d_i(x,i,j) + \sum_{j'=1}^{J} q_{j'j}\cdot d_{j'j} = \sum_{i=1}^{n} q_{ji}\sum_{x=0}^{m} p_i(x)\cdot d_i(x,j,i) + \sum_{j'=1}^{J} q_{jj'}\cdot d_{jj'}.$$

Proof. First, $d_{jj'} = \operatorname{argmax}_{d\in[0,1]}\Big\{r_{jj'}(d) + d\cdot(\mu_{j'} - \mu_j)\Big\} = d^*_{jj'}$ when $\mu_j = 0$ for all $j \in [J]$. Next, let $p_i(x)$ be an optimal probability distribution to the spoke i problem. The controls $d_i(x,j,i)$ and $d_i(x+1,i,j)$ are optimal to the concave problem

$$\begin{aligned}
\max_{d_{ij},d_{ji}\in[0,1]} \quad & p_i(x)\cdot\sum_{j=1}^{J} q_{ji}\cdot r_{ji}(d_{ji}) + p_i(x+1)\cdot\sum_{j=1}^{J} q_{ij}\cdot r_{ij}(d_{ij})\\
\text{s.t.} \quad & p_i(x)\cdot\sum_{j=1}^{J} q_{ji}\cdot d_{ji} = p_i(x+1)\cdot\sum_{j=1}^{J} q_{ij}\cdot d_{ij}.
\end{aligned}$$

Since $r_{ij}(d)$ and $r_{ji}(d)$ are strictly concave by Assumption 2.1, the solution is unique. Moreover, since these revenue functions are identical across hubs by Definition E.1, Jensen's inequality implies that $d_i(x,i,j)$ and $d_i(x,j,i)$ are identical across hubs. Otherwise the average controls $\frac{\sum_{j=1}^{J} q_{ji}\cdot d_{ji}}{\sum_{j=1}^{J} q_{ji}}$ and $\frac{\sum_{j=1}^{J} q_{ij}\cdot d_{ij}}{\sum_{j=1}^{J} q_{ij}}$ are feasible and yield a strictly better objective. Finally, summing up the flow balance constraint in (15):

$$p_i(x)\cdot\sum_{j=1}^{J} q_{ji}\cdot d_i(x,j,i) = p_i(x+1)\cdot\sum_{j=1}^{J} q_{ij}\cdot d_i(x+1,i,j)$$

on both sides over $x \in [0:m-1]$ and noting that $d_i(0,i,j) = 0$ and $d_i(m,j,i) = 0$ for all $j \in [J]$, we have

$$\sum_{x=0}^{m} p_i(x) \cdot \sum_{j=1}^{J} q_{ji} \cdot d_i(x,j,i) = \sum_{x=0}^{m} p_i(x) \cdot \sum_{j=1}^{J} q_{ij} \cdot d_i(x,i,j),$$

which implies that the flow is balanced for each spoke. Since the controls are identical across hubs and the request rates satisfy $q_{ij} = c_i \cdot q_{ji}$ for all $j \in [J]$ and some constant $c_i > 0$ by Definition E.1, the flow between each hub-spoke pair is balanced, i.e., for any $j \in [J]$ and $i \in [n]$

$$\sum_{x=0}^{m} p_i(x) \cdot q_{ji} \cdot d_i(x,j,i) = \sum_{x=0}^{m} p_i(x) \cdot q_{ij} \cdot d_i(x,i,j).$$

Summing both sides over spokes plus the fact that $d_{jj'} = d^*_{jj'}$ and the third condition of Definition E.1 implies that the flow is balanced at each hub. $\square$

To show the flow balance constraint of (16) is redundant, it suffices to show the constraint holds by itself if we solve (16) ignoring this constraint. This problem corresponds to minimizing the objective of (16) over $\lambda \geq 0$ while keeping $\mu_j = 0$ fixed for all $j \in [J]$. Lemma E.4 indicates that the flow balance constraint indeed holds by itself, and all the properties of the controls stated in parts (a) and (b) are satisfied.

Part (c): The properties of the controls stated in parts (a) and (b) imply that the resulting Markov chain under the Lagrangian policy enjoys some helpful properties. Let $\mathbf{x}_{\mathrm{H}} = (x_j)_{j \in [J]} \in \mathbb{N}^J$ and $\mathbf{x}_{\mathrm{S}} = (x_i)_{i \in [n]} \in \mathbb{N}^n$ denote the number of resources in the hubs and the spokes, respectively, and let the system state be the resource levels $(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}) \in \mathbb{N}^{J+n}$. Due to the assumption that the network topology is strongly connected and the fact that the controls for requests between a hub and a spoke are identical across the hubs, it is easy to see the Markov chain with the policy $\pi(\delta)$ has a single recurrent class $\mathcal{C} = \Big\{(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}) \in \mathbb{N}^{n+J} : \sum_{i \in [n]} x_i + \sum_{j \in [J]} x_j = m,\ x_i \in I_i \text{ for all } i \in [n]\Big\}$ and is aperiodic. Hence the Lagrangian policy is a unichain policy, and with this policy the limiting distribution converges to a unique stationary distribution, which we denote by $\mathbb{P}\big(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}\big)$, independent of the initial state.

Part (d): We say that $\mathbb{P}(\cdot)$ is reversible in $\mathcal{C}$ if for any two states $(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}), (\mathbf{x}'_{\mathrm{H}}, \mathbf{x}'_{\mathrm{S}}) \in \mathcal{C}$,

$$\mathbb{P}\big(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}\big) \cdot p\big(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}, \mathbf{x}'_{\mathrm{H}}, \mathbf{x}'_{\mathrm{S}}\big) = \mathbb{P}\big(\mathbf{x}'_{\mathrm{H}}, \mathbf{x}'_{\mathrm{S}}\big) \cdot p\big(\mathbf{x}'_{\mathrm{H}}, \mathbf{x}'_{\mathrm{S}}, \mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}\big),$$

where $p\big(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}, \mathbf{x}'_{\mathrm{H}}, \mathbf{x}'_{\mathrm{S}}\big)$ denotes the transition probability from a state $(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}})$ to a state $(\mathbf{x}'_{\mathrm{H}}, \mathbf{x}'_{\mathrm{S}})$ with the Lagrangian policy. This part is a direct consequence of Theorem 6.5.1 in Durrett (2010), which provides a necessary and sufficient condition for an irreducible Markov chain to have a reversible measure.

Lemma E.5 (Theorem 6.5.1 in Durrett 2010). *Consider an irreducible Markov chain with states denoted by* $\mathbf{x}$ *and transition probabilities denoted by* p. *A necessary and sufficient condition for the existence of a reversible measure is that*

1. $p(\mathbf{x}, \mathbf{x}') > 0$ *implies* $p(\mathbf{x}', \mathbf{x}) > 0$ *for any two states* $\mathbf{x}$ *and* $\mathbf{x}'$*; and*
2. *for any loop of states* $\mathbf{x}_0, \mathbf{x}_1, \cdots, \mathbf{x}_N = \mathbf{x}_0$ *with* $\prod_{k=1}^{N} p(\mathbf{x}_k, \mathbf{x}_{k-1}) > 0$, $\prod_{k=1}^{N} \frac{p(\mathbf{x}_{k-1}, \mathbf{x}_k)}{p(\mathbf{x}_k, \mathbf{x}_{k-1})} = 1$.

We show the transition probabilities p induced by the Lagrangian policy satisfy the conditions in Lemma E.5 within the recurrent class $\mathcal{C}$. The first part simply follows from the fact that

$q_{jj'} \cdot d^*_{jj'} = q_{j'j} \cdot d^*_{j'j}$ from part 3 of Definition E.1, and $d_i(x-1,j,i), d_i(x,i,j) > 0$ for any $x \in I_i$ due to part (b).

We now prove the second part. Let s_k denote the request that induces the transition from state $\mathbf{x}_k$ to $\mathbf{x}_{k+1}$, $\tilde{s}_k$ denote the reverse trip of s_k, and $x_{k,i}$ denote the number of resources in spoke i when the state is $\mathbf{x}_k$. Let $\mathcal{S}_h = \{s_k = (j,j') : j,j' \in [J]\}$ denote the set of requests between the hubs and $\mathcal{S}_i = \{s_k = (i,j) \text{ or } (j,i) : j \in [J]\}$ denote the set of requests between a hub and spoke i. By partitioning requests into these sets, it is equivalent to show that

$$\underbrace{\prod_{k:s_k \in \mathcal{S}_h} q_{s_k} \cdot d^*_{s_k}}_{\varepsilon_h} \cdot \prod_{i \in [n]} \underbrace{\prod_{k:s_k \in \mathcal{S}_i} q_{s_k} \cdot d_i\big(x_{k,i}, s_k\big)}_{\varepsilon_i} = \underbrace{\prod_{k:s_k \in \mathcal{S}_h} q_{\tilde{s}_k} \cdot d^*_{\tilde{s}_k}}_{\varepsilon'_h} \cdot \prod_{i \in [n]} \underbrace{\prod_{k:s_k \in \mathcal{S}_i} q_{\tilde{s}_k} \cdot d_i\big(x_{k+1,i}, \tilde{s}_k\big)}_{\varepsilon'_i}.$$

It suffices to show that $\varepsilon_h = \varepsilon'_h$ and $\varepsilon_i = \varepsilon'_i$ for all $i \in [n]$. $\varepsilon_h = \varepsilon'_h$ is clear from the fact that $q_{jj'} \cdot d^*_{jj'} = q_{j'j} \cdot d^*_{j'j}$ in part 3 of Definition E.1. To see that $\varepsilon_i = \varepsilon'_i$ for any $i \in [n]$, note that $\mathbf{x}_0 = \mathbf{x}_N$ implies that $\mathbf{x}_{0,i} = \mathbf{x}_{N,i}$, i.e., the resources transiting out of spoke i equals the resources transiting into spoke i. Based on this, it is easy to show that $\prod_{k:s_k \in \mathcal{S}_i} q_{s_k} = \prod_{k:s_k \in \mathcal{S}_i} q_{\tilde{s}_k}$ from part 2 of Definition E.1, and $\prod_{k:s_k \in \mathcal{S}_i} d_i\big(x_{k,i}, s_k\big) = \prod_{k:s_k \in \mathcal{S}_i} d_i\big(x_{k+1,i}, \tilde{s}_k\big)$ from part (b) of Proposition E.1; thus $\varepsilon_i = \varepsilon'_i$.

Part (e): Suppose spoke i is connected to hubs j and j' and $x_j \geq 1$. From part (d), the Markov chain is reversible and hence we have

$$\mathbb{P}\big(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}\big) \cdot q_{ji} \cdot d_i\big(x_i, j, i\big) = \mathbb{P}\big(\mathbf{x}_{\mathrm{H}} - \mathbf{e}_j, \mathbf{x}_{\mathrm{S}} + \mathbf{e}_i\big) \cdot q_{ij} \cdot d_i\big(x_i + 1, i, j\big),$$

and

$$\mathbb{P}\big(\mathbf{x}_{\mathrm{H}} - \mathbf{e}_j + \mathbf{e}_{j'}, \mathbf{x}_{\mathrm{S}}\big) \cdot q_{j'i} \cdot d_i\big(x_i, j', i\big) = \mathbb{P}\big(\mathbf{x}_{\mathrm{H}} - \mathbf{e}_j, \mathbf{x}_{\mathrm{S}} + \mathbf{e}_i\big) \cdot q_{ij'} \cdot d_i\big(x_i + 1, i, j'\big).$$

Since the controls $d_i(x,i,j)$ and $d_i(x,j,i)$ are identical across hubs by part (b) and the request rates satisfy $q_{ij}/q_{ji} = q_{ij'}/q_{j'i} = c_i > 0$ from Definition E.1, $\mathbb{P}\big(\mathbf{x}_{\mathrm{H}}, \mathbf{x}_{\mathrm{S}}\big) = \mathbb{P}\big(\mathbf{x}_{\mathrm{H}} - \mathbf{e}_j + \mathbf{e}_{j'}, \mathbf{x}_{\mathrm{S}}\big)$. The general result follows from the fact that the network topology is strongly connected.

E.2 Proof of Proposition E.2

First, for the high multiplicity model, since the expected number of resources in the spokes of the relaxed system is no larger than $m - \delta$, we have $\mathbb{E}\big[\tilde{X}_i(\delta)\big] \leq \frac{m}{\alpha \cdot n}$ for all $i \in [n]$. Let random variables $X_0(\delta) = \sum_{j'=1}^{J} X_{j'}(\delta)$ and $\tilde{X}_0(\delta)$ denote the sum of resources in the hubs of the original system and the relaxed system under the stationary distributions of the Lagrangian policy $\pi(\delta)$. Proposition E.1(e) implies that conditional on the total number of resources in the hubs, the distribution of resources across the hubs is uniform. Therefore, we have

$$\mathbb{P}\left[X_j(\delta) = 0 \middle| X_0(\delta) = k\right] = \frac{\binom{k+J-2}{J-2}}{\binom{k+J-1}{J-1}} = \frac{J-1}{k+J-1}.$$

Thus,

$$\begin{aligned}
\mathbb{P}\big[X_j(\delta)=0\big] &= \sum_{k=0}^{m} \mathbb{P}\big[X_0(\delta)=k\big]\cdot \mathbb{P}\Big[X_j(\delta)=0\Big|X_0(\delta)=k\Big] \\
&= \sum_{k=0}^{m} \mathbb{P}\big[X_0(\delta)=k\big]\cdot \frac{J-1}{k+J-1} = \mathbb{E}\left[\frac{J-1}{X_0(\delta)+J-1}\right] \\
&\le \mathbb{E}\left[\mathbb{1}[X_0(\delta)\le c] + \mathbb{1}[X_0(\delta)>c]\cdot \frac{J-1}{X_0(\delta)+J-1}\right] \\
&\le \mathbb{P}[\tilde{X}_0(\delta)\le c] + \frac{J-1}{c+J-1}\,,
\end{aligned} \tag{87}$$

for any constant $c \ge 0$, where the last inequality is because $X_0(\delta) \succeq_{FOSD} \tilde{X}_0(\delta)$ from Lemma 4.4, dropping the second indicator, and using that $(J-1)/(x+J-1)$ is decreasing for $x \ge 0$.

Let $\mu = \mathbb{E}\Big[\sum_{i=1}^{n} \tilde{X}_i(\delta)\Big]$ be the expected number of resources in the spokes of the relaxed system. Since the policy $\pi(\delta)$ is solved from the perturbed Lagrangian relaxation, we have $0 < \mu \le m-\delta$. For any $\gamma \in [0,1]$, set $c = \gamma\cdot(m-\mu)$ in (87) and we have

$$\begin{aligned}
\mathbb{P}\big[X_j(\delta)=0\big] &\le \mathbb{P}\Big[\tilde{X}_0(\delta)\le \gamma\cdot(m-\mu)\Big] + \frac{J-1}{\gamma\cdot(m-\mu)+J-1} \\
&\le \mathbb{P}\Big[\tilde{X}_0(\delta)\le \gamma\cdot(m-\mu)\Big] + \frac{J-1}{\gamma\delta+J-1}.
\end{aligned} \tag{88}$$

We can bound the first term in (88) using the concentration inequality developed in Lemma A.12 of Appendix A.10. Specifically, applying Lemma A.12 with $\lambda = \frac{m-\gamma\cdot(m-\mu)}{\mu}$ and $b = \frac{1}{1+m/(\underline{\alpha}\cdot n)}$ gives

$$\begin{aligned}
\mathbb{P}\Big[\tilde{X}_0(\delta)\le \gamma\cdot(m-\mu)\Big] &= \mathbb{P}\left[\sum_{i=1}^{n}\tilde{X}_i(\delta) \ge m-\gamma\cdot(m-\mu)\right] \\
&\le \exp\left\{-b\cdot\left(\lambda\mu-\mu+(n+\mu)\cdot\ln\left(1-\frac{\lambda\mu-\mu}{\lambda\mu+n}\right)\right)\right\} \\
&= \exp\left\{-b\cdot\left((1-\gamma)\cdot(m-\mu)+(n+\mu)\cdot\ln\left(\frac{n+\mu}{m+n-\gamma\cdot(m-\mu)}\right)\right)\right\} \\
&= \exp\Bigg\{b\cdot\Bigg(\underbrace{(n+\mu)\cdot\ln\left(\frac{m+n-\gamma\cdot(m-\mu)}{n+\mu}\right)-(1-\gamma)\cdot(m-\mu)}_{\clubsuit}\Bigg)\Bigg\}.
\end{aligned} \tag{89}$$

Since $\ln x \le \frac{x-1}{\sqrt{x}}$ for $x \ge 1$, we have

$$\begin{aligned}
\clubsuit &\le (n+\mu)\cdot\frac{(1-\gamma)\cdot(m-\mu)}{n+\mu}\cdot\sqrt{\frac{n+\mu}{m+n-\gamma\cdot(m-\mu)}}-(1-\gamma)\cdot(m-\mu)\\
&=(1-\gamma)\cdot(m-\mu)\cdot\left(\sqrt{1-\frac{(1-\gamma)\cdot(m-\mu)}{m+n-\gamma\cdot(m-\mu)}}-1\right)\\
&\le-\frac{(1-\gamma)^2\cdot(m-\mu)^2}{2\cdot(m+n-\gamma\cdot(m-\mu))}\\
&\le-\frac{(1-\gamma)^2\cdot\delta^2}{2\cdot(m+n-\gamma\delta)},
\end{aligned}$$

where the second-to-last inequality is due to $\sqrt{1-x}-1\le-\frac{x}{2}$ for $x\le 1$. Thus from (89) we have

$$\mathbb{P}\Big[\tilde{X}_0(\delta)\le\gamma\cdot(m-\mu)\Big]\le\exp\left(-\frac{b}{2}\cdot\frac{(1-\gamma)^2\cdot\delta^2}{m+n-\gamma\delta}\right). \tag{90}$$

Combining (88) and (90) we have

$$\mathbb{P}\big[X_j(\delta)=0\big]\le\exp\left(-\frac{b}{2}\cdot\frac{(1-\gamma)^2\cdot\delta^2}{m+n-\gamma\delta}\right)+\frac{J-1}{\gamma\delta+J-1}.$$

Letting $\gamma=\frac{1}{2}$ gives the desired result.

F Exponential Relocation Times in a Single Hub Network

With exponential relocation times, the Lagrangian simplifies greatly as we only need to track the number of resources on each route and in each location. Suppose we have one hub and n spokes, and the relocation times for requests $(i,0)$ and $(0,i)$ follow independent exponential distributions with mean values τ_{i0} and τ_{0i}, respectively. Let $\Lambda=\sum_{i\in[n]}\eta_{i0}+\eta_{0i}$ denote the total request rate.

For each spoke i problem, we let x_i denote the number of resources in the spoke and x_{0i} denote the number of resources in transit from the hub to the spoke; we require that $x_i+x_{0i}\le m$ to bound the state space. The resources that are leaving the spoke are irrelevant to the spoke problem. Suppose the current state is (x_i,x_{0i}) and no request is fulfilled in the current period. Let $\rho(x_i,x_{0i},\tilde{x}_i,\tilde{x}_{0i})$ denote the transition probability that the next period starts with a state $(\tilde{x}_i,\tilde{x}_{0i})$. We have

$$\rho(x_i,x_{0i},\tilde{x}_i,\tilde{x}_{0i})=\begin{cases}\binom{x_{0i}}{\tilde{x}_{0i}}\left(\frac{\Lambda}{\Lambda+1/\tau_{0i}}\right)^{\tilde{x}_{0i}}\left(\frac{1/\tau_{0i}}{\Lambda+1/\tau_{0i}}\right)^{x_{0i}-\tilde{x}_{0i}} & \text{if } \tilde{x}_{0i}\le x_{0i} \text{ and } \tilde{x}_i+\tilde{x}_{0i}=x_i+x_{0i},\\ 0 & \text{otherwise,}\end{cases}$$

because each resource on $(0,i)$ will reach spoke i with probability $\frac{1/\tau_{0i}}{\Lambda+1/\tau_{0i}}$ and keep relocating with probability $\frac{\Lambda}{\Lambda+1/\tau_{0i}}$ by the end of the current period. The spoke problem is:

$$h_i^\lambda=\max_{\substack{d_i(x_i,x_{0i},i,0)\in[0,1],\\ d_i(x_i,x_{0i},0,i)\in[0,1],\\ p_i(x_i,x_{0i})\ge 0}}\sum_{(x_i,x_{0i})}p_i(x_i,x_{0i})\cdot\Big\{q_{i0}\cdot\Big[r_{i0}\Big(d_i(x_i,x_{0i},i,0)\Big)-\lambda\cdot d_i(x_i,x_{0i},i,0)\cdot\Lambda\cdot\tau_{i0}\Big]$$

$$+ q_{0i} \cdot r_{0i}\Big(d_i(x_i, x_{0i}, 0, i)\Big) - \lambda \cdot (x_i + x_{0i})\Big\}$$

s.t.

$$\sum_{(x_i, x_{0i}) \in \mathbb{N}^2 : x_i + x_{0i} \le m} p_i(x_i, x_{0i}) = 1,$$

$$p_i(x_i, x_{0i}) = \sum_{(\tilde{x}_i, \tilde{x}_{0i})} p_i(\tilde{x}_i, \tilde{x}_{0i}) \cdot \Big[q_{i0} \cdot d_i(\tilde{x}_i, \tilde{x}_{0i}, i, 0) \cdot \rho(\tilde{x}_i - 1, \tilde{x}_{0i}, x_i, x_{0i})$$

$$+ q_{0i} \cdot d_i(\tilde{x}_i, \tilde{x}_{0i}, 0, i) \cdot \rho(\tilde{x}_i, \tilde{x}_{0i} + 1, x_i, x_{0i})$$

$$+ \Big(1 - q_{i0} \cdot d_i(\tilde{x}_i, \tilde{x}_{0i}, i, 0) - q_{0i} \cdot d_i(\tilde{x}_i, \tilde{x}_{0i}, 0, i)\Big) \cdot \rho(\tilde{x}_i, \tilde{x}_{0i}, x_i, x_{0i})\Big],$$

$$d_i(x_i, x_{0i}, i, 0) = 0, \ \forall\, x_i = 0,$$

$$d_i(x_i, x_{0i}, 0, i) = 0, \ \forall\, x_i + x_{0i} = m,$$

where in the objective, the term $\sum_{(x_i, x_{0i})} \Lambda \cdot q_{i0} \cdot p_i(x_i, x_{0i}) \cdot d_i(x_i, x_{0i}, i, 0) \cdot \tau_{i0}$ equals the expected number of resources moving from the spoke to the hub by Little's law.

G More on Numerical Examples

In this section, we plot the stationary distributions of the number of resources in the single hub example (Section 7.1) in Section G.1, we consider another synthetic example with two hubs in Section G.2, and we provide more numerical results for the RideAustin example (Section 7.2) in Section G.3.

G.1 Stationary Distributions in the Single Hub Examples

Figure 6 shows the stationary distributions of the number of resources in the hub of the one hub examples in Section 7.1 under the Lagrangian policy $\pi(\delta)$ and the static policy π^{F}. Based on the results in Whitt (1984), we actually have analytical expressions for the marginal distributions with the static policy: the probability that there are x resources in any location $i \in [0:n]$ is $\binom{m+n-1-k}{n-1}/\binom{m+n}{n}$ (note that $n+1$ locations are in the example), with the mode being that the location has zero resources.

G.2 Two Hub Examples

In this section, we consider examples with two hubs as illustrated in Figure 7; these hubs are not uniformly related as defined in Definition E.1. We let the arrival rates be $q_{i1} = \frac{1}{3n}$ and $q_{1i} = \frac{1}{6n}$ for hub 1, and $q_{i2} = \frac{1}{6n}$ and $q_{2i} = \frac{1}{3n}$ for hub 2, for all spokes $i \in [n]$, and let all the other arrival rates be zero; thus without the flow "balancing" constraints captured by the dual variables $\boldsymbol{\mu}$, hub 1 tends to accumulate resources whereas hub 2 tends to lose resources.

For each fixed n, we calculate the same quantities as in the one hub examples (Section 7.1), and we additionally calculate: the Lagrangian relaxation upper bound $\tilde{V}^{\mathrm{R}}$ omitting the flow balance constraints at the hubs, and the performance $\tilde{V}^{\pi}(\delta)$ of the policy derived from the perturbed problem, with $\delta = \sqrt{n \ln n}$, omitting the flow balance constraints at the hubs. Figure 8 shows the simulation results for the two-hub case. From Figure 8, when there are multiple hubs and the hubs are asymmetric, omitting the flow balance constraints at the hubs leads to a loose upper bound and that the corresponding Lagrangian policy does not perform well.

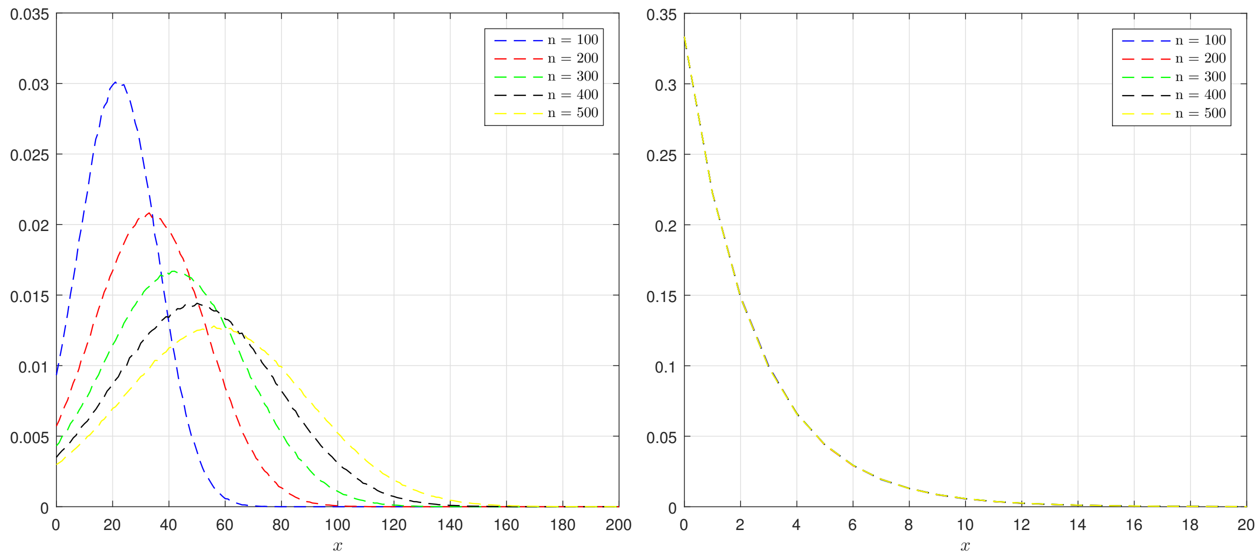

(a)

(b)

Figure 6: Stationary distributions of resources in the hub for the one hub examples (Section 7.1), with varying number of spokes n. (a) is with the Lagrangian policy $\pi(\delta)$ and (b) is with the static policy π^{F}.

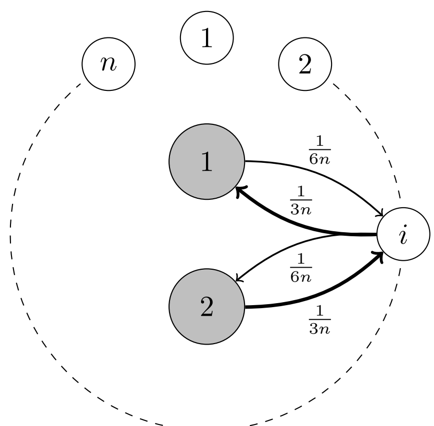

Figure 7: A hub-and-spoke network with 2 hubs (grey) and n symmetric spokes. We only draw the connections between spoke i and the hubs. Edge widths illustrate relative values of request rates.

Figures 9 and 10 show the stationary distributions under the Lagrangian policy that omits the flow balance constraints of the hubs and the Lagrangian policy that incorporates these constraints, respectively. Since hub 1 tends to accumulate resources whereas hub 2 tends to lose resources, without flow balancing, hub 1 have excessive resources and hub 2 is essentially depleted.

G.3 More on the RideAustin Example

Figure 11 shows the partition of Austin, Texas with $n = 100$ locations and the ride flow of the city based on the partition. In the ride flow figure, each node represents a location of the city. The radius of a node is proportional to the amount of the requests that leave the location, and the width of an edge is proportional to the size of the requests on the edge. An edge has the same color as its origin location.

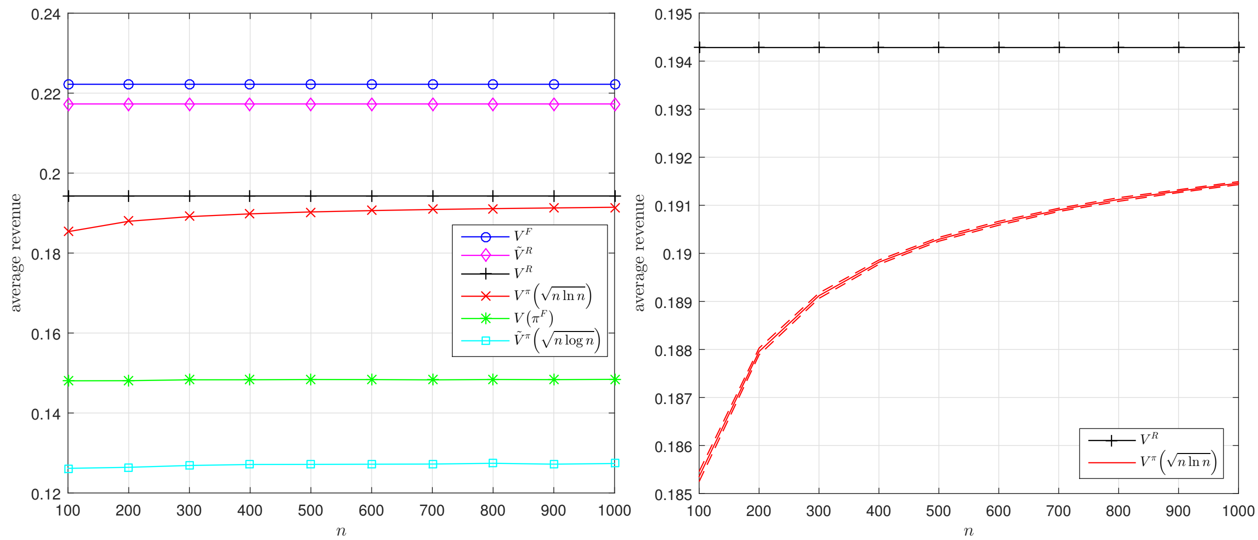

(a) (b)

Figure 8: Simulation results of the two-hub case. (b) is magnified versions of (a), highlighting the performance of our policy and the Lagrangian relaxation upper bound. 95% confidence intervals around $V^\pi\left(\sqrt{n \ln n}\right)$ are plotted with dashed lines in (b).

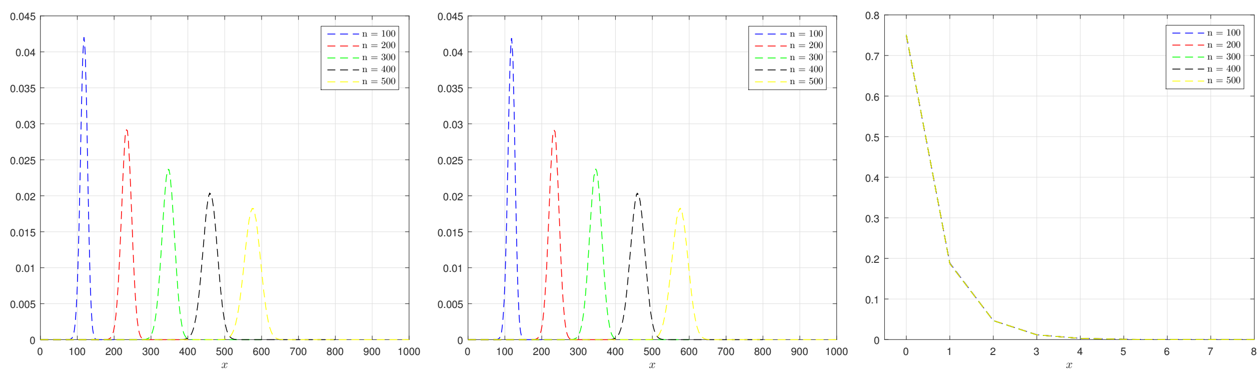

(a) (b) (c)

Figure 9: Stationary distributions with the Lagrangian policy that omits the flow balance constraints in the hubs, for the two hub examples (Section G.2). (a) is for the total number of resources in the hubs, (b) is for the number of resources in hub one, and (c) is for the number of resources in hub two.

Figure 12 illustrates the locations of hubs obtained from solving (18) with different values of J from one to six. For each value of J, the locations of the J hubs are the nodes labelled from one to J.

Figure 13 demonstrates how the performances of the Lagrangian policy $\pi(\delta)$ and the Lagrangian-based static policy $\pi^{\mathrm{S}}(\delta)$ varies with δ for each value of J.

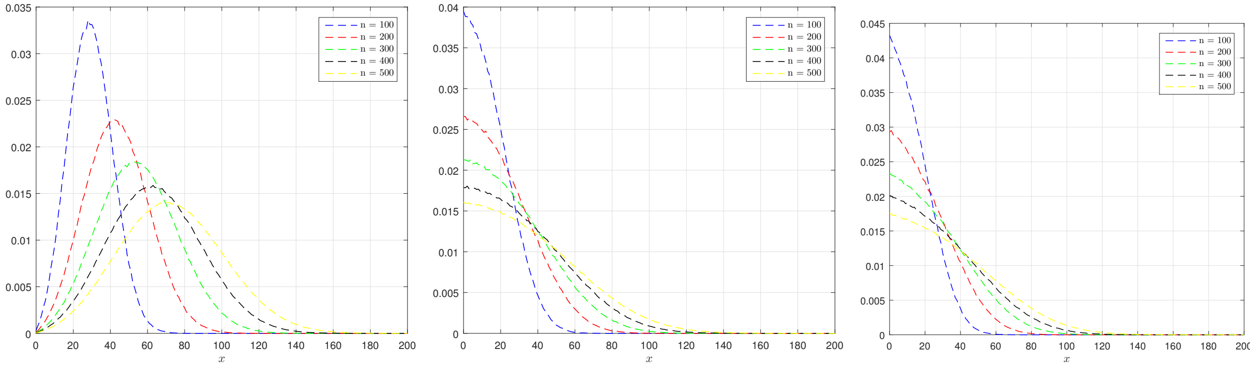

(a) (b) (c)

Figure 10: Stationary distributions with the Lagrangian policy $\pi(\delta)$ that incorporates the flow balance constraints in the hubs, for the two hub examples (Section G.2). (a) is for the total number of resources in the hubs, (b) is for the number of resources in hub one, and (c) is for the number of resources in hub two.

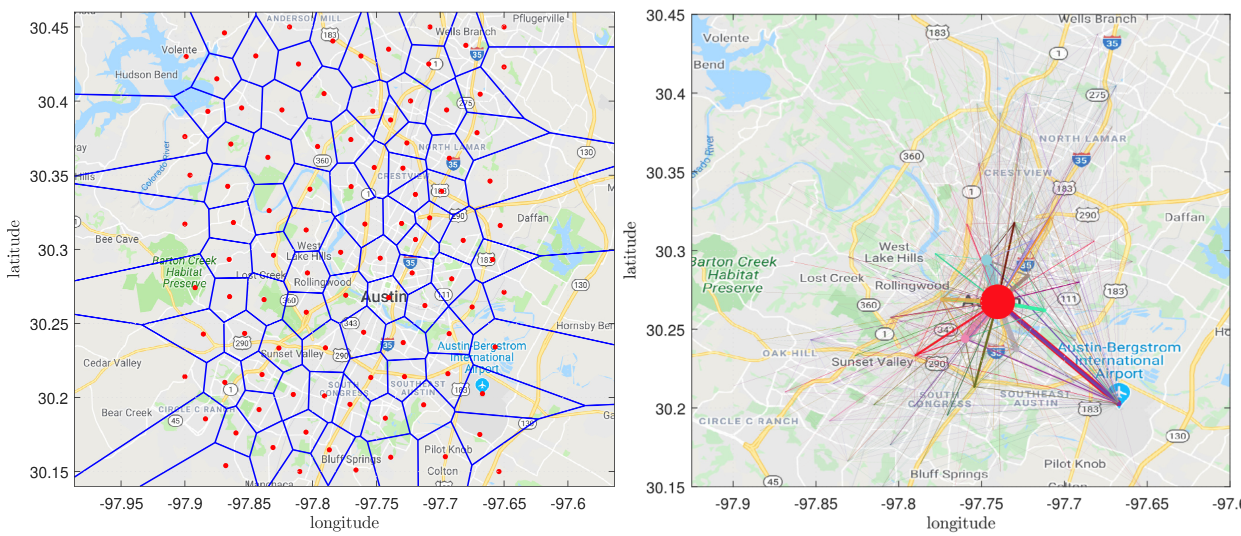

(a) (b)

Figure 11: (a) The Voronoi diagram of the cluster centers from solving the k-center problem with $k = 100$ and the first few centers initialized with k-means clustering centers. (b) The ride flow of the city based on data from RideAustin and the partition.

References

Adelman, D. (2007), ‘Price-directed control of a closed logistics queueing network’, *Operations Research* **55**(6), 1022–1038.

Adelman, D. and Mersereau, A. J. (2008), ‘Relaxations of weakly coupled stochastic dynamic programs’, *Operations Research* **56**(3), 712–727.

An, Y., Zhang, Y. and Zeng, B. (2015), ‘The reliable hub-and-spoke design problem: Models and algorithms’, *Transporation Research Part B* **77**, 103–122.

Anderson, E. J. and Nash, P. (1987), *Linear programming in infinite-dimensional spaces: theory and applications*, John Wiley & Sons.

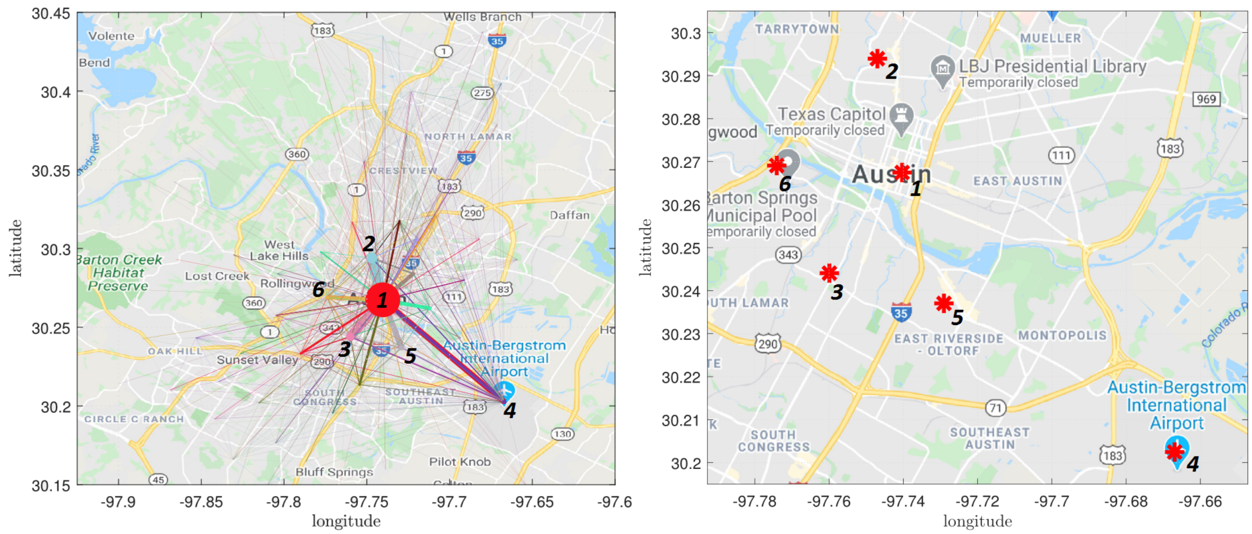

(a) (b)

Figure 12: Locations of hubs obtained from solving (18) with different values of J from one to six. For each value of J, the locations of the J hubs are the nodes labelled from one to J. (b) is simply a zoom-in of (a).

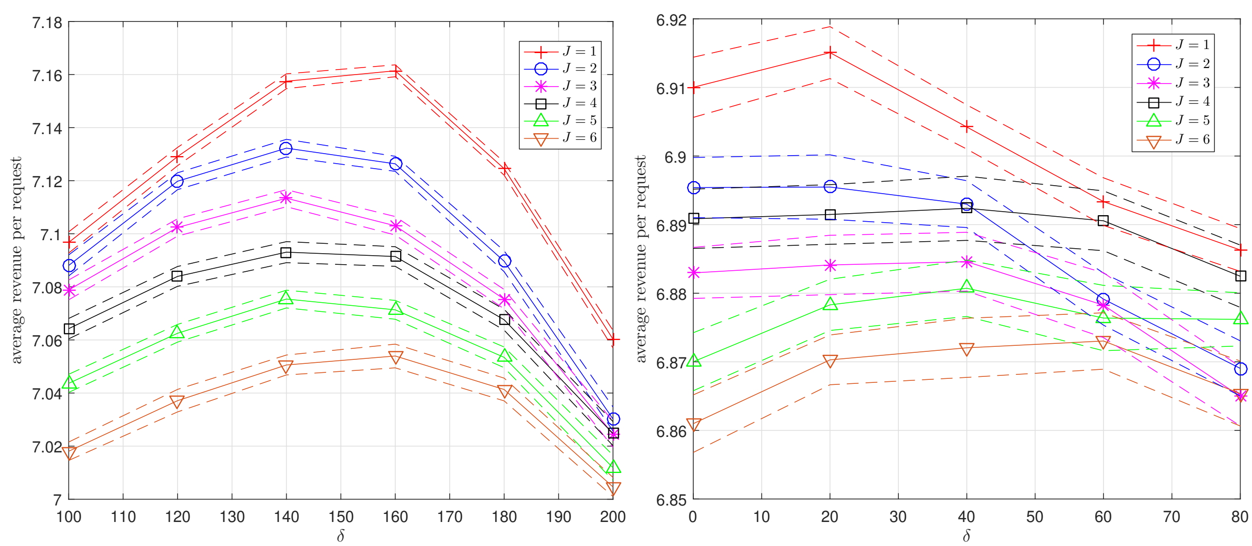

(a) (b)

Figure 13: (a) The performance (average revenue per request) $V^{\pi}(\delta)$ of the Lagrangian policy $\pi(\delta)$ for different values of δ and J. The optimal choice δ^* is roughly 160, 140, 140, 140, 140 and 160 for J from one to six. (b) The performance $V\big(\pi^{\mathrm{S}}(\delta)\big)$ of the Lagrangian-based static policy $\pi^{\mathrm{S}}(\delta)$ for different values of δ and J. The optimal choice δ^*_{S} is roughly 20, 20, 40, 40, 40 and 60 for J from one to six. We estimate the values $V^{\pi}(\delta)$ and $V\big(\pi^{\mathrm{S}}(\delta)\big)$ with 50 sample paths and for each sample path, we approximate the average revenue with an average of the total revenue of the first 10^6 requests. 95% confidence intervals are plotted with dashed lines.

Banerjee, S., Freund, D. and Lykouris, T. (2016), 'Pricing and optimization in shared vehicle systems: An approximation framework', *arXiv preprint arXiv:1608.06819* .

Banerjee, S., Kanoria, Y. and Qian, P. (2018), State dependent control of closed queueing networks, *in* 'Abstracts of the 2018 ACM International Conference on Measurement and Modeling of Computer Systems', ACM, pp. 2–4.

Baskett, F., Chandy, K. M., Muntz, R. R. and Palacios, F. G. (1975), 'Open, closed, and mixed networks of queues with different classes of customers', *Journal of the ACM (JACM)* **22**(2), 248–260.

Bertsekas, D. (2012), *Dynamic programming and optimal control*, Vol. 2, 4 edn, Athena scientific Belmont, MA.

Bertsekas, D. P., Nedic, A. and Ozdaglar, A. E. (2003), 'Convex analysis and optimization'.

Bertsimas, D. and Mersereau, A. J. (2007), 'A learning approach for interactive marketing to a customer segment', *Operations Research* **55**(6), 1120–1135.

Bertsimas, D. and Mišić, V. V. (2016), 'Decomposable markov decision processes: A fluid optimization approach', *Operations Research* **64**(6), 1537–1555.

Besbes, O., Castro, F. and Lobel, I. (2018), 'Surge pricing and its spatial supply response'. Available at `https://ssrn.com/abstract=3124571`.

Besbes, O., Castro, F. and Lobel, I. (2019), 'Spatial capacity planning'. Available at `https://ssrn.com/abstract=3292651`.

Bimpikis, K., Candogan, O. and Daniela, S. (2016), 'Spatial pricing in ride-sharing networks'. *Operations Research*, to appear.

Boyd, S. and Vandenberghe, L. (2004), *Convex optimization*, Cambridge university press.

Braverman, A., Dai, J. G., Liu, X. and Ying, L. (2016), 'Empty-car routing in ridesharing systems', *arXiv preprint arXiv:1609.07219* .

Brown, D. B. and Smith, J. E. (2017), 'Index policies and performance bounds for dynamic selection problems'. *Management Science*, to appear.

Caro, F. and Gallien, J. (2007), 'Dynamic assortment with demand learning for seasonal consumer goods', *Management Science* **53**(2), 276–292.

Du, Y. and Hall, R. (1997), 'Fleet sizing and empty equipment redistribution for center-terminal transportation networks', *Management Science* **43**(2), 145–157.

Durrett, R. (2010), *Probability: theory and examples*, 4 edn, Cambridge university press.

George, D. K. and Xia, C. H. (2011), 'Fleet-sizing and service availability for a vehicle rental system via closed queueing networks', *European journal of operational research* **211**(1), 198–207.

Gordon, W. J. and Newell, G. F. (1967), 'Closed queuing systems with exponential servers', *Operations research* **15**(2), 254–265.

Harrison, J. M. and Wein, L. M. (1990), 'Scheduling networks of queues: Heavy traffic analysis of a two-station closed network', *Operations research* **38**(6), 1052–1064.

Hawkins, J. T. (2003), A Langrangian decomposition approach to weakly coupled dynamic optimization problems and its applications, PhD thesis, Massachusetts Institute of Technology.

Hu, W. and Frazier, P. (2017), 'An asymptotically optimal index policy for finite-horizon restless bandits', *arXiv preprint arXiv:1707.00205* .

Janson, S. (2018), ‘Tail bounds for sums of geometric and exponential variables’, *Statistics & Probability Letters* **135**, 1 – 6.

Kanoria, Y. and Qian, P. (2019), ‘Near optimal control of a ride-hailing platform via mirror backpressure’, *arXiv preprint arXiv:1903.02764* .

Keilson, J. (1972), ‘A threshold for log-concavity for probability generating functions and associated moment inequalities’, *The Annals of Mathematical Statistics* **43**(5), 1702–1708.

Keilson, J. and Gerber, H. (1971), ‘Some results for discrete unimodality’, *Journal of the American Statistical Association* **66**(334), 386–389.

Kumar, S. and Kumar, P. (1996), Closed queueing networks in heavy traffic: Fluid limits and efficiency, *in* ‘Stochastic Networks’, Springer, pp. 41–64.

Kunnumkal, S. and Talluri, K. (2016), ‘On a piecewise-linear approximation for network revenue management’, *Mathematics of Operations Research* **41**(1), 72–91.

Luenberger, D. G. (1997), *Optimization by vector space methods*, John Wiley & Sons.

Marklund, J. and Rosling, K. (2012), ‘Lower bounds and heuristics for supply chain stock allocation’, *Operations Research* **60**(1), 92–105.

Miao, S., Jasin, S. and Chao, X. (2020), ‘Asymptotically optimal lagrangian policies for one-warehouse multi-store system with lost sales’. Available at `https://ssrn.com/abstract=3552995`.

Milgrom, P. and Segal, I. (2002), ‘Envelope theorems for arbitrary choice sets’, *Econometrica* **70**(2), 583–601.

Milgrom, P. and Shannon, C. (1994), ‘Monotone comparative statics’, *Econometrica: Journal of the Econometric Society* pp. 157–180.

Ozkan, E. and Ward, A. (2017), ‘Dynamic matching for real-time ridesharing’. Working Paper.

Pirkul, H. and Schilling, D. A. (1998), ‘An efficient procedure for designing single allocation hub and spoke systems’, *Management Science* **44**(2), 235–242.

RideAustin (2017), ‘Ride-austin-june6-april13 [data file and code book]’. Available at `https://data.world/ride-austin/ride-austin-june-6-april-13`.

Royden, H. L. and Fitzpatrick, P. (2010), *Real analysis*, 4 edn, Boston : Prentice Hall.

Shaked, M. and Shanthikumar, J. G. (2007), *Stochastic orders*, Springer Science & Business Media.

Song, D.-P. and Carter, J. (2008), ‘Optimal empty vehicle redistribution for hub-and-spoke transportation systems’, *Naval Research Logistics (NRL)* **55**(2), 156–171.

Talluri, K. T. and Van Ryzin, G. J. (2006), *The theory and practice of revenue management*, Vol. 68, Springer Science & Business Media.

Talluri, K. and van Ryzin, G. (1998), ‘An analysis of bid-price controls for network revenue management’, *Management Science* **44**(11), 1577–1593.

Topaloglu, H. (2009), ‘Using lagrangian relaxation to compute capacity-dependent bid prices in network revenue management’, *Operations Research* **57**(3), 637–649.

Topkis, D. M. (1978), ‘Minimizing a submodular function on a lattice’, *Operations research* **26**(2), 305–321.

Topkis, D. M. (2011), *Supermodularity and complementarity*, Princeton university press.

Tran, T. H., O'Hanley, J. R. and Scaparra, M. P. (2017), 'Reliable hub network design: Formulation and solution techniques', *Transportation Science* **51**(1), 358–375.

Vera, A. and Banerjee, S. (2018), 'The bayesian prophet: A low-regret framework for online decision making'.

Waserhole, A. and Jost, V. (2016), 'Pricing in vehicle sharing systems: Optimization in queuing networks with product forms', *EURO Journal on Transportation and Logistics* **5**(3), 293–320.

Whitt, W. (1984), 'Open and closed models for networks of queues', *AT&T Bell Laboratories Technical Journal* **63**(9), 1911–1979.

Zayas-Cabán, G., Jasin, S. and Wang, G. (2019), 'An asymptotically optimal heuristic for general nonstationary finite-horizon restless multi-armed, multi-action bandits', *Advances in Applied Probability* **51**(3), 745–772.