This article was downloaded by: [132.236.27.111] On: 31 May 2019, At: 12:17
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

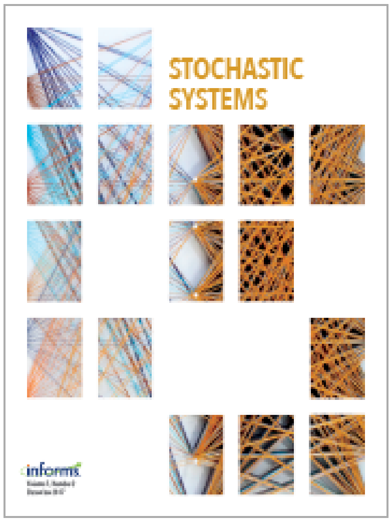

Stochastic Systems

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Mean Field Equilibria for Resource Competition in Spatial Settings

Pu Yang, Krishnamurthy Iyer, Peter Frazier

To cite this article:
Pu Yang, Krishnamurthy Iyer, Peter Frazier (2018) Mean Field Equilibria for Resource Competition in Spatial Settings. Stochastic Systems 8(4):307-334. https://doi.org/10.1287/stsy.2018.0018

Full terms and conditions of use: https://pubsonline.informs.org/page/terms-and-conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2018, The Author(s)

Please scroll down for article—it is on subsequent pages

INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

informs
http://pubsonline.informs.org/journal/stsy/

STOCHASTIC SYSTEMS
Vol. 8, No. 4, December 2018, pp. 307–334
ISSN 1946-5238 (online)

Mean Field Equilibria for Resource Competition in Spatial Settings

Pu Yang,[a] Krishnamurthy Iyer,[a] Peter Frazier[a]

[a] School of Operations Research and Information Engineering, Cornell University, Ithaca, New York 14850

Contact: py75@cornell.edu, http://orcid.org/0000-0002-2093-3719 (PY); ki82@cornell.edu, http://orcid.org/0000-0002-5538-1432 (KI); pf98@cornell.edu, http://orcid.org/0000-0002-3501-3341 (PF)

Received: July 16, 2017
Revised: April 11, 2018; July 26, 2018
Accepted: July 28, 2018
Published Online in Articles in Advance: December 27, 2018

https://doi.org/10.1287/stsy.2018.0018

Copyright: © 2018 The Author(s)

Abstract. We study a model of competition among nomadic agents for time-varying and location-specific resources arising in crowdsourced transportation services, online communities, and traditional location-based economic activity. This model comprises a group of agents and a single location endowed with a dynamic stochastic resource process. Periodically, each agent derives a reward determined by the location's resource level and the number of other agents there and has to decide whether to stay at the location or move. On moving, the agent arrives at a different location whose dynamics are independent of and identical to the original location. Using the methodology of mean field equilibrium, we study the equilibrium behavior of the agents as a function of the dynamics of the stochastic resource process and the nature of the competition among colocated agents. We show that an equilibrium exists in which each agent decides whether to switch locations based only on the agent's current location's resource level and the number of other agents there. We additionally show that when an agent's payoff is decreasing in the number of other agents at the agent's location, equilibrium strategies obey a simple threshold structure. We show how to exploit this structure to compute equilibria numerically and use these numerical techniques to study how system structure affects the agents' collective ability to explore their domain to find and effectively utilize resource-rich areas.

Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as "*Stochastic Systems*. Copyright © 2018 The Author(s). https://doi.org/10.1287/stsy.2018.0018, used under a Creative Commons Attribution License: https://creativecommons.org/licenses/by/4.0/."

Funding: K. Iyer gratefully acknowledges support from the National Science Foundation [Grants CMMI-1462592 and CMMI-1633920]. P. Frazier gratefully acknowledges support from the National Science Foundation [Grant CMMI-1254298] and the Air Force Office of Scientific Research [Grant FA9550-15-1-0038].

Keywords: resource sharing • mean field equilibrium • exploration • spatial economics

1. Introduction

We consider a model of nomadic agents exploring and competing for time-varying, stochastic, location-specific resources. Such multiagent systems arise in many real-world settings as illustrated herein.

They arise in the sharing economy in crowdsourced transportation services, such as Uber and Lyft, and in crowdsourced food delivery services, such as GrubHub and DoorDash, in which drivers choose neighborhoods and then earn money based on the number of riders or eaters requesting service within that neighborhood (the location-specific resource) and the number of other drivers working there. This overall resource level varies stochastically as demand rises and falls, and the resource derived by a driver decreases as more drivers drive in the neighborhood.

They also arise in the traditional economy, for example, in mobile food vendors deciding where to locate their trucks, in pastoralists deciding where to graze their livestock, and in fishermen deciding where to fish. In these examples, the level of resource derived by each agent from the agent's location (whether profit from hungry passersby, food for livestock provided by the range land, or profit from the catch) depends on both the number of other agents at the location and the location's stochastically varying resource level.

They also arise in online communities, such as Reddit and Twitch, in which participants choose subcommunities or channels and then derive enjoyment depending on both some underlying but transitory societal interest in the subcommunity's topic of focus (the overall resource) and the number of other participants in the subcommunity. When the number of other participants is too small, lack of social interaction prevents enjoyment; when the number of other participants is too large, crowding diminishes the sense of community.

They even arise among scientific researchers, who choose a research area in which to work and derive value based on the underlying level of societal interest and funding in their chosen area and the number of other researchers working in it. As with online communities, the number of other researchers should be neither too large nor too small to maximize the value derived.

In each of these examples, the overall welfare of the system is determined by how agents explore their domain to find and exploit resource-rich locations. This willingness to explore, in turn, depends on the level of competition or cooperation among agents at the same location and the distribution of agents and resources across locations.

In this paper, we develop a formal model to analyze such spatiotemporal competition among agents and the equilibrium behavior of such systems. The model we study comprises a single location and a group of agents. This location represents one in a large collection of locations between which the agents move. It has a resource level that varies stochastically with time. Each agent at the location periodically obtains a payoff whose amount is determined by the number of other agents currently at the location and the location's current resource level. Based on these quantities, the agent then decides whether to stay at the same location or leave. On leaving, the agent receives a reward that represents the expected future discounted payoff that would be obtained by moving to another randomly chosen location in the system. The agents are fully strategic and seek to maximize the total expected payoff over their lifetime.

Using the methodology of mean field equilibrium (MFE), we study the equilibrium behavior of the agents in this system as a function of the dynamics of the spatiotemporal resource process and the level of competition in the agents' sharing of a location's resources. We prove the existence of an equilibrium for general resource-sharing functions. For the specific case in which the resource-sharing function is nonincreasing in the number of agents at the location, we further show that the equilibrium strategy has a simple threshold structure in which it is optimal for an agent to leave a location when the number of other agents there exceeds a threshold that depends on the location's resource level. This result enables a simple description of equilibrium strategies, and allows us to efficiently compute an equilibrium.

Using numerical analysis of a setting with two resource levels and decreasing resource-sharing function, we investigate how the equilibrium welfare depends on resource levels' rate of change and the density of agents. Here, the equilibrium welfare is the sum of payoffs earned across all agents in equilibrium, normalized to the length of time over which these payoffs have accrued and either the number of agents or the number of locations. Using this methodology, we show qualitatively different system behavior when the single-location welfare function (the contribution to welfare from all agents at one location) increases with the number of agents at the location as compared with when it decreases. Our ability to derive these and other insights discussed in detail in Section 5 provide evidence that our model and equilibrium notion lend themselves to analysis through simple numerical methods. Specifically, our methodology presents a promising approach to evaluate engineering interventions, such as providing subsidies to or imposing costs on agents to promote or discourage exploration to improve welfare.

1.1. Related Work

Our work contributes to the literature on MFE (Jovanovic and Rosenthal 1988, Huang et al. 2007, Lasry and Lions 2007, Weintraub et al. 2008, Adlakha et al. 2015) that studies complex systems under a large system limit and obtains insights about agent behavior that are hard to obtain from analyzing finite models. The main insight behind this literature—that, in the large system limit, agents' behavior is characterized by their private state and an aggregate distribution of the rest of the system—has been used to study settings including industry dynamics and oligopoly models (Hopenhayn 1992; Weintraub et al. 2008, 2011), repeated dynamics auctions (Iyer et al. 2014, Balseiro et al. 2015), online labor markets (Arnosti et al. 2014), queueing (Xu and Hajek 2013, Manjrekar et al. 2014), content sharing (Li et al. 2017), and pedestrian motion (Lachapelle and Wolfram 2011), among others. In these papers, the unit of analysis is a single agent's decision problem, assuming the behavior of all other agents together constitutes a mean field distribution. In contrast, in our work, the unit of analysis is the game among the agents at a single location, assuming that the behavior of agents and the resource level at all other locations constitutes a mean field distribution.

Our work also contributes to the literature on spatial models of ride sharing and crowdsourced transportation (Banerjee et al. 2015, 2016; Braverman et al. 2016). In this literature, the paper most closely related to ours is Bimpikis et al. (2016), who consider a ride-sharing platform with a continuum of riders and drivers spread across a finite network of locations and study how the platform should set origin-based prices to maximize profits. In particular, the drivers' decision of where, when, and whether to provide service is explicitly modeled. The paper studies the impact of the underlying network structure of the locations on the platform's profits and consumers' surplus under the assumption that the demand at each location is stationary. In contrast, in our model, the resources at each location (analogous to demand) are stochastic and time varying. However, in our model, agents decide whether to stay or switch from their current location and not to which location to switch.

Our model is also related to congestion games (Rosenthal 1973, Nisan et al. 2007), in which agents choose paths on which to travel and then incur costs that depend on the number of other agents who have chosen the same path. One may view paths as being synonymous with locations in our model and observe that, in both cases, the utility/cost derived from a path/location depends on the number of other agents using that path or portion thereof. The main difference between our model and congestion games is the stochastic time-varying nature of our overall level of resource (making our model more complex) and the lack of interaction between locations contrasting with the interaction between paths (making our model simpler).

Another related strand of literature studies ecological models of metapopulations in static and dynamic habitats (Levin 1970; Durrett and Levin 1994a, b; Molofsky 1994). Keymer et al. (2000) consider a set of habitats, arranged on a lattice, each containing a subpopulation of a species, and the landscape structure of each habitat is stochastic and dynamic. Using a mean field analysis and through numerical simulations, the authors study the dependence of persistence and extinction rates of the species across habitats as a function of the rate of change of the landscape. In such models, the species dynamics are exogenously specified, whereas we are interested in the equilibrium behavior of agents.

Our work can be seen as an extension of the Kolkata Paise restaurant problem (Chakrabarti et al. 2009), a generalization of the El Farol bar problem (Arthur 1994, Chakrabarti 2007). In this game, each agent chooses (simultaneously) a restaurant to visit and earns a reward that depends on both the restaurant's fixed rank, which is common across agents, and the number of other agents at that restaurant. This reward is inversely proportional to the number of agents visiting the restaurant. The Kolkata Paise restaurant problem is studied in both the one-shot and repeated settings with results on the limiting behavior of myopic (Chakrabarti et al. 2009) and other strategies (Ghosh et al. 2010) although we are not aware of existing results on mean field equilibria in this model. The model we consider is both more general in that we allow general reward functions and allow a location's resource to vary stochastically and more specific in that our locations are homogeneous. Our model also differs in that our agents' decisions are made asynchronously.

2. Model

Our formal model is motivated by considering a system of locations occupied by strategic agents. Each location has a resource level that varies over time according to a finite-state, continuous-time Markov chain and is occupied by a time-varying collection of agents. Each agent has an associated sequence of independent decision epochs separated by exponential times. At the start of an agent's decision epoch, the agent receives a payoff that depends on the number of other agents at the agent's current location and the resource level there. The agent then decides whether to stay at the location or to leave and move to another location. When moving, the agent's destination is chosen uniformly at random from the set of all locations other than the origin. Each agent seeks to maximize the agent's expected payoff over an independent, exponentially distributed lifetime. When an agent's lifetime expires, the agent exits the system and is replaced by a new agent who arrives to a new location chosen uniformly at random. In Appendix A, we provide a detailed description of this system with a finite number of agents and locations.

Because the payoff obtained by an agent at any location is determined by the number of agents at that location, each agent's decision to stay in the agent's current location or to move to a new one depends on all the other agents' behavior. Consequently, the interaction among the agents in this finite model is a dynamic game, and describing the agents' behavior requires an equilibrium analysis. Because the agents are not fully informed about the resource levels at other locations, the standard equilibrium concept to analyze the induced dynamic game is a *perfect Bayesian equilibrium* (PBE). A PBE consists of a strategy ξ^i and a belief system μ^i for each player i. A belief system μ^i for agent i specifies a belief $\mu^i(h_t^i)$ after any history h_t^i over all aspects of the system of which the agent is uncertain and that influence the agent's expected payoff. A PBE then requires two conditions to hold: (1) each agent i's strategy ξ^i is a best response after any history h_t^i given their belief system and given all other agents' strategies, and (2) each agent i's beliefs $\mu^i(h_t^i)$ are updated via Bayes' rule whenever possible (see Fudenberg and Tirole 1991a, b for more details).

A PBE supposes a complex model of agent behavior. Each agent keeps track of the agent's entire history and maintains complex beliefs about the rest of the system. Although this behavioral model may be plausible in small settings, in large systems an agent's history may not contain too much information about the state of all other locations because the agent would typically only visit a small fraction of the locations. In such settings, it is more plausible that each agent would base the decision to stay or switch solely on the current state of the location the agent is in—specifically on its level of resource and congestion—and on the aggregate features of the entire system. Moreover, we expect that an agent would prefer to stay at a location with a high resource level and few other agents.

We seek to uncover this intuitive behavioral model as an equilibrium in large systems by letting the number of agents and the number of locations both increase proportionally to infinity and studying the limiting infinite system.

As the number of locations and agents grows to infinity proportionally (with the proportionality constant $\beta > 0$ defined as the *agent density*), it is reasonable to suppose that the dynamics at any fixed finite collection of locations is independent asymptotically and that the rewards experienced by an agent can be described by modeling the dynamics at a single location and then supposing that upon leaving that location the agent moves to another location whose dynamics are independent and identically distributed ad infinitum until the agent's lifetime expires. Thus, to analyze a large finite system, we posit a formal model for the dynamic of a single location and treat each agent who leaves this location as returning to an independent copy.

2.1. Formal Model of a Single Location in the Limiting Infinite System

Here we state our formal model of a single location k. Let Z_t^k denote the resource level at the location at time $t \geq 0$. We assume the resource process $\{Z_t^k : t \geq 0\}$ is a finite-state, continuous-time Markov chain. We let $\mathbb{Z}$ denote the set of values the resource process can take. Furthermore, we let $\mu_{zy} > 0$ denote the transition rate of Z_t^k from a state $z \in \mathbb{Z}$ to a state $y \in \mathbb{Z}$. We assume that the process Z_t^k is irreducible and positive recurrent with a unique invariant distribution $\{\pi_{\mathsf{res}}(z) : z \in \mathbb{Z}\}$.

We let N_t^k denote the number of agents at the location k at time t. The stochastic process (Z_t^k, N_t^k) will evolve according to arrivals to this location and the decisions made by agents at this location. Toward that end, we suppose that new agents arrive to this location according to a Poisson process with rate κ, and we describe the agents' decision process herein. The rate κ models both arrivals of agents switching from other locations in a finite system and new arrivals of agents to the system following the exit of other agents from the system, but here it is taken to be an input to the formal model of a single location, and later it is required to satisfy consistency conditions in equilibrium.

Associated with each agent i at location k is a Poisson clock with rate λ such that each time the clock rings, the agent decides whether to stay in the location or leave. We refer to each clock ring of agent i as the agent's *decision epoch* and let τ_i^ℓ denote the time of the agent's ℓ^{th} decision epoch.

At each decision epoch τ_i^ℓ, the agent i receives a payoff $F(Z_t^k, N_t^k)|_{t=\tau_i^\ell}$ that depends on the resource level Z_t^k and the number of agents N_t^k. We refer to the function F as the *resource-sharing function*. We assume that the resource-sharing function is nonnegative; that is, $F(z, n) \geq 0$ for each $z \in \mathbb{Z}$ and $n \geq 1$. To avoid trivialities, we require that there exists a (z_0, n_0) such that $F(z_0, n_0) > 0$. Finally, to model the competitive nature of interaction among the agents, we assume that as the number of agents at a location increases, the payoff an agent receives approaches zero: $\lim_{n\to\infty} F(z, n) = 0$ for each $z \in \mathbb{Z}$.

To model agents with finite lifetimes, we assume that subsequent to receiving a payoff at time τ_i^ℓ with probability $1 - \gamma \in (0, 1)$, the agent's lifetime expires, and the agent exits the system permanently. Thus, each agent i can exist in the system for at most a random time interval distributed exponentially with rate $\lambda(1 - \gamma)$. We refer to $\gamma \in (0, 1)$ as the *survival probability*.

If the agent's lifetime does not expire, then agent i decides whether to stay at the location or move. Agents are free to make this choice based on their history of past observations. If the agent stays, then the dynamics and payoffs described continue forward for another decision epoch. If the agent leaves, then the agent is awarded a one-time payoff of $V_{\mathsf{sw}} > 0$ and no subsequent payoffs. Here V_{sw} is taken simply to be a constant input to our model for a single location, and later it is required to satisfy a condition at equilibrium. This condition corresponds to V_{sw} being the conditional expected payoff experienced by an agent when moving to a new location whose current number of agents and resource level are distributed according to the stationary distribution induced by equilibrium agent behavior.

2.2. The Single-Location Decision Problem When Other Agents Follow Markovian Strategies

Having specified the arrival process and agents' decision process in a single location, we are interested in characterizing a symmetric equilibrium among agents. For a given arrival rate κ and the switching payoff V_{sw}, the particular notion of equilibrium we consider is a Markov perfect equilibrium (Fudenberg and Tirole 1991a), in which, in equilibrium, each agent finds it optimal to base the decision only on the current state of the location at the agent's decision epoch and not on the agent's past (although the agent is not restricted from doing so). Formally, let $\mathbb{S} = \mathbb{Z} \times \mathbb{N}_0$ denote the set of possible states of the process (Z_t^k, N_t^k). A Markovian strategy for an agent is a function $\xi : \mathbb{S} \to [0, 1]$, where $\xi(z, n)$ denotes the probability with which the agent chooses to stay if the state of the location at the agent's decision epoch is $(z, n) \in \mathbb{S}$. (Note that $\xi(z, 0)$ is not well defined; by convention, we let $\xi(z, 0) = 1$ for all $z \in \mathbb{Z}$).

As a step toward formulating the game among the agents, we first study the dynamics at a location when all agents in location k adopt a Markovian strategy ξ. Given the arrival rate κ and the Markovian strategy ξ, the process (Z_t^k, N_t^k) for any location k evolves as a continuous-time Markov chain on the state space $\mathbb{S}$ with the following transition-rate matrix $\mathsf{Q}^{\xi,\kappa}$:

$$\begin{aligned}\mathsf{Q}^{\xi,\kappa}((z,n)\to(x,m)) &= \mathbf{I}\{x\neq z, m=n\}\mu_{z,x} + \mathbf{I}\{x=z, m=n+1\}\kappa \\ &+ \mathbf{I}\{x=z, m=n-1\}\lambda n(1-\gamma\xi(z,n)) \\ &- \mathbf{I}\{x=z, m=n\}\left(\sum_{y\neq z}\mu_{z,y} + \kappa + \lambda n(1-\gamma\xi(z,n))\right),\end{aligned} \tag{1}$$

where $z, x \in \mathbb{Z}$ and $n, m \in \mathbb{N}_0$. Here the first term on the right-hand side represents the transition in the resource level Z_t^k at the location, which is an independent Markov chain with rates $\mu_{z,x}$. The second term on the right-hand side represents the arrival of an agent to location k at rate κ. The third term on the right-hand side represents the departure of one of the n agents from location k. Such a departure can only occur at a decision epoch of one of these agents. At any such decision epoch, an agent stays with probability $\xi(z,n)$ times the survival probability γ. Thus, with probability $1-\gamma\xi(z,n)$, the agent leaves location k. Because there are n agents at the location, each of whose decision epoch occurs at rate λ, the total rate for a departure at the location is given by $\lambda n(1-\gamma\xi(z,n))$. Finally, the last term on the right-hand side represents the rate of no transition. We denote this continuous-time Markov chain describing the dynamics of a single location, in which all agents adopt the Markovian strategy ξ and the rate of arrival of agents is κ by $\mathsf{MC}(\xi,\kappa)$.

Now consider the decision problem faced by a single agent i at location k, assuming all other agents (current as well as in future) at the location follow strategy ξ. For any fixed switching payoff $V_{\mathsf{sw}} > 0$ and arrival rate κ, the decision problem faced by an agent i can be described as follows. As long as the agent stays at location k, at each decision epoch τ_i^ℓ, the agent receives a payoff $F(Z_i^\ell, N_i^\ell)$ and must choose whether to stay in location k or switch. Also, irrespective of this decision, the agent's lifetime expires with probability $1-\gamma$. On choosing to stay with survival probability γ, the agent continues until the agent's next decision epoch $\tau_i^{\ell+1}$. On choosing to switch with survival probability γ, the agent immediately receives the switching payoff V_{sw}. From this description, it follows that the decision problem facing an agent i in location k is an optimal stopping problem. Denote this optimal stopping problem by $\mathsf{DEC}(\xi,\kappa,V_{\mathsf{sw}})$. In the following, we develop the dynamic-programming formulation of this problem.

We begin by defining the value functions for the agent. Let $V(z,n)$ denote the value function of agent i at the agent's decision epoch prior to the agent making a decision or receiving payoffs given resource level $z\in\mathbb{Z}$ and number of agents $n\in\mathbb{N}$ at location k. Similarly, we let $V_{\mathsf{st}}(z,n)$ denote the continuation payoff of the agent at the agent's decision epoch, subsequent to the agent making the decision to stay and conditional on the agent not leaving the system given resource level z and number of agents n at location k. We have the following Bellman's equation for the optimal stopping problem $\mathsf{DEC}(\xi,\kappa,V_{\mathsf{sw}})$ faced by the agent:

$$\begin{aligned}V(z,n) &= F(z,n) + \gamma\max\{V_{\mathsf{st}}(z,n), V_{\mathsf{sw}}\}\\ V_{\mathsf{st}}(z,n) &= \mathbf{E}^{\xi}[V(Z_\tau, N_\tau)\,|\,z,n],\end{aligned} \tag{2}$$

where $\mathbf{E}^{\xi}[\cdot\,|\,z,n]$ denotes the expectation with respect to the process defined by (1) subject to $(Z_0,N_0)=(z,n)$, and τ denotes the time of the first decision epoch of the agent i. Here the first equation follows from the fact that at the decision epoch the agent receives an immediate payoff equal to $F(z,n)$ and has to make the decision whether to stay or switch. Subsequent to the decision, the agent survives in the system with probability γ. On choosing to switch and surviving, the agent receives a continuation payoff equal to V_{sw}. By contrast, on choosing to stay and surviving, the agent receives a continuation payoff equal to $V_{\mathsf{st}}(z,n)$. The second equation relates $V_{\mathsf{st}}(z,n)$ to the expectation of the agent's value function at the next decision epoch.

For value functions V and V_{st} satisfying the Bellman's equation (2) and any optimal strategy ξ_i, agent i chooses to stay if the resource level z and the number of agents n in the location satisfies $V_{\mathsf{st}}(z,n) > V_{\mathsf{sw}}$, to switch if $V_{\mathsf{st}}(z,n) < V_{\mathsf{sw}}$, and any mixed action if $V_{\mathsf{st}}(z,n) = V_{\mathsf{sw}}$. We let $\mathsf{OPT}(\xi,\kappa,V_{\mathsf{sw}})$ denote the set of all optimal strategies for the agent's decision specified by (2). Specifically, for any Markovian strategy $\hat{\xi}$, we have $\hat{\xi}\in\mathsf{OPT}(\xi,\kappa,V_{\mathsf{sw}})$ if and only if the following conditions hold: $\hat{\xi}(z,n)=1$ if $V_{\mathsf{st}}(z,n) > V_{\mathsf{sw}}$, $\hat{\xi}(z,n)=0$ if $V_{\mathsf{st}}(z,n) < V_{\mathsf{sw}}$, and $\hat{\xi}(z,n)\in(0,1)$ only if $V_{\mathsf{st}}(z,n)=V_{\mathsf{sw}}$.

2.3. Mean Field Equilibrium

With the description of the model in place, we are now ready to formally define the notion of equilibrium on which we focus. First, for any arrival rate κ and switching payoff V_{sw}, we require the agents to play a Markov perfect equilibrium at location k. In other words, we require the strategy ξ to satisfy the following requirement: assuming all agents other than an agent i follow the strategy ξ, agent i maximizes the agent's payoff (across all possible history-dependent strategies) by following the strategy ξ. This leads us to the following condition:

$$\xi \in \mathsf{OPT}(\xi, \kappa, V_{\mathsf{sw}}). \tag{3}$$

Now suppose, for a given κ and V_{sw}, a Markov perfect equilibrium ξ is being played at location k. Then the dynamics of the location's state are given by $\mathsf{MC}(\xi, \kappa)$. Let $\pi(\xi, \kappa)$ denote the steady-state distribution of this process. In particular, for $z \in \mathbb{Z}$ and $n \geq 0$, we let $\pi_{z,n}(\xi, \kappa)$ denote the probability that the location has a resource level z and number of agents n in steady state. (We drop the explicit dependence of the steady-state distribution on ξ and κ when the context is clear.) Thus, $\pi(\xi, \kappa)$ is an invariant distribution under $\mathsf{Q}^{\xi,\kappa}$ and satisfies

$$\sum_{z\in\mathbb{Z}} \sum_{n\in\mathbb{N}_0} \pi_{z,n}(\xi, \kappa)\mathsf{Q}^{\xi,\kappa}((z, n) \to (x, m)) = 0, \quad \text{for all } x \in \mathbb{Z}, m \in \mathbb{N}_0. \tag{4}$$

Now consider an agent arriving to the location k in steady state $\pi(\xi, \kappa)$. We denote the total expected payoff that this agent receives over the agent's lifetime on following the strategy ξ by V_{arr}. Using the definition of the value function V_{st}, we obtain

$$V_{\mathsf{arr}} = \sum_{(z,n)\in\mathbb{S}} \pi_{z,n}(\xi, \kappa)V_{\mathsf{st}}(z, n + 1).$$

Here the right-hand side is obtained by observing that after the agent arrives to the location in state (z, n), which happens with probability $\pi_{z,n}(\xi, \kappa)$, the number of agents at that location becomes $n + 1$, and the agent's continuation payoff is then $V_{\mathsf{st}}(z, n + 1)$.

Our second condition on equilibrium requires that the total expected payoff V_{arr} to an agent arriving at location k equals the total expected payoff an agent at the location receives on switching V_{sw}. Intuitively, we expect this condition to hold in any symmetric equilibrium of a system with a large but finite number of homogeneous locations, in which agents choose whether to stay in their current location or switch to a different location (chosen uniformly at random). In such a model, the switching decisions of the agents will force the switching payoffs of all populated locations to have the same value. Because our model of a single location does not endogenously capture these considerations, we impose this explicitly. In particular, we require that the switching payoff satisfies the following equation:

$$V_{\mathsf{sw}} = \sum_{(z,n)\in\mathbb{S}} \pi_{z,n}(\xi, \kappa)V_{\mathsf{st}}(z, n + 1). \tag{5}$$

The final condition we impose on the equilibrium is a requirement on arrival rate κ. Again, intuitively, in a symmetric equilibrium of a large finite model with homogeneous locations, we expect the expected number of agents at each location to be the same given by the agent density $\beta > 0$. To capture this in our model, we require that for a given agent density β, the arrival rate κ satisfies the following condition:

$$\sum_{(z,n)\in\mathbb{S}} n\, \pi_{z,n}(\xi, \kappa) = \beta. \tag{6}$$

Given these three conditions, we are now ready to define an MFE:

Definition 1 (Mean Field Equilibrium). An MFE consists of a strategy ξ, an arrival rate κ, and a switching payoff V_{sw} satisfying (3), (5), and (6).

Note that in comparison with a PBE, an MFE adopts a fairly natural and vastly simpler model of agent behavior. In a PBE of a finite model, an agent's strategy depends on the state of the agent's current location, history, and belief about the state of all other locations. Moreover, the agent constantly updates this belief based on the agent's observations of the arrival process at the agent's current location. For example, if an agent sees a high volume of arrivals at the agent's current location, the agent's updated belief would attribute lower resource levels at other locations, thereby lowering the agent's expected payoff for switching. Such complex considerations do not arise in an MFE, in which the payoff from switching is assumed to be fixed and independent of the state dynamics of the current location. In a large market, this assumption is reasonable because the fluctuations in the empirical distribution of the states of other locations are expected to cancel each other, analogous to a law of large numbers result.[1] In the next section, we show existence of an MFE.

3. Existence of a Mean Field Equilibrium

Here we state the main result of this paper, proving the existence of an MFE for general resource-sharing functions. Subsequently, in Section 4, we analyze the structure and properties of an MFE under specific assumptions on the resource-sharing function. We have the following main theorem.

Theorem 1. *For any $\lambda > 0$, $\beta > 0$, and $\{\mu_{z,y} > 0 : z, y \in \mathbb{Z}\}$, there exists an MFE $(\xi, \kappa, V_{\mathsf{sw}})$, where $\xi(z, n) = 0$ for all $z \in \mathbb{Z}$ and all large enough n.*

The underlying argument behind the proof is to carefully construct a correspondence $\mathcal{R}$ and show that the existence of an MFE is equivalent to the existence of a fixed point of $\mathcal{R}$. The latter is obtained by an application of the Fan–Glicksberg fixed-point theorem (Aliprantis and Border 2006). Here we first sketch the steps involved and highlight the technical challenges in each of those steps. Using these intermediate results, we then provide the proof of Theorem 1. (The complete proof is provided in Appendices B–G.)

1. We first show that for any Markovian strategy ξ and arrival rate $\kappa > 0$, the Markov chain $\mathsf{MC}(\xi, \kappa)$ has a unique invariant distribution π satisfying (4). This involves showing that the chain $\mathsf{MC}(\xi, \kappa)$ is irreducible and positive recurrent, which we accomplish by using coupling arguments to bound the chain between two $M/M/\infty$ queues. The proof of this result is provided in Appendix B.

Denote the (unique) invariant distribution of $\mathsf{MC}(\xi, \kappa)$ by $\pi(\xi, \kappa)$. In Appendix C, by applying Berge's maximum theorem (Berge 1963), we show that the invariant distribution $\pi(\xi, \kappa)$ is jointly continuous in (ξ, κ).

2. Second, we establish that for any strategy ξ, there exists a unique value of $\kappa > 0$ such that the invariant distribution $\pi(\xi, \kappa)$ satisfies (6). This result is achieved by showing that the quantity $\sum_{(z,n)\in\mathbb{S}} n\pi(z, n)$, where $\pi = \pi(\xi, \kappa)$ is strictly increasing and continuous for $\kappa \in [\beta\lambda(1 - \gamma), \beta\lambda]$ for any fixed ξ and using the intermediate value theorem. The proof of this result is provided in Appendix D.

Let $\kappa(\xi)$ denote the unique value of arrival rate κ for which $\pi(\xi, \kappa)$ satisfies (6). The first two steps together then define an injective map from the strategy ξ to an arrival rate $\kappa(\xi)$ and a steady-state distribution $\pi(\xi, \kappa(\xi))$ such that $\pi(\xi, \kappa(\xi))$ is the (unique) invariant distribution of the Markov chain $\mathsf{MC}(\xi, \kappa(\xi))$ and satisfies (6).

3. Third, we consider the decision problem $\mathsf{DEC}(\xi, \kappa(\xi), V_{\mathsf{sw}})$ for a given strategy ξ and switching payoff V_{sw}. We let $\mathcal{V}(\xi, V_{\mathsf{sw}})$ denote the value function satisfying the corresponding Bellman equation (2) and let $\mathcal{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}})$ denote the corresponding continuation payoff function. Finally, we let $\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})$ denote the right-hand side of (5):

$$\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}}) = \sum_{(z,n)\in\mathbb{S}} \pi_{z,n} V_{\mathsf{st}}(z, n + 1),$$

where $\pi = \pi(\xi, \kappa(\xi))$ and $V_{\mathsf{st}} = \mathcal{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}})$.

In Appendix E, we show that these functions are uniformly bounded. In particular, we show that there exists $0 < \underline{\mathsf{V}} \le \overline{\mathsf{V}}$ such that for all Markovian strategy ξ and $V_{\mathsf{sw}} > 0$, we have the switching payoff $\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}}) \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$. The proof of the uniform bounds makes extensive use of the strong Markov property for the chain $\mathsf{MC}(\xi, \kappa(\xi))$.

4. Fourth, we let $\mathcal{X}(\xi, V_{\mathsf{sw}})$ denote the set of all optimal strategies for the agent's decision problem $\mathsf{DEC}(\xi, \kappa(\xi), V_{\mathsf{sw}})$. Note that $\mathcal{X}(\xi, V_{\mathsf{sw}}) = \mathsf{OPT}(\xi, \kappa(\xi), V_{\mathsf{sw}})$. In Appendix F, we identify a convex, compact set $\widehat{\Pi}$ of Markovian strategies such that if $\xi \in \widehat{\Pi}$ and $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$, then $\mathcal{X}(\xi, V_{\mathsf{sw}}) \subseteq \widehat{\Pi}$. Let $\Upsilon = \widehat{\Pi} \times [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$.

Finally, we construct the correspondence $\mathcal{R} : \Upsilon \rightrightarrows \Upsilon$ defined as

$$\begin{aligned}\mathcal{R}(\xi, V_{\mathsf{sw}}) &= \mathcal{X}(\xi, V_{\mathsf{sw}}) \times \{\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})\} \\ &= \{(\zeta, \mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})) : \zeta \in \mathcal{X}(\xi, V_{\mathsf{sw}})\}.\end{aligned}$$

We depict the map pictorially in Figure 1. In Appendix G, we show that the correspondence $\mathcal{R}$ is upper hemicontinuous. This requires showing the continuity of the value functions in (ξ, V_{sw}), which is achieved using the continuity in ξ of the process $\mathsf{MC}(\xi, \kappa(\xi))$ under the topology of weak convergence (Ethier and Kurtz 1986).

We then obtain the following proof for the existence of an MFE.

Proof of Theorem 1. The steps outlined show that $\mathcal{R}$ is an upper-hemicontinuous correspondence on a convex, compact subset Υ of a metric space with values that are nonempty and convex. From an application of the Fan–Glicksberg fixed-point theorem (Aliprantis and Border 2006), we obtain that $\mathcal{R}$ has a fixed point; that is, there exists $(\xi, V_{\mathsf{sw}}) \in \Upsilon$ such that $(\xi, V_{\mathsf{sw}}) \in \mathcal{R}(\xi, V_{\mathsf{sw}})$.

Thus, by definition of $\mathcal{R}$, we have $\xi \in \mathcal{X}(\xi, V_{\mathsf{sw}})$. This implies that ξ satisfies (3) for the decision problem $\mathsf{DEC}(\xi, \kappa(\xi), V_{\mathsf{sw}})$. Second, by definition of $\kappa(\xi)$, we obtain that the steady-state distribution $\pi(\xi, \kappa(\xi))$ satisfies (6). Finally, from $V_{\mathsf{sw}} = \mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})$, we obtain that (5) holds. From this, we conclude that $(\xi, \kappa(\xi), V_{\mathsf{sw}})$ constitutes an MFE. □

Figure 1. Illustration of the Correspondence $\mathcal{R}(\xi, V_{\mathsf{sw}}) = X(\xi, V_{\mathsf{sw}}) \times \{V_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})\}$

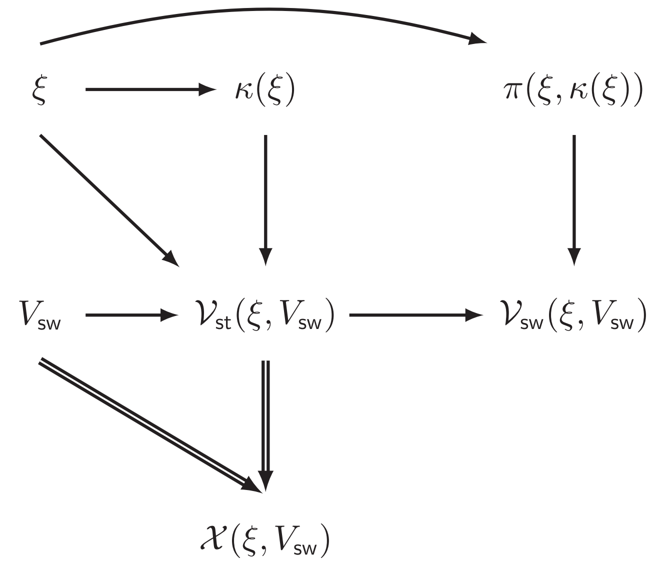

Note. Single arrows denote functions, and double arrows denote correspondences.

4. Equilibrium Analysis for Decreasing Resource-Sharing Functions

Having shown the existence of an MFE for general resource-sharing functions, we now characterize the equilibrium strategy for the specific case in which the resource-sharing function is nonincreasing in the number of agents at the location. Under this assumption, we show existence of an MFE in which the equilibrium strategies have a threshold structure. We then use this structural characterization in Section 5 to compute this MFE and analyze its welfare. We define decreasing resource-sharing functions as follows.

Definition 2. We say that a resource-sharing function F is *decreasing* if $F(z, n+1) \le F(z, n)$ for each $z \in \mathbb{Z}$ and all $n \in \mathbb{N}$.

Decreasing resource-sharing functions appear when agents' interactions are competitive rather than cooperative. In Section 5, we consider these three examples of decreasing resource-sharing functions:

- As a first example of a decreasing resource-sharing function, consider $F(z, n) = f(z)/n$ for some function f. This models settings in which all agents at a location equally share the resource there. In particular, given resource level z at a location, the n agents at the location would collectively obtain total payoffs at rate $\lambda nF(z, n) = \lambda f(z)$, a quantity independent of n. We refer to the quantity $W(z, n) \triangleq \lambda nF(z, n)$ as a *single-location welfare function.*
- Next, consider $F(z, n) = f(z)/\sqrt{n}$. Here the agents collectively receive payoffs at rate $\lambda\sqrt{n}f(z)$, which is increasing in n. Although agents compete with each other, the single-location welfare function increases with the number of agents there.
- Finally, consider $F(z, n) = f(z)/n^{3/2}$. This models extremely competitive settings in which the single-location welfare function decreases with the number of agents.

Before providing our result, we define threshold strategies. Formally, for $\mathbf{x} = (x_z : z \in \mathbb{Z})$, where $x_z \in \mathbb{R}_+$ for each $z \in \mathbb{Z}$, define the threshold strategy $\xi^{\mathbf{x}}$ as follows:

$$\xi^{\mathbf{x}}(z, n) = \begin{cases} 1 & \text{if } n < \lfloor x_z \rfloor; \\ x_z - \lfloor x_z \rfloor & \text{if } n = \lfloor x_z \rfloor; \\ 0 & \text{otherwise,} \end{cases}$$

for each $z \in \mathbb{Z}$ and $n \ge 0$. In particular, under strategy $\xi^{\mathbf{x}}$, an agent, at the agent's decision epoch, will stay at the agent's current location with resource level $z \in \mathbb{Z}$ if the number of agents n at the location is strictly below $\lfloor x_z \rfloor$, will switch to a different location if $n > \lfloor x_z \rfloor$, and will stay with probability $x_z - \lfloor x_z \rfloor$ and switch with remaining probability if $n = \lfloor x_z \rfloor$. We say that a strategy is a *threshold strategy* if it is of this form.

We now state our main result of this section.

Theorem 2. *If F is a decreasing resource-sharing function, there exists an MFE $(\xi, \kappa, V_{\mathsf{sw}})$, where ξ is a threshold strategy.*

The proof of this theorem makes essential use of the following lemma, which states that with decreasing resource-sharing functions, the continuation values are nonincreasing.

Lemma 1. *Let ξ be a Markovian strategy, $\kappa > 0$, and $V_{\mathsf{sw}} > 0$. If F is a decreasing resource-sharing function, then, for each $z \in \mathbb{Z}$, the continuation payoff $V_{\mathsf{st}}(z, n)$ for the decision problem $\mathsf{DEC}(\xi, \kappa, V_{\mathsf{sw}})$ is nonincreasing in n.*

The proof of this lemma, provided in Appendix H, shows that the decision problem $\mathsf{DEC}(\xi, \kappa, V_{\mathsf{sw}})$ has a dynamic program that satisfies closed convex cone properties defined in Smith and McCardle (2002). With the lemma in place, the proof of Theorem 2 follows from minor modifications of the argument in the proof of Theorem 1 and is omitted.

5. Computation of MFE and Numerical Equilibrium Analysis

The implications of Theorem 2 are of substantial practical importance: when the resource-sharing function is decreasing, the equilibrium behavior of the agents can be fully described by $|\mathbb{Z}|$ nonnegative real numbers $\{x_z : z \in \mathbb{Z}\}$. This parsimony allows simple computational methods to numerically identify an equilibrium, especially when $|\mathbb{Z}|$ is small. We use this fact to analyze the equilibrium numerically for several representative decreasing resource-sharing functions. We first describe our approach for computing an equilibrium in more detail.

5.1. Computation of MFE

To simplify notation in this section, we use $\mathbf{x}$ to denote the threshold strategy $\xi^{\mathbf{x}}$. Recall that an MFE is a fixed point of the correspondence $\mathcal{R}(\mathbf{x}, V_{\mathsf{sw}}) = \mathcal{X}(\mathbf{x}, V_{\mathsf{sw}}) \times \mathcal{V}_{\mathsf{sw}}(\mathbf{x}, V_{\mathsf{sw}})$. For any $(\mathbf{x}, V_{\mathsf{sw}})$, we define the distance metric $\mathsf{dist}_{\mathcal{R}}$ as follows:

$$\mathsf{dist}_{\mathcal{R}}(\mathbf{x}, V_{\mathsf{sw}}) = |V_{\mathsf{sw}} - \mathcal{V}_{\mathsf{sw}}(\mathbf{x}, V_{\mathsf{sw}})| + \inf_{\mathbf{y} \in \mathcal{X}(\mathbf{x}, V_{\mathsf{sw}})} \|\mathbf{x} - \mathbf{y}\|_2,$$

where $\|\cdot\|_2$ denotes the Euclidean norm. The second term on the right-hand side denotes the distance between $\mathbf{x}$ and the set $\mathcal{X}(\mathbf{x}, V_{\mathsf{sw}})$, which is compact and convex. To find a fixed point of $\mathcal{R}$, we identify a value of $(\mathbf{x}, V_{\mathsf{sw}})$ such that $\mathsf{dist}_{\mathcal{R}}(\mathbf{x}, V_{\mathsf{sw}}) = 0$. We implement two relaxations to this exact problem. First, we consider an approximation $\mathsf{dist}^{\epsilon}_{\mathcal{R}}$ to the metric $\mathsf{dist}_{\mathcal{R}}$, obtained primarily by truncating the state space to a finite set. Second, we perform an adaptive search method to find a (approximate) minimizer of the function $\mathsf{dist}^{\epsilon}_{\mathcal{R}}$. We choose this approximate minimizer as the value of the (approximate) MFE strategy and the corresponding switching payoff. We describe the steps in detail:

1. We truncate the state space $\mathbb{S}$ of the agent's decision problem to $\mathbb{S}_L = \mathbb{Z} \times \{0, 1, \cdots, L-1\}$ for some $L \in \mathbb{N}$. For each $\mathbf{x} \in [0, L-1]^{|\mathbb{Z}|}$, we let $\mathsf{MC}_L(\mathbf{x}, \kappa)$ denote the Markov chain obtained by restricting the transitions of the chain $\mathsf{MC}(\mathbf{x}, \kappa)$ to lie in the set $\mathbb{S}_L$ and let $\pi_L(\mathbf{x}, \kappa)$ denote its steady-state distribution. For any $\mathbf{x} \in [0, L]^{|\mathbb{Z}|}$, the distribution $\pi_L(\mathbf{x}, \kappa)$ can be obtained by solving a set of $L \cdot |\mathbb{Z}|$ linear equations analogous to (4).
2. For any given $\mathbf{x} \in [0, L-1]^{|\mathbb{Z}|}$, we perform a binary search over the interval $[\beta\lambda(1-\gamma), \beta\lambda]$ to find a value $\kappa = \kappa_L(\mathbf{x})$ for which

$$\left| \sum_{z \in \mathbb{Z}} \sum_{n=0}^{L} n\pi_{z,n} - \beta \right| \le \epsilon_1,$$

where $\pi = \pi_L(\mathbf{x}, \kappa_L(\mathbf{x}))$, and $\epsilon_1 > 0$ denotes the tolerance level within which we seek to satisfy (6).
3. For any given $\mathbf{x} \in [0, L-1]^{|\mathbb{Z}|}$ and $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$, we then consider the decision problem $\mathsf{DEC}(\mathbf{x}, \kappa_L(\mathbf{x}), V_{\mathsf{sw}})$ (with state space restricted to $\mathbb{S}_L$). We perform value iteration to compute approximate value functions $\mathcal{V}^{\epsilon}_{\mathsf{st}}(\mathbf{x}, V_{\mathsf{sw}})$ and $\mathcal{V}^{\epsilon}_{\mathsf{sw}}(\mathbf{x}, V_{\mathsf{sw}})$, where we iterate until $\mathcal{V}^{\epsilon}_{\mathsf{st}}(\mathbf{x}, V_{\mathsf{sw}})$ is within $\epsilon_0 > 0$ (in sup-norm) of the limit. Using these approximate value functions, we identify the set of approximately optimal thresholds $\mathcal{X}^{\epsilon}(\mathbf{x}, V_{\mathsf{sw}})$. Define $\mathsf{dist}^{\epsilon}_{\mathcal{R}}$ by replacing $\mathcal{V}_{\mathsf{sw}}$ and $\mathcal{X}$ in the definition of $\mathsf{dist}_{\mathcal{R}}$ with $\mathcal{V}^{\epsilon}_{\mathsf{sw}}$ and $\mathcal{X}^{\epsilon}$.
4. We seek to minimize $\mathsf{dist}^{\epsilon}_{\mathcal{R}}(\mathbf{x}, V_{\mathsf{sw}})$ over all values of $\mathbf{x} \in [0, L-1]^{|\mathbb{Z}|}$ and $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$. We use the Nelder–Mead neighborhood search method (Nelder and Mead 1965) to find the minimizer of the distance function. To locate the global minimum, we run the method in parallel with multiple initial values of $\mathbf{x}$ and V_{sw}, chosen among a discretized set of threshold strategies $\Pi^k_L = \{0, (L-1)/k, 2(L-1)/k, \cdots, (k-1)(L-1)/k, L-1\}^{|\mathbb{Z}|}$ for some $k \in \mathbb{N}$ and a discretized subset of $[\underline{\mathsf{V}}, \overline{\mathsf{V}}]$ constructed in a similar way.
5. After obtaining $(\mathbf{x}^*, V^*_{\mathsf{sw}})$ that attains the minimum of $\mathsf{dist}^{\epsilon}_{\mathcal{R}}$ over all runs, we do a validation check by comparing $\mathsf{dist}^{\epsilon}_{\mathcal{R}}(\mathbf{x}^*, V^*_{\mathsf{sw}})$ with a threshold ϵ_2 to see if this distance is close enough to zero for $(\mathbf{x}^*, V^*_{\mathsf{sw}})$ to be an equilibrium. We accept $(\mathbf{x}^*, V^*_{\mathsf{sw}})$ as an approximate MFE strategy and the corresponding switching payoff if $\mathsf{dist}^{\epsilon}_{\mathcal{R}}(\mathbf{x}^*, V^*_{\mathsf{sw}}) \le \epsilon_2$. If the validation check fails, a larger k is chosen to provide more fine-grained initial starting points until a maximum number of iterations is reached. Although our method does not guarantee finding an approximate equilibrium on terminating, in all our computations in Section 5.2, we obtain an approximate equilibrium with corresponding $\mathsf{dist}^{\epsilon}_{\mathcal{R}}$ smaller than 10^{-10}.

We also note that there may be multiple equilibria in our model for general model parameters and resource-sharing functions; we have not shown uniqueness. Such instances of nonuniqueness may arise, for example, when the resource-sharing function is multimodal because, in those settings, coordination concerns dominate, and an agent may prefer to stay at a location if other agents do so and prefer to switch if others switch. In such instances, the preceding numerical procedure selects for a particular (approximate) equilibrium, and our comparative statics results in the following section correspond to the equilibrium[2] selected by this algorithm.

5.2. Comparative Statics

In this section, we present the results of our numerical investigations of the agents' behavior in an MFE using the computational approach described in the preceding section. We study the setting in which $\mathbb{Z} = \{0, 1\}$ with transition rates $\mu_{0,1} = \mu_{1,0} = \mu$. Because our model is invariant to proportional scaling of the transition rate μ and the agents' interepoch rate λ, we fix $\lambda = 1$. We set the survival probability to $\gamma = 0.95$. We consider decreasing resource-sharing functions of the form $F(z,n) = zn^{-\alpha}$, where $\alpha \in \{0.5, 1, 1.5\}$. In this setting, some locations have resource (those with $z = 1$), and others do not ($z = 0$), and the single-location welfare function is increasing for $\alpha = 0.5$, constant for $\alpha = 1$, and decreasing for $\alpha = 1.5$ in the number of agents there. Finally, our approximation scheme uses parameters $L = 200$, $k = 20$, $\epsilon_0 = 10^{-4}$, $\epsilon_1 = 10^{-6}$, and $\epsilon_2 = 10^{-8}$.

In our computational study, we study how the model's parameters influence both agent behavior as quantified by the equilibrium thresholds and system efficiency as quantified by the welfare per location. The *welfare per location* is defined as the rate of total expected payoff obtained in equilibrium by all the agents at a location in steady state. At a location with resource level $z \in \mathbb{Z}$ and n agents, the total payoff rate to those n agents is given by $W(z,n) = \lambda n F(z,n)$. Because, in steady state, the state (z,n) is distributed according to the mean field distribution π, the agents' welfare per location equals

$$W_L = \mathbf{E}_\pi[W(Z,N)] = \sum_{z,n} \lambda n F(z,n) \pi_{z,n}.$$

We also analyze the *welfare per agent*, defined as the rate at which a randomly chosen agent receives payoff in equilibrium. Because the agent density is equal to β, the welfare per agent W_A is given by $W_A = W_L/\beta$. When β is held fixed, the two welfare measures are proportional, and thus we study W_A in addition to W_L only when we vary β.

Figure 2 shows how the equilibrium thresholds and the welfare per location vary as the resource process changes more frequently, that is, as μ increases, for a fixed value of $\beta = 20$. For each resource-sharing function, for small values of μ, the difference between the thresholds x_1 and x_0 is substantial. Because the resource level changes slowly, an agent in a location with resources is willing to suffer significant competition (in the form of other agents)

Figure 2. Equilibrium Thresholds and Welfare Under Different Resource Transition Rates μ with Agent Density Fixed at $\beta = 20$

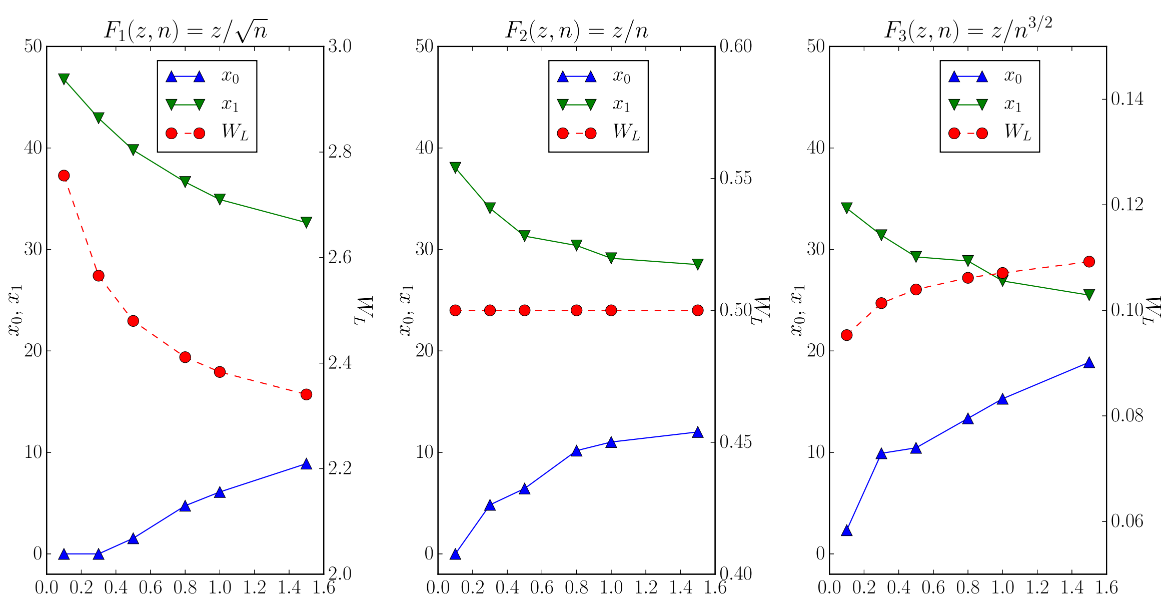

before choosing to switch locations. Note that as α increases, the level of competition at which agents switch decreases, consistent with our observation that as α increases, competition becomes more severe. By contrast, as μ increases, the difference in the two thresholds diminishes. This is because increasing μ diminishes the benefit of staying in a location. As the resource levels change more frequently, the resource process mixes more readily, and thus, future resource levels are less correlated with current levels.

Figure 2 also shows that the welfare per location depends crucially on the resource-sharing function. When the single-location welfare function increases with the number of agents at that location ($\alpha = 1/2$), the welfare per location decreases as resource levels change more frequently, that is, as μ increases. By contrast, when the single-location welfare function decreases with the number of agents there ($\alpha = 3/2$), the welfare per location increases as μ increases. To understand this, observe that when μ is small, the thresholds x_1 and x_0 are well separated, implying that the agents will be concentrated in locations with positive resource levels. By contrast, when μ is large, the two thresholds are similar, and agents are more equitably distributed between locations with and without resource. When $\alpha < 1$, the former distribution of the agents obtains more welfare per location because single-location welfare function is increasing with the number of agents at a location with resource, and having more agents at these locations increases welfare. By contrast, when $\alpha > 1$, the former distribution incurs lower welfare per location because of severe competition among the agents at the location with resource. (When $\alpha = 1$, the distribution of the agents between locations with or without resource does not substantially affect the welfare per location. In particular, as long as a location with resource has at least one agent present, the total payoff at that location is the same.)

Figure 3 shows equilibrium properties as a function of agent density β when resource levels change slowly ($\mu = 0.25$). The difference between the thresholds x_1 and x_0 widens as β increases for each resource-sharing function. This is because increasing β for any fixed state (z, n) at the current location diminishes an agent's expected payoff from switching because there are more agents to compete against. Thus, when the current location has resources, the agents become more likely to stay as β gets larger.

We further observe that as β increases, the welfare per location increases when $\alpha = 0.5$, decreases when $\alpha = 1.5$, and is essentially constant when $\alpha = 1$. As in Figure 2, this relation is explained by the equilibrium distribution of agents between locations with and without resource, arising from the dependence of the equilibrium thresholds on β: as the difference between the two thresholds increases, the welfare per location increases when $\alpha = 0.5$ and decreases when $\alpha = 1.5$. However, because the degree of competition increases as β increases, we observe that, irrespective of the resource-sharing function, the welfare per agent decreases.

The preceding comparative statics reveals an important feature of our dynamic model and its equilibrium that is lacking in a static analysis: our analysis captures the joint distribution of the agents and the resource levels across locations. Figure 2 demonstrates this by showing that agents' strategies change as the resource transition rate μ

Figure 3. Equilibrium Thresholds and Welfare Under Different Agent Densities β

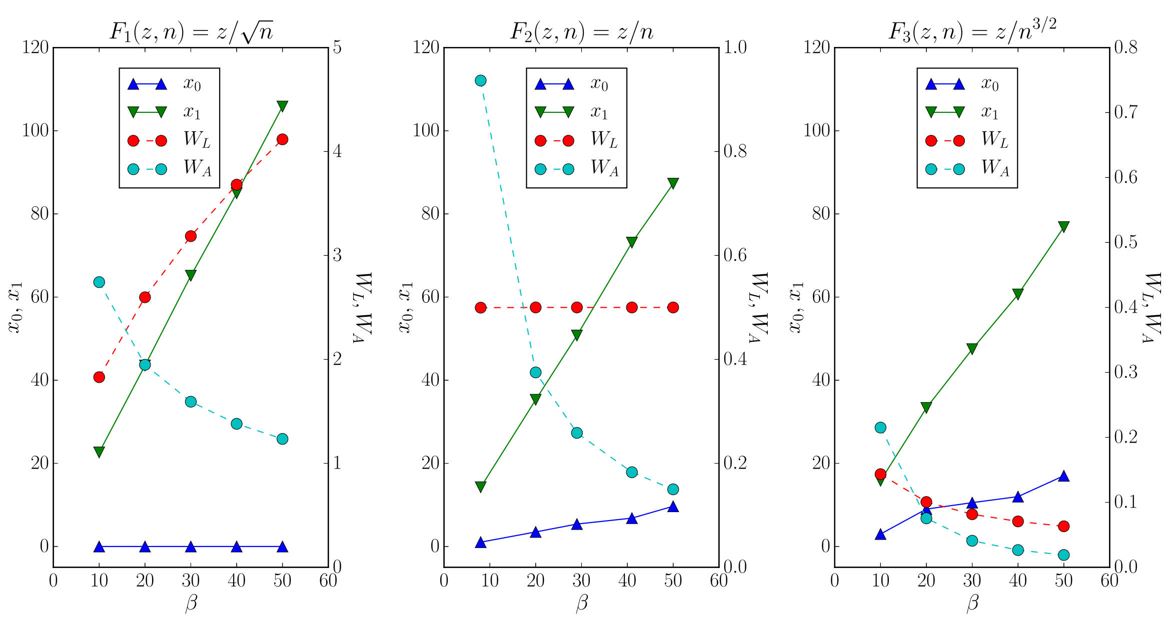

Notes. W_A is multiplied by 15 for all values. Note that the resource transition rate is given by $\mu = 0.25$.

changes. By contrast, because all values of μ result in the same steady-state proportion (50%) of locations in each resource state, a static analysis that only tracks the stationary resource state distribution would generate the same market outcomes for all values of μ. Furthermore, the welfare also changes with μ for resource-sharing functions other than z/n, where the total payoff rate in a location $\lambda nF(z,n)$ depends nontrivially on n. Such an effect would not materialize in a static model, which ignores the dynamics of the resource process and tracks only the steady state.

5.3. Case Study: Setting Platform Commission

In this section, we provide a case study to illustrate how our model can be used to evaluate engineering interventions. Specifically, we apply our model to the ride-hailing market in Manhattan. Ride-hailing platforms charge a commission when they transfer rider payments to their driver partners, and consequently, the drivers' behavior in the market is influenced by this commission rate. In this case study, we investigate how different commission rates affect the aggregate revenue of the drivers and the platform (and how it is split between the two); the outcome of this analysis provides a reference for platforms when an adjustment of commission rate is under consideration.

We view taxi drivers as agents, different neighborhoods of Manhattan as locations, and taxi trip demand as the resource in our model. We assume the drivers, at the end of each day, decide for the next day whether to stay in the same neighborhood or switch to another one. We also assume that a driver makes this decision based on the trip demand in the driver's current neighborhood as well as the driver's estimate of the number of competing drivers in the same neighborhood.

We describe how the model parameters are estimated and further describe the assumptions. We use the yellow cab trip records from the New York City Taxi and Limousine Commission (2017) data set to estimate these parameters. The data limitations prevent us from performing a full-blown analysis; in such instances, we use our judgment to assign parameter values. We set the parameter values as follows:

- Agent density β: We divide Manhattan into 12 regions, with the diameter of each region approximately equal to the average taxi trip length in Manhattan. The agent density is then estimated as $\beta = 400$ drivers per location, following an estimate of 4,800 active taxi drivers obtained by averaging across different times of day.
- Resource process $\{\mu_{z,y}\}$: We assume a resource model with binary states, with zero denoting the typical resource state and one denoting a high-resource state. Such a high resource may describe local conditions (such as local events, weather patterns, etc.) that temporarily lead to high demand for rides. To estimate the transition rates between the two states, we use weather as a proxy and estimate the transition between rainy and nonrainy days using historical weather data from Manhattan (Weather Underground 2017). This yields a transition rate of $\mu_{0,1} = 1/3.86$ and $\mu_{1,0} = 1/1.93$ with units day^{-1}. These values are a reasonable proxy for state transitions, indicating a high-resource state approximately every four days for a duration of about two consecutive days.
- Payment function F: Most ride-hailing platforms use dynamic pricing mechanisms to improve market efficiency, and such mechanisms can be designed to increase the aggregate revenue with the number of drivers (Chen 2016, Castillo et al. 2017) because increased driver availability allows more trips to happen. However, at the same time, higher competition among the drivers decreases the revenue received by an individual driver. To model these aspects, we let the aggregate revenue rate from riders at a location with resource state z and n drivers equal $n^{1-\alpha}f(z)$ for some parameter $\alpha \in (0,1)$, where $f(z)$ captures the dependence on the resource state. This entails the revenue rate per driver to equal $f(z)/n^{\alpha}$, and hence, the rate of payment to an individual driver in the location takes the following form:

$$F(z,n) = \frac{(1-c(z))f(z)}{n^{\alpha}}, \quad z \in \{0,1\}, n \in \mathbb{N},$$

where $c(z)$ denotes the (resource-dependent) commission rate charged by the platform. For our analysis, we choose $\alpha = 0.5$.

To estimate $f(0)$, we use the average daily rider payment on nonrainy days in Manhattan from the New York City Taxi and Limousine Commission (2017), which yields an estimate of 1.2×10^4 dollars per hour per location. We do not, however, estimate $f(1)$ using rider payments on rainy days because our data come from yellow cab data with fixed prices, whereas modern ride-hailing platforms typically increase price as demand increases. We therefore assume the average total rider payment when the resource is high ($z = 1$) to be 20% higher and set $f(1) = 1.2f(0)$.

- Decision rate: We choose $\lambda = 1\text{day}^{-1}$.
- Survival probability: We choose $\gamma = 0.995$, indicating a planning horizon of $1/\lambda(1-\gamma) = 200$ days.

Assuming a baseline commission rate of 15% in both resource states, we investigate how the revenue of drivers and the platform would vary under a number of commission-rate scenarios. For each such combination of $c(0)$ and $c(1)$ (and under the parameter values described), we numerically compute the resulting MFE in our model and the

Table 1. Revenue of the Ride-Sharing Platform Under Different Commission Rates

$c(0)$	$c(1)$	DriRev	ΔDriRev	PlatRev	ΔPlatRev	AggRev	ΔAggRev
0.15	0.15	26.121	—	4.610	—	30.731	—
0.175	0.175	25.353	−2.94%	5.378	16.66%	30.731	0.00%
0.15	0.20	25.504	−2.36%	5.219	13.20%	30.723	−0.02%
0.20	0.15	25.210	−3.49%	5.507	19.46%	30.718	−0.04%
0.2	0.2	24.584	−5.88%	6.146	33.32%	30.730	0.00%

Notes. Here $c(z)$ is the commission rate at resource state $z \in \{0, 1\}$. DriRev, PlatRev, and AggRev denote the revenue (in units 10^5 dollars per hour) for drivers, for the platform, and in aggregate, respectively. ΔDriRev denotes the change in drivers' revenue compared with the base case ($c(0) = c(1) = 0.15$) with ΔPlatRev and ΔAggRev defined similarly.

driver and platform revenues in the computed equilibrium. We share these results in Table 1. These results can be used to access the magnitude of the impact and to decide whether commission should be raised in aggregate or if it would be better to selectively raise it based on demand (resource states). A table such as this could be shared with decision makers as part of a larger decision process.

As discussed earlier, our dynamic model allows us to capture the joint distribution of the drivers and the aggregate revenue across locations. The distribution of the drivers across locations is important because it influences a driver's payoff on switching, which influences the driver's switching decisions. Our model enables us to include this endogenous effect of the driver distribution on the drivers' switching decisions in evaluating different commission rates. Without the dynamics (and the tractable equilibrium concept of an MFE), such effects would be hard to incorporate in a static analysis, rendering it incomplete.

6. Conclusion

Our results establish that, in equilibrium, the agents in our model base their decision to explore solely on the state of the location they currently reside in and on its steady-state distribution. In particular, our results justify analyzing spatiotemporal models under simple yet optimal models of agent behavior.

Our model and analysis raise many topics for future research. First, we have used the notion of an MFE to analyze a single location in isolation, assuming that the other locations are described by the mean field distribution. A natural question is whether the resulting strategy constitutes an approximate equilibrium in the system with a large but finite number of agents and locations. Such approximation results for MFE have been obtained in other contexts (see, e.g., Iyer et al. 2014, Adlakha et al. 2015, Balseiro et al. 2015). In the finite system, a single agent visits multiple locations over the agent's lifetime, inducing correlations among the states of those locations. The analytical challenge in obtaining an approximation result involves showing that as the system size increases, such correlations vanish, and in the limit, the dynamics of a location in the finite system approaches the dynamics of the single location in our model.

On the modeling front, we have assumed that each location is homogeneous. In particular, we assume that the resource process is distributed independently and identically across different locations. One consequence of this homogeneity is that agents do not choose their destination when they switch. It is straightforward to extend our model and the analysis to incorporate location heterogeneity and to let agents choose their destination when switching. Such a model would better represent the settings we study. For example, in ride-hailing settings, residential neighborhoods have different demand characteristics than business districts, and drivers choose the neighborhood to operate in based on these characteristics. A formal finite model of this extension has multiple *types* of locations, and each location has type-dependent resource dynamics, resource-sharing function, and agent density. Agents choose not only whether to stay or switch but also which location type to switch to, whereupon the destination is chosen uniformly among locations of that type. Using similar arguments for the homogeneous setting, we can obtain the corresponding limiting infinite system meant to capture the limiting behavior as the number of agents and locations increases with the type distribution and the agent density fixed. Our proof of the existence of an MFE applies in this setting with minor modifications; we omit the details because of space considerations. By contrast, we have assumed that the resource process at a location is exogenously specified, whereas an extension could allow for the resource transitions at a location to depend on the number of agents therein.

Finally, our work also sets the stage for analyzing engineering interventions and their economic impact. One such intervention involves altering the resource-sharing function at each location through subsidies or penalties to induce the agents to stay in or switch from a location, thereby affecting their welfare. A further question is whether sharing information about locations' states would benefit or harm the agents and how such an information-sharing mechanism should be designed. Answers to these questions would help platforms, such as Uber or Airbnb, to increase their efficiency.

Appendix A. Description of the Finite System

In this appendix, we provide a formal description of the system with a finite number of locations and agents (the *finite system*), which motivates our mean field model. The finite system has a set of locations $\mathcal{K}$, where each location $k \in \mathcal{K}$ contains a stochastic time-varying resource. We use Z_t^k to denote the resource level at location k at time $t \geq 0$. We assume that the resource process $\{Z_t^k : t \geq 0\}$ is a finite-state, continuous-time Markov chain and further assume that the resource processes across different locations in the system are distributed identically and independently. We let $\mathbb{Z}$ denote the set of values the resource process can take and let $\mu_{zy} > 0$ denote the transition rate of Z_t^k from a state $z \in \mathbb{Z}$ to a state $y \in \mathbb{Z}$. Furthermore, we make the assumption that each process Z_t^k is irreducible and positive recurrent with a unique invariant distribution given by $\{\pi_{\mathsf{res}}(z) : z \in \mathbb{Z}\}$.

Spread across this set of locations are N agents. Each agent may switch between locations in search of resources and less competition as we detail. Each agent i is associated with a Poisson clock with rate λ such that each time the clock rings, the agent decides whether to stay in the location or switch to another one. We refer to each clock ring of agent i as the agent's *decision epoch* and let τ_i^ℓ and k_i^ℓ denote the time and location of the agent's ℓ^{th} decision epoch, respectively.

We let N_t^k denote the number of agents at location k at time t. At each decision epoch $t = \tau_i^\ell$, agent i at location $k = k_i^\ell$ receives a payoff $F(Z_t^k, N_t^k)$ that depends on the resource level Z_t^k and the number of agents N_t^k at that location. We make the same assumptions on F as in Section 2.1.

Subsequent to receiving the payoff, agent i makes the decision whether to continue at the location or move to a different location. On choosing to move to a different location, agent i instantaneously arrives at a new location $k_i^{\ell+1}$. We make the assumption that the new location $k_i^{\ell+1}$ is drawn independently and uniformly from the set of all locations other than the agent's current location. Note that this assumption precludes us from modeling an agent's strategic choice of to *which* location to move. Nevertheless, we make this assumption because, even under this restrictive assumption, the analysis of the agent's decision problem turns out to be challenging. In Section 6, we discussed a few extensions and modifications that align closer to practical settings.

Similar to the mean field model, we assume that agents in the finite system are short lived: after each decision epoch τ_i^ℓ, subsequent to making the decision regarding whether to stay in the agent's current location or move to a different location, agent i departs the system independently with probability $1-\gamma$, never to return, and we denote as τ_i the time the agent leaves the system. We also assume that for each agent that departs, a new agent arrives to the system at a location chosen uniformly at random to maintain constant system size, same as in the mean field model.

Finally, we describe the utility and the information structure of each agent in the model. We assume that each agent i at each time t at the agent's current location k observes the resource level Z_t^k and the number of agents N_t^k. By contrast, the agent cannot observe the resource level and the number of agents at any other location. We assume that the agents have perfect recall, and hence, at any decision epoch τ_i^ℓ, agent i bases the decision to stay or move on the entire history (namely, the resource levels and the number of agents at each location the agent has visited) the agent has observed until that time.

Given this informational assumption, each agent i is risk neutral and wants to maximize the total expected payoff accrued over the agent's lifetime. Formally, each agent i seeks to maximize

$$\mathbf{E}\left[\sum_{\ell=1}^{\infty} F(Z_i^\ell, N_i^\ell)\mathbf{I}\{\tau_i^\ell \leq \tau_i\}\right],$$

where the expectation is over the randomness in the resource levels, the arrival and departure process of the agents, and their (and their competitors') strategies. Because the departure of an agent is independent of the rest of the system, it is straightforward to show that the agent's expected payoff can be equivalently written as

$$\mathbf{E}\left[\sum_{\ell=1}^{\infty} \gamma^{\ell-1} F(Z_i^\ell, N_i^\ell)\right].$$

Thus, each agent i's decision problem is equivalent to the decision problem faced by a persistent agent (who never departs the system) seeking to maximizer the agent's total expected discounted payoff.

Appendix B. Existence and Uniqueness of Invariant Distribution of $\mathsf{MC}(\xi, \kappa)$

In this appendix, we show that for any Markovian strategy ξ and arrival rate $\kappa > 0$, the Markov chain $\mathsf{MC}(\xi, \kappa)$ has a unique steady-state distribution.

Lemma B.1. *For any Markovian strategy ξ and arrival rate $\kappa \geq 0$, there exists a unique steady-state distribution for* $\mathsf{MC}(\xi, \kappa)$ *satisfying* (4).

Proof of Lemma B.1. Fix a Markovian strategy ξ and an arrival rate $\kappa > 0$. We prove the lemma by showing that the Markov chain $\mathsf{MC}(\xi, \kappa)$ is irreducible and positive recurrent. The fact that $\mathsf{MC}(\xi, \kappa)$ is irreducible follows straightforwardly from (1) and the fact that the resource process is independent and ergodic. Thus, it only remains to show that the chain is positive recurrent.

Let $\mathbb{S}_0 \triangleq \{(z, 0) : z \in \mathbb{Z}\}$, and define $T_z(\mathbb{S}_0)$ as the first return time of the chain to $\mathbb{S}_0$ given that it starts at $(z, 0)$:

$$T_z(\mathbb{S}_0) \triangleq \inf\{t > 0 : (Z_t, N_t) \in \mathbb{S}_0, (Z_s, N_s) \notin \mathbb{S}_0 \text{ for some } 0 < s < t, \\ \text{given } (Z_0, N_0) = (z, 0)\}.$$

In the following, we show that $T_z(\mathbb{S}_0)$ has a finite expectation for each $z \in \mathbb{Z}$. From this, using the ergodicity of the resource process, it follows that the return time to a particular state $(z_0, 0) \in \mathbb{S}_0$ also has a finite expectation, and hence the chain is ergodic.

To show that $T_z(\mathbb{S}_0)$ has finite expectation, we use a coupling argument. Given a Markov chain $(Z_t, N_t) \sim \mathsf{MC}(\xi, \kappa)$ with $(Z_0, N_0) = (z, 0)$, we construct a coupled process $N_t^{(1)} \sim \mathsf{M/M/\infty}((1-\gamma)\lambda, \kappa)$ with $N_t^{(1)} = 0$ as in the proof of Lemma I.2. Define $\tilde{T}_z(0)$ to be the first return time to zero of the chain $N_t^{(1)}$. From the construction of the coupling, it follows that $N_t \leq N_t^{(1)}$ for all $t \geq 0$, and hence, $T_z(\mathbb{S}_0) \leq \tilde{T}_z(0)$. Thus, we have $\mathbf{E}[T_z(\mathbb{S}_0)] \leq \mathbf{E}[\tilde{T}_z(0)]$. The result then follows immediately from the fact that an $\mathsf{M/M/\infty}((1-\gamma)\lambda, \kappa)$ queue is ergodic, and hence, $\mathbf{E}[\tilde{T}_z(0)] < \infty$ for all $z \in \mathbb{Z}$. □

Appendix C. Joint Continuity of the Invariant Distribution of $\mathsf{MC}(\xi, \kappa)$

In the following, we show that the steady-state distribution $\pi(\xi, \kappa)$ of the Markov chain $\mathsf{MC}(\xi, \kappa)$ is jointly (and uniformly) continuous in its parameters. This continuity result plays an important role in subsequent results that constitute our proof of existence of an MFE.

To prove the continuity of $\pi(\xi, \kappa)$, we adopt an approach similar to Le Van and Stachurski (2007), in which we characterize the invariant distribution of $\mathsf{MC}(\xi, \kappa)$ as a maximizer of a continuous function and apply Berge's maximum theorem. Before we present the formal argument, we specify the topologies (and the metric) we impose on the set of Markovian strategies and the set of invariant probability distributions and specify the continuous function Λ that we consider. First, we endow the state space $\mathbb{S} = \mathbb{Z} \times \mathbb{N}_0$ with the discrete topology. Let $\mathscr{C}_b(\mathbb{S})$ denote the set of bounded function $h : \mathbb{S} \to \mathbb{R}$. (Note that because we impose the discrete topology on $\mathbb{S}$, any such h is also continuous.) We endow $\mathscr{C}_b(\mathbb{S})$ with the sup-norm:

$$\|h_1 - h_2\|_\infty \triangleq \sup_{(z,n)\in\mathbb{S}} |h_1(z,n) - h_2(z,n)|, \quad \text{for } h_1, h_2 \in \mathscr{C}_b(\mathbb{S}). \tag{C.1}$$

Let $\Pi \subseteq \mathscr{C}_b(\mathbb{S})$ denote the set of Markovian strategies with the topology induced from $\mathscr{C}_b(\mathbb{S})$.

We let $\mathscr{M}(\mathbb{S})$ denote the set of finite signed measures on $\mathbb{S}$, and we endow $\mathscr{M}(\mathbb{S})$ with the weak topology, which is equivalent to the topology induced by the ℓ_1-norm because $\mathbb{S}$ is countable:

$$\|\mu - \nu\|_1 = \sum_{(z,n)\in\mathbb{S}} |\mu(z,n) - \nu(z,n)|, \quad \text{for } \mu, \nu \in \mathscr{M}(\mathbb{S}). \tag{C.2}$$

Let $\Gamma = \{\pi(\xi, \kappa) : \xi \in \Pi, \kappa \in [\beta\lambda(1-\gamma), \beta\lambda]\} \subseteq \mathscr{M}(\mathbb{S})$ denote the set of invariant distributions (with the induced topology) for all Markovian strategies and arrival rates. Let $\overline{\Gamma}$ denote the closure of Γ.

For $\xi \in \Pi$, $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda]$, and $\nu \in \overline{\Gamma}$, define $\Lambda(\xi, \kappa, \nu)$ as follows:

$$\Lambda(\xi, \kappa, \nu) \triangleq -\sum_{(z,n)\in\mathbb{S}} \frac{1}{n+1} |(\nu\mathsf{Q})(z,n)|, \tag{C.3}$$

where $\mathsf{Q} = \mathsf{Q}^{\xi,\kappa}$ denotes the transition kernel of $\mathsf{MC}(\xi, \kappa)$, and $\nu\mathsf{Q}$ is defined as

$$(\nu\mathsf{Q})(z,n) = \sum_{(y,m)\in\mathbb{S}} \nu(y,m)\mathsf{Q}((y,m) \to (z,n)).$$

With the preliminaries in place, we are now ready to state the main lemma of this section.

Lemma C.1. *The map $(\xi, \kappa) \mapsto \pi(\xi, \kappa)$ is jointly (and uniformly) continuous in (ξ, κ) for $\xi \in \Pi$ and $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda]$.*

Proof of Lemma C.1. In Lemma C.2, we show that the set of distributions Γ is uniformly tight. Then, from Prohorov's theorem (Billingsley 2013), we obtain that $\overline{\Gamma}$ is compact. Observe that

$$\arg\max_{\nu\in\overline{\Gamma}} \Lambda(\xi, \kappa, \nu) = \{\pi(\xi, \kappa)\}. \tag{C.4}$$

This follows from the fact that $\pi(\xi, \kappa)$ is the unique probability distribution over $\mathbb{S}$ for which (4) holds.

In Lemma C.3, we show that $\Lambda(\xi, \kappa, \nu)$ is jointly (and uniformly) continuous in its parameters for $\xi \in \Pi$, $\nu \in \overline{\Gamma}$, and $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda]$. The result then follows from a direct application of Berge's maximum theorem (Berge 1963) to (C.4).

The following two auxiliary lemmas are used in the proof of Lemma C.1. □

Lemma C.2. *The set Γ of invariant distributions is tight.*

Proof of Lemma C.2. We prove this lemma using a coupling argument. For any $\xi \in \Pi$ and $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda]$, let $(Z_t, N_t) \sim \mathsf{MC}(\xi, \kappa)$ with $(Z_0, N_0) = (z, n)$ for some $(z, n) \in \mathbb{S}$. Let π denote the invariant distribution of $\mathsf{MC}(\xi, \kappa)$. Independently, let $\widehat{N}_t \sim \mathsf{M/M/\infty}(\lambda(1-\gamma), \beta\lambda)$ with $\widehat{N}_0 = n$. Let $\widehat{\pi}$ denote the invariant distribution of $\mathsf{M/M/\infty}(\lambda(1-\gamma), \beta\lambda)$; it is straightforward to show that $\widehat{\pi}$ is Poisson with mean $\beta/(1-\gamma)$.

Using Lemmas I.1 and I.2, we obtain that N_t is (first-order) stochastically dominated by $\widehat{N}_t$ for all $t \geq 0$. From this, we obtain (by taking limits and using ergodicity) that for all $k > 0$, we have

$$\sum_{z\in\mathbb{Z}} \sum_{n>k} \pi(z,n) \leq \sum_{n>k} \widehat{\pi}(n).$$

For any $\epsilon > 0$, choose a $k^\epsilon > 0$ such that $\sum_{n>k^\epsilon} \widehat{\pi}(n) < \epsilon$. (Such a k^ϵ exists given that $\widehat{\pi}$ is Poisson with finite mean.) This implies that

$$\sum_{z\in\mathbb{Z}} \sum_{n>k^\epsilon} \pi(z,n) < \epsilon, \quad \text{for all } \epsilon > 0.$$

Because k^ϵ is independent of the choice of (ξ, κ), we obtain that Γ is tight.

The following lemma proves the joint continuity of Λ. □

Lemma C.3. *The function Λ as defined in (C.3) is jointly (and uniformly) continuous.*

Proof of Lemma C.3. Consider $\xi_i \in \Pi$, $\kappa_i \in [\beta\lambda(1-\gamma), \beta\lambda]$, and $\nu_i \in \overline{\Gamma}$ for $i = 1, 2$. We let Q_i denote the transition kernel of $\mathsf{MC}(\xi_i, \kappa_i)$. We have

$$\begin{aligned} &|\Lambda(\xi_1,\kappa_1,\nu_1) - \Lambda(\xi_2,\kappa_2,\nu_2)| \\ &\leq |\Lambda(\xi_1,\kappa_1,\nu_1) - \Lambda(\xi_2,\kappa_2,\nu_1)| + |\Lambda(\xi_2,\kappa_2,\nu_1) - \Lambda(\xi_2,\kappa_2,\nu_2)|. \end{aligned} \tag{C.5}$$

Now note that

$$\begin{aligned} &|\Lambda(\xi_1,\kappa_1,\nu_1) - \Lambda(\xi_2,\kappa_2,\nu_1)| \\ &\leq \sum_{(z,n)\in\mathbb{S}} \frac{1}{n+1} |(\nu_1\mathsf{Q}_1)(z,n) - (\nu_1\mathsf{Q}_2)(z,n)| \\ &\leq \sum_{(z,n)\in\mathbb{S}} \sum_{(y,m)\in\mathbb{S}} \frac{1}{n+1} \nu_1(y,m) \big|\mathsf{Q}_1((y,m)\to(z,n)) - \mathsf{Q}_2((y,m)\to(z,n))\big| \\ &\leq \sum_{(y,m)\in\mathbb{S}} \nu_1(y,m) \sum_{(z,n)\in\mathbb{S}} \frac{1}{n+1} \big|\mathsf{Q}_1((y,m)\to(z,n)) - \mathsf{Q}_2((y,m)\to(z,n))\big| \\ &\leq \sup_{(y,m)\in\mathbb{S}} \sum_{(z,n)\in\mathbb{S}} \frac{1}{n+1} \big|\mathsf{Q}_1((y,m)\to(z,n)) - \mathsf{Q}_2((y,m)\to(z,n))\big|. \end{aligned} \tag{C.6}$$

Now, using (1), we obtain that

$$\begin{aligned} &\mathsf{Q}_1((y,m)\to(z,n)) - \mathsf{Q}_2((y,m)\to(z,n)) \\ &= \mathbf{I}\{z=y, n=m+1\}(\kappa_1-\kappa_2) + \mathbf{I}\{z=y, n=m-1\}\lambda\gamma m(\xi_1(y,m) - \xi_2(y,m)) \\ &\quad - \mathbf{I}\{z=y, n=m\}\big(\kappa_1-\kappa_2 + \lambda m\gamma(\xi_1(y,m)-\xi_2(y,m))\big), \end{aligned}$$

and hence

$$\begin{aligned} &\sum_{(z,n)\in\mathbb{S}} \frac{1}{n+1} \big|\mathsf{Q}_1((y,m)\to(z,n)) - \mathsf{Q}_2((y,m)\to(z,n))\big| \\ &\leq \left(\frac{1}{m+2} + \frac{1}{m+1}\right)|\kappa_1-\kappa_2| + \lambda\gamma\left(1 + \frac{m}{m+1}\right)|\xi_1(y,m) - \xi_2(y,m)| \\ &\leq 2\big(|\kappa_1-\kappa_2| + \lambda\gamma|\xi_1(y,m)-\xi_2(y,m)|\big). \end{aligned}$$

Thus, from (C.6), we obtain

$$|\Lambda(\xi_1,\kappa_1,\nu_1) - \Lambda(\xi_2,\kappa_2,\nu_1)| \leq 2|\kappa_1-\kappa_2| + 2\lambda\gamma\|\xi_1-\xi_2\|_\infty. \tag{C.7}$$

Next, observe that

$$\begin{aligned} &|\Lambda(\xi_2,\kappa_2,\nu_1) - \Lambda(\xi_2,\kappa_2,\nu_2)| \\ &\leq \sum_{(z,n)\in\mathbb{S}} \frac{1}{n+1} |(\nu_1\mathsf{Q}_2)(z,n) - (\nu_2\mathsf{Q}_2)(z,n)| \\ &\leq \sum_{(z,n)\in\mathbb{S}} \frac{1}{n+1} \sum_{(y,m)\in\mathbb{S}} |\mathsf{Q}_2((y,m)\to(z,n))||\nu_1(y,m) - \nu_2(y,m)| \\ &\leq \sum_{(y,m)\in\mathbb{S}} |\nu_1(y,m) - \nu_2(y,m)| \sum_{(z,n)\in\mathbb{S}} \frac{1}{n+1} |\mathsf{Q}_2((y,m)\to(z,n))|. \end{aligned} \tag{C.8}$$

Now, again from (1) and after some straightforward algebra, we obtain that

$$\begin{aligned}
&\sum_{(z,n)\in\mathbb{S}} \frac{1}{n+1}\left|\mathsf{Q}_2((y,m)\to(z,n))\right| \\
&\le \frac{2}{m+1}\sum_{z\neq y}\mu_{y,z} + \kappa_2\left(\frac{1}{m+2}+\frac{1}{m+1}\right) + \lambda(1-\gamma\xi_2(y,m))\left(1+\frac{m}{m+1}\right) \\
&\le \sum_{z\neq y}\mu_{y,z} + 2\kappa_2 + 2\lambda \\
&\le \max_{y\in\mathbb{Z}}\sum_{z\neq y}\mu_{y,z} + 2(\beta+1)\lambda,
\end{aligned}$$

where we have used the fact that $\kappa_2 \le \beta\lambda$ in the last inequality. Thus, from (C.8), we obtain

$$|\Lambda(\xi_2,\kappa_2,\nu_1) - \Lambda(\xi_2,\kappa_2,\nu_2)| \le \left(\max_{y\in\mathbb{Z}}\sum_{z\neq y}\mu_{y,z} + 2(\beta+1)\lambda\right)\|\nu_1-\nu_2\|_1. \tag{C.9}$$

Therefore, combining (C.5), (C.7), and (C.9), we obtain

$$\begin{aligned}
&|\Lambda(\xi_1,\kappa_1,\nu_1) - \Lambda(\xi_2,\kappa_2,\nu_2)| \\
&\le 2|\kappa_1-\kappa_2| + 2\lambda\gamma\|\xi_1-\xi_2\|_\infty + \left(\max_{y\in\mathbb{Z}}\sum_{z\neq y}\mu_{y,z} + 2(\beta+1)\lambda\right)\|\nu_1-\nu_2\|_1.
\end{aligned}$$

Thus, Λ is Lipschitz and, hence, jointly and (uniformly) continuous in its parameters. □

Appendix D. Existence of κ Satisfying Equilibrium Conditions

In this appendix, we show for any Markovian strategy ξ that there exists a unique arrival rate $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda]$ for which the steady-state distribution π of the Markov chain $\mathsf{MC}(\xi,\kappa)$ satisfies (6).

Toward that goal, for any Markovian strategy ξ and arrival rate $\kappa > 0$, define

$$\phi(\xi,\kappa) \triangleq \sum_{(z,n)\in\mathbb{S}} n\pi^{\xi,\kappa}(z,n),$$

where $\pi^{\xi,\kappa}$ is the unique steady-state distribution of $\mathsf{MC}(\xi,\kappa)$. We seek to show that there exists a $\kappa \in [\beta\lambda(1-\gamma),\beta\lambda]$ such that $\phi(\xi,\kappa)=\beta$. We prove this result using the intermediate value theorem. First, we show that $\phi(\xi,\kappa)$ is a strictly increasing function of κ for any given $\xi \in \Pi$. Second, we show that $\phi(\xi,\beta\lambda(1-\beta)) \le \beta$ and $\phi(\xi,\beta\lambda) \ge \beta$, which imply that any κ such that $\phi(\xi,\kappa)=\beta$ must lie in $[\beta\lambda(1-\gamma),\beta\lambda]$. The result then follows once we show that $\phi(\xi,\kappa)$ is a continuous function of κ.

In the rest of this appendix, we assume that the strategy ξ is fixed and drop the explicit dependence on ξ from notation wherever convenient. We now proceed with the first step.

D.1. Strict Monotonicity of $\phi(\cdot)$

Lemma D.1. *Given any Markovian strategy ξ, $\phi(\kappa)$ is a strictly increasing function of κ on $[\beta\lambda(1-\gamma),\beta\lambda]$.*

Proof of Lemma D.1. For any $\kappa_1,\kappa_2 \in [\beta\lambda(1-\gamma),\beta\lambda]$ with $\kappa_1<\kappa_2$, consider two coupled chains $(Z_t^{(i)},N_t^{(i)}) \sim \mathsf{MC}(\xi,\kappa_i)$ for $i=1,2$ as in the proof of Lemma I.2, in which $Z_t^{(1)}=Z_t^{(2)}$ and $N_t^{(1)}\le N_t^{(2)}$ for all $t\ge 0$.

For $i=1,2$, we have

$$\frac{1}{t}\int_0^t N_s^{(i)}ds \to \sum_{z,n} n\pi_i(z,n) = \phi(\kappa_i)$$

almost surely as $t\to\infty$, where we write π_i for π^{κ_i}. Because $N_t^{(1)}\le N_t^{(2)}$ for all t, we have $\phi(\kappa_1)\le\phi(\kappa_2)$.

Next, suppose for the sake of contradiction that $\phi(\kappa_1)=\phi(\kappa_2)$. Because $Z_t^{(1)}=Z_t^{(2)}$ and $N_t^{(1)}\le N_t^{(2)}$ for all $t\ge0$, we have

$$\mathbf{I}\{Z_t^{(1)}=z, N_t^{(1)}\ge n\} \le \mathbf{I}\{Z_t^{(2)}=z, N_t^{(2)}\ge n\}, \tag{D.1}$$

for all $(z,n)\in\mathbb{S}$, $t\ge 0$.

For any $(z,n)\in\mathbb{S}$, we have

$$\frac{1}{t}\int_0^t \mathbf{I}\{Z_s^{(i)}=z, N_s^{(i)}\ge n\}ds \to \sum_{n'\ge n}\pi_i(z,n') \tag{D.2}$$

almost surely as $t \to \infty$ for $i = 1, 2$. By (D.1) and (D.2), we have

$$\sum_{n' \geq n} \pi_1(z, n') \leq \sum_{n' \geq n} \pi_2(z, n'),$$

for any (z, n), and

$$\phi(\kappa_1) = \sum_{n \geq 0} n \left(\sum_{z \in \mathbb{Z}} \pi_1(z, n) \right) = \sum_{n \geq 0} \sum_{z \in \mathbb{Z}, n' > n} \pi_1(z, n') \leq \sum_{n \geq 0} \sum_{z \in \mathbb{Z}, n' > n} \pi_2(z, n') = \phi(\kappa_2).$$

Because, by our assumption $\phi(\kappa_1) = \phi(\kappa_2)$, the inequality in the preceding equation is actually an equality. This implies

$$\sum_{n' \geq n} \pi_1(z, n') = \sum_{n' \geq n} \pi_2(z, n'),$$

for all (z, n), which further implies that π_1 and π_2 are the same distribution.

For $i = 1, 2$, Equation (4) implies

$$\begin{aligned} &\pi_i(z, n) \left(\kappa_i + \sum_{y \neq z} \mu_{z,y} + \lambda n (1 - \gamma \xi(z, n)) \right) \\ &= \pi_i(z, n-1)\kappa_i + \sum_{y \neq z} \mu_{y,z} \pi_i(y, n) + \pi_i(z, n+1)\lambda(n+1)(1 - \gamma \xi(z, n+1)), \end{aligned}$$

which leads to

$$\begin{aligned} &\kappa_i(\pi_i(z, n) - \pi_i(z, n-1)) \\ &= \sum_{y \neq z} \mu_{y,z} \pi_i(y, n) + \pi_i(z, n+1)\lambda(n+1)(1 - \gamma \xi(z, n+1)) \\ &\quad - \pi_i(z, n) \left(\sum_{y \neq z} \mu_{z,y} + \lambda n (1 - \gamma \xi(z, n)) \right). \end{aligned} \tag{D.3}$$

Because $\pi_1 = \pi_2$, the right-hand side of (D.3) is the same for $i = 1, 2$; hence, we have

$$\kappa_1(\pi_1(z, n) - \pi_1(z, n-1)) = \kappa_2(\pi_2(z, n) - \pi_2(z, n-1)).$$

But $\kappa_1 < \kappa_2$ and $\pi_1 = \pi_2$ imply that, for $i = 1, 2$, $\pi_i(z, n) = \pi_i(z, n-1)$ for all (z, n); hence, π_i cannot be a probability distribution over $\mathbb{S}$, and this contradiction completes the proof. □

D.2. Bounds for $\phi(\cdot)$

In this section, we provide bounds on the function $\phi(\kappa)$ for any $\kappa > 0$. These bounds immediately imply that, for $\kappa = \beta\lambda$, $\phi(\kappa) \geq \beta$, and for $\kappa = \beta\lambda(1-\gamma)$, $\phi(\kappa) \leq \beta$. Together with Lemma D.1, this implies that any κ for which $\phi(\kappa) = \beta$ must lie in the interval $[\beta\lambda(1-\gamma), \beta\lambda]$.

Lemma D.2. *For any Markovian strategy ξ and arriving rate $\kappa \geq 0$, $\phi(\kappa)$ satisfies*

$$\frac{\kappa}{\lambda} \leq \phi(\kappa) \leq \frac{\kappa}{\lambda(1-\gamma)}.$$

Proof of Lemma D.2. Let $(Z_t, N_t) \sim \mathsf{MC}(\xi, \kappa)$, and $(Z_0, N_0) = (z, n)$. Denote as $\mathsf{M/M/}\infty\,(\lambda, \kappa)$ an (independent) $M/M/\infty$ queue with arrival rate κ and service rate λ, and let $N_t^{(i)} \sim \mathsf{M/M/}\infty\,(\lambda_i, \kappa)$ for $i = 1, 2$ be two independent processes with $N_0^{(1)} = N_0^{(2)} = n$, where $\lambda_1 = \lambda$ and $\lambda_2 = (1-\gamma)\lambda$.

Let π^κ be the steady-state distribution of $\mathsf{MC}(\xi, \kappa)$ and π_i be the steady-state distribution of $\mathsf{M/M/}\infty\,(\lambda_i, \kappa)$ for $i = 1, 2$. We have

$$\frac{1}{t} \int_0^t N_s ds \to \sum_{z,n} n \pi^\kappa(z, n),$$

and

$$\frac{1}{t} \int_0^t N_s^{(i)} ds \to \sum_n n \pi_i(n), \quad i = 1, 2$$

almost surely as $t \to \infty$. From Lemma I.2, we have $N_t^{(1)} \leqslant_{\mathsf{sd}} N_t \leqslant_{\mathsf{sd}} N_t^{(2)}$ for all $t \geq 0$; therefore, we have

$$\sum_n n \pi_1(n) \leq \sum_{z,n} n \pi^\kappa(z, n) \leq \sum_n n \pi_2(n).$$

The result then follows from the fact that, for $i = 1, 2$, π_i is a Poisson distribution with mean κ/λ_i. □

D.3. Continuity of $\phi(\cdot)$

Observe that the existence of a $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda]$ such that $\phi(\kappa) = \beta$ would follow immediately once we prove the continuity of $\phi(\cdot)$ in κ for any fixed Markovian strategy ξ. In this section, we prove a stronger statement, namely that $\phi(\xi, \kappa)$ is jointly continuous in (ξ, κ).

Lemma D.3. *The map $\phi(\xi, \kappa)$ is jointly and uniformly continuous in (ξ, κ) for Markovian ξ and for $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda]$.*

Proof of Lemma D.3. Given Markovian strategies ξ_1 and ξ_2 and arriving rates $\kappa_1, \kappa_2 \in [\beta\lambda(1-\gamma), \beta\lambda]$, let π_i be the steady-state distribution of $\mathsf{MC}(\xi_i, \kappa_i)$ for $i = 1, 2$. We have, for any arbitrary $k \geq 0$,

$$
\begin{aligned}
&\left|\phi(\xi_1, \kappa_1) - \phi(\xi_2, \kappa_2)\right| \\
&= \left|\sum_{z,n} n(\pi_1(z,n) - \pi_2(z,n))\right| \\
&\leq \left|\sum_{z\in\mathbb{Z}} \sum_{n\leq k} n(\pi_1(z,n) - \pi_2(z,n))\right| + \left|\sum_{z\in\mathbb{Z}} \sum_{n>k} n(\pi_1(z,n) - \pi_2(z,n))\right| \\
&\leq \left|\sum_{z\in\mathbb{Z}} \sum_{n\leq k} n(\pi_1(z,n) - \pi_2(z,n))\right| + \sum_{z\in\mathbb{Z}} \sum_{n>k} n\pi_1(z,n) + \sum_{z\in\mathbb{Z}, n>k} n\pi_2(z,n).
\end{aligned}
\tag{D.4}
$$

Now, bounding the first term, we obtain

$$
\begin{aligned}
&\left|\sum_{z\in\mathbb{Z}} \sum_{n\leq k} n(\pi_1(z,n) - \pi_2(z,n))\right| \\
&\leq k \sum_{z\in\mathbb{Z}} \sum_{n\leq k} |\pi_1(z,n) - \pi_2(z,n)| \leq k\|\pi_1 - \pi_2\|_1.
\end{aligned}
\tag{D.5}
$$

To bound the other terms, we use a coupling argument. Let $(Z_t^{(i)}, N_t^{(i)}) \sim \mathsf{MC}(\xi_i, \kappa_i)$ with $(Z_0^{(i)}, N_0^{(i)}) = (z, n)$ for $i = 1, 2$. Let $\widehat{N}_t \sim \mathsf{M/M/}\infty\,(\lambda(1-\gamma), \beta\lambda)$, with $\widehat{N}_0 = n$ denoting the number of agents in an (independent) $M/M/\infty$ queue at time with arrival rate $\beta\lambda$ and service rate $\lambda(1-\gamma)$. Let $\widehat{\pi}$ denote the steady-state distribution of $\widehat{N}_t$. By Lemmas I.1 and I.2, we have $N_t^{(i)} \preccurlyeq_{\mathsf{sd}} \widehat{N}_t$ for all $t \geq 0$ and for each $i = 1, 2$. From this stochastic dominance, it is straightforward to obtain that

$$
\sum_{z\in\mathbb{Z}, n>k} n\pi_i(z,n) \leq \sum_{n>k} n\widehat{\pi}(n), \quad i = 1, 2.
\tag{D.6}
$$

Thus, from (D.4) and (D.5), we have

$$
|\phi(\xi_1, \kappa_1) - \phi(\xi_2, \kappa_2)| \leq k\|\pi_1 - \pi_2\|_1 + 2\sum_{n>k} n\widehat{\pi}(n).
$$

Now, for any $\epsilon > 0$, choose k such that $\sum_{n>k} n\widehat{\pi}(n) < \epsilon/4$. (Note that this choice of k is independent of (ξ_i, κ_i) and depends only on the steady state $\widehat{\pi}$ of $\mathsf{M/M/}\infty(\lambda(1-\gamma), \beta\lambda)$, which is Poisson with mean $\beta/(1-\gamma)$.) Second, from Lemma C.1, we obtain that for any $\epsilon > 0$, there exists a $\delta > 0$ such that for all (ξ_1, κ_1) and (ξ_2, κ_2) such that $\|\xi_1 - \xi_2\|_\infty < \delta$ and $|\kappa_1 - \kappa_2| < \delta$, we have $\|\pi_1 - \pi_2\|_1 < \epsilon/2k$. Taken together, we obtain that for any $\epsilon > 0$, there exists a $\delta > 0$ such that for all (ξ_1, κ_1) and (ξ_2, κ_2) such that $\|\xi_1 - \xi_2\| < \delta$ and $|\kappa_1 - \kappa_2| < \delta$, we have $|\phi(\xi_1, \kappa_1) - \phi(\xi_2, \kappa_2)| < \epsilon$. Thus, we obtain that $\phi(\cdot)$ is jointly and uniformly continuous. □

D.4. Continuity

For any $\xi \in \Pi$, let $\kappa(\xi)$ denote the unique value of κ for which $\pi(\xi, \kappa)$ satisfies (6). We show that $\kappa(\xi)$ is a continuous function of ξ.

Lemma D.4. *The map $\xi \mapsto \kappa(\xi)$ is continuous.*

Proof of Lemma D.4. Define $W(\xi, \kappa) = -|\beta - \phi(\xi, \kappa)|$. Note that from Lemma D.1, we obtain

$$
\underset{\kappa\in[\beta\lambda(1-\gamma),\beta\lambda]}{\arg\max}\; W(\xi, \kappa) = \{\kappa(\xi)\}.
$$

From Lemma D.3, we obtain that $\phi(\xi, \kappa)$ is jointly continuous in (ξ, κ), and hence, so is $W(\xi, \kappa)$. The result then follows from Berge's maximum theorem (Berge 1963). □

Appendix E. Uniform Bounds on Value Functions

For a given $\xi \in \Pi$ and $V_{\mathsf{sw}} > 0$, we seek to study the decision problem $\mathsf{DEC}(\xi, \kappa(\xi), V_{\mathsf{sw}})$. Before we proceed, we need some definitions. Let $(Z_t, N_t) \sim \mathsf{MC}(\xi, \kappa(\xi))$, and let $\mathbf{E}^\xi(\cdot|z, n)$ denote the expectation operator with respect to $\{(Z_t, N_t) : t \geq 0\}$ conditioned on $(Z_0, N_0) = (z, n)$. Fix an agent, say agent 1, among all the agents at time 0, and let τ be the agent's first decision epoch.

Let $\mathbf{T} : \Pi \times \mathbb{R}_+ \times \mathscr{C}_b(\mathbb{S}) \to \mathscr{C}_b(\mathbb{S})$ denote the Bellman operator for the agent's decision problem $\mathsf{DEC}(\xi, \kappa(\xi), V_{\mathsf{sw}})$, where, for any $\xi \in \Pi$, $V_{\mathsf{sw}} > 0$, and $U \in \mathscr{C}_b(\mathbb{S})$, the function $W = \mathbf{T}(\xi, V_{\mathsf{sw}}, U)$ is defined as follows:

$$W(z,n) = F(z,n) + \gamma \max\left\{\mathbf{E}^{\xi}\left[U(Z_\tau, N_\tau)|z,n\right], V_{\mathsf{sw}}\right\}, \quad \text{for all } (z,n) \in \mathbb{S}. \tag{E.1}$$

The following lemma states that the map $\mathbf{T}(\xi, V_{\mathsf{sw}}, \cdot)$ is a contraction. The proof follows from standard arguments and is omitted.

Lemma E.1. *For any $\xi \in \Pi$ and $V_{\mathsf{sw}} > 0$, we have $\mathbf{T}(\xi, V_{\mathsf{sw}}, U) \in \mathscr{C}_b(\mathbb{S})$ for all $U \in \mathscr{C}_b(\mathbb{S})$. Furthermore, the map $\mathbf{T}(\xi, V_{\mathsf{sw}}, \cdot) : \mathscr{C}_b(\mathbb{S}) \to \mathscr{C}_b(\mathbb{S})$ is a contraction (with contraction parameter γ) for any $\xi \in \Pi$ and $V_{\mathsf{sw}} > 0$.*

Let $\mathscr{V}(\xi, V_{\mathsf{sw}}) \in \mathscr{C}_b(\mathbb{S})$ be the unique fixed point of $\mathbf{T}(\xi, V_{\mathsf{sw}}, \cdot)$. Define $\mathscr{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}}) \in \mathscr{C}_b(\mathbb{S})$ and $\mathscr{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}}) \in \mathbb{R}_+$ as follows:

$$\begin{aligned}
\mathscr{V}_{\mathsf{st}}(z,n;\xi,V_{\mathsf{sw}}) &= \mathbf{E}^{\xi}\left[\mathscr{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}})|z,n\right] \\
\mathscr{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}}) &= \sum_{(z,n)\in\mathbb{S}} \pi(z,n)\mathscr{V}_{\mathsf{st}}(z,n+1;\xi,V_{\mathsf{sw}}),
\end{aligned} \tag{E.2}$$

where $\pi = \pi(\xi, \kappa(\xi))$. Here $\mathscr{V}(z,n;\xi,V_{\mathsf{sw}})$ (and $\mathscr{V}_{\mathsf{st}}(z,n;\xi,V_{\mathsf{sw}})$) denotes the value taken by $\mathscr{V}(\xi, V_{\mathsf{sw}})$ (respectively, $\mathscr{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}})$) at $(z,n) \in \mathbb{S}$.

We begin this section by providing bounds on $\mathscr{V}_{\mathsf{sw}}$, $\mathscr{V}_{\mathsf{st}}$, and $\mathscr{V}$. Define

$$\begin{aligned}
\overline{\mathsf{V}} &= \frac{1}{1-\gamma}\|F\|_\infty, \\
\underline{\mathsf{V}} &= \exp\left(-\frac{\beta}{1-\gamma}\right) \sum_{(z,n)\in\mathbb{S}} \frac{\beta^n(1-\gamma)^n}{(1+\beta+\Psi)^{n+1}(n+1)!}\pi_{\mathsf{res}}(z)F(z,n+1) > 0,
\end{aligned}$$

where $\Psi = \frac{1}{\lambda}\max_{z\in\mathbb{Z}} \sum_{y\neq z} \mu_{zy} \in (0,\infty)$, and π_{res} is the steady-state distribution of the resource process.

The following lemma, providing a uniform upper bound on the value functions, follows immediately from the definition.

Lemma E.2. *For any $\xi \in \Pi$ and $V_{\mathsf{sw}} > 0$, the value functions satisfy $|\mathscr{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})| \leq \|\mathscr{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}})\|_\infty \leq \|\mathscr{V}(\xi, V_{\mathsf{sw}})\|_\infty \leq \overline{\mathsf{V}} = \|F\|_\infty/(1-\gamma)$.*

Proof of Lemma E.2. Observe that from (E.2), we have $|\mathscr{V}_{\mathsf{sw}}(\xi, \mathsf{sw})| \leq \|\mathscr{V}_{\mathsf{st}}(\xi, \mathsf{sw})\|_\infty \leq \|\mathscr{V}(\xi, \mathsf{sw})\|_\infty$. Also, from the fact that $\mathscr{V}(\xi, V_{\mathsf{sw}})$ is the fixed point of $\mathbf{T}(\xi, V_{\mathsf{sw}}, \cdot)$, we obtain

$$\|\mathscr{V}(\xi, V_{\mathsf{sw}})\|_\infty \leq \|F\|_\infty + \gamma \max\{\|\mathscr{V}(\xi, V_{\mathsf{sw}})\|_\infty, |\mathscr{V}_{\mathsf{sw}}(\xi, \mathsf{sw})|\} = \|F\|_\infty + \gamma\|\mathscr{V}(\xi, V_{\mathsf{sw}})\|_\infty.$$

Rearranging, we obtain that $\|\mathscr{V}(\xi, V_{\mathsf{sw}})\|_\infty \leq \|F\|_\infty/(1-\gamma) = \overline{\mathsf{V}}$.

The next lemma provides a uniform lower bound on the value functions. The proof makes extensive use of the strong Markovian property for the chain $\mathsf{MC}(\xi, \kappa(\xi))$. □

Lemma E.3. *For any $\xi \in \Pi$ and $V_{\mathsf{sw}} > 0$, we have $\mathscr{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}}) \geq \underline{\mathsf{V}}$.*

Proof of Lemma E.3. Observe that

$$\mathscr{V}(z,n;\xi,V_{\mathsf{sw}}) = F(z,n) + \gamma\max\{\mathscr{V}_{\mathsf{st}}(z,n;\xi,V_{\mathsf{sw}}), V_{\mathsf{sw}}\} \geq F(z,n).$$

Recalling the definition of $\mathscr{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}})$ and using the (strong) Markov property, we obtain

$$\begin{aligned}
\mathscr{V}_{\mathsf{st}}(z,n;\xi,V_{\mathsf{sw}}) = &\ \frac{\lambda}{n\lambda + \kappa(\xi) + \sum_{y\neq z}\mu_{zy}}\mathscr{V}(z,n;\xi,V_{\mathsf{sw}}) \\
&+ \frac{\kappa(\xi)}{n\lambda + \kappa(\xi) + \sum_{y\neq z}\mu_{zy}}\mathscr{V}_{\mathsf{st}}(z,n+1;\xi,V_{\mathsf{sw}}) \\
&+ \sum_{w\neq z}\frac{\mu_{wz}}{n\lambda + \kappa(\xi) + \sum_{y\neq z}\mu_{zy}}\mathscr{V}_{\mathsf{st}}(w,n;\xi,V_{\mathsf{sw}}) \\
&+ \frac{(n-1)\lambda(1-\gamma\xi(z,n))}{n\lambda + \kappa(\xi) + \sum_{y\neq z}\mu_{zy}}\mathscr{V}_{\mathsf{st}}(z,n-1;\xi,V_{\mathsf{sw}}) \\
&+ \frac{(n-1)\lambda\gamma\xi(z,n)}{n\lambda + \kappa(\xi) + \sum_{y\neq z}\mu_{zy}}\mathscr{V}_{\mathsf{st}}(z,n;\xi,V_{\mathsf{sw}}) \\
\geq &\ \frac{\lambda}{n\lambda + \kappa(\xi) + \sum_{y\neq z}\mu_{zy}}\mathscr{V}(z,n;\xi,V_{\mathsf{sw}}) \\
\geq &\ \frac{\lambda}{\lambda(n+\beta) + \sum_{y\neq z}\mu_{zy}}F(z,n),
\end{aligned}$$

where the last line follows from the fact that $\kappa(\xi) \le \beta\lambda$. Using the definition of Ψ, we obtain

$$\mathcal{V}_{\mathsf{st}}(z, n; \xi, V_{\mathsf{sw}}) \ge \frac{1}{n + \beta + \Psi} F(z, n) \ge \frac{1}{n(1 + \beta + \Psi)} F(z, n). \tag{E.3}$$

Next, observe that $\pi = \pi(\xi, \kappa(\xi))$ satisfies the steady-state (4):

$$\sum_{(y,m)\in\mathbb{S}} \pi(y, m)\mathsf{Q}^{\xi}((y, m) \to (z, n)) = 0,$$

where Q^{ξ} denotes the transition kernel of the Markov chain $\mathsf{MC}(\xi, \kappa(\xi))$. Using expression (1) for Q^{ξ}, we obtain

$$\begin{aligned}\pi(z, n)&\Big(\kappa(\xi) + \sum_{y\neq z} \mu_{zy} + \lambda n(1 - \gamma\xi(z, n))\Big)\\ &= \pi(z, n - 1)\kappa(\xi) + \sum_{w\neq z} \pi(w, n)\mu_{wz} + \pi(z, n + 1)\lambda(n + 1)(1 - \gamma\xi(z, n + 1)).\end{aligned}$$

This implies that

$$\begin{aligned}\pi(z, n) &\ge \pi(z, n - 1)\frac{\kappa(\xi)}{\kappa(\xi) + \sum_{y\neq z} \mu_{zy} + \lambda n(1 - \gamma\xi(z, n))}\\ &\ge \pi(z, n - 1)\frac{\beta\lambda(1 - \gamma)}{\beta\lambda(1 - \gamma) + \sum_{y\neq z} \mu_{zy} + \lambda n}\\ &\ge \pi(z, n - 1)\frac{\beta(1 - \gamma)}{\beta(1 - \gamma) + \Psi + n}\\ &\ge \pi(z, n - 1)\frac{\beta(1 - \gamma)}{(1 + \beta + \Psi)n}.\end{aligned}$$

Thus, we obtain

$$\pi(z, n) \ge \pi(z, 0)\frac{\beta^n(1 - \gamma)^n}{(1 + \beta + \Psi)^n n!}.$$

Now, from Lemma I.2, we obtain that the process $(Z_t, N_t) \sim \mathsf{MC}(\xi, \kappa(\xi))$ with $(Z_0, N_0) = (z, n)$ is stochastically dominated by $(Z_t, N_t^{(1)})$, where $N_t^{(1)}$ is an (independent) $\mathsf{M/M/\infty}$ $(\lambda(1 - \gamma), \beta\lambda)$ process with $N_0^{(1)} = n$. Hence, we have $\pi(z, 0) \ge \pi_{\mathsf{res}}(z)\mathbf{P}(N_\infty^{(1)} = 0)$, where the steady-state $N_\infty^{(1)}$ is given by a Poisson distribution with parameter $\beta/(1 - \gamma)$, implying that $\mathbf{P}(N_\infty^{(1)} = 0) = \exp(-\beta/(1 - \gamma))$. Thus, we obtain

$$\pi(z, n) \ge \pi_{\mathsf{res}}(z)\exp\left(-\frac{\beta}{1 - \gamma}\right)\frac{\beta^n(1 - \gamma)^n}{(1 + \beta + \Psi)^n n!}. \tag{E.4}$$

Finally, from (E.2), we have

$$\begin{aligned}\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}}) &= \sum_{(z,n)\in\mathbb{S}} \pi(z, n)\mathcal{V}_{\mathsf{st}}(z, n + 1; \xi, V_{\mathsf{sw}})\\ &\ge \sum_{(z,n)\in\mathbb{S}} \pi(z, n)\frac{F(z, n + 1)}{(1 + \beta + \Psi)(n + 1)}\\ &\ge \exp\left(-\frac{\beta}{1 - \gamma}\right)\sum_{(z,n)\in\mathbb{S}} \frac{\beta^n(1 - \gamma)^n}{(1 + \beta + \Psi)^{n+1}(n + 1)!}\pi_{\mathsf{res}}(z)F(z, n + 1)\\ &= \underline{\mathsf{V}},\end{aligned}$$

where we use (E.3) in the first inequality and (E.4) in the second. □

Appendix F. A Compact Set of Markovian Strategies

For $\xi \in \Pi$ and $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$, denote the set of all optimal Markovian strategies for the decision problem $\mathsf{DEC}(\xi, \kappa(\xi), V_{\mathsf{sw}})$ by $\mathcal{X}(\xi, V_{\mathsf{sw}}) \subseteq \Pi$. In particular, $\mathcal{X}(\xi, V_{\mathsf{sw}})$ is the set of all $\zeta \in \Pi$ such that

$$\zeta(z, n) = \begin{cases} 1 & \text{if } \mathcal{V}_{\mathsf{st}}(z, n; \xi, V_{\mathsf{sw}}) > V_{\mathsf{sw}};\\ 0 & \text{if } \mathcal{V}_{\mathsf{st}}(z, n; \xi, V_{\mathsf{sw}}) < V_{\mathsf{sw}}.\end{cases}$$

It is straightforward to show that the set $\mathcal{X}(\xi, V_{\mathsf{sw}})$ is nonempty and convex.

In this appendix, we provide characterization of a compact set $\widehat{\Pi} \subseteq \Pi$ of strategies such that if $\xi \in \widehat{\Pi}$ and $V_{\mathsf{sw}} \in [\underline{V}, \overline{V}]$, then $\mathscr{X}(\xi, V_{\mathsf{sw}}) \subseteq \widehat{\Pi}$. This characterization is later used to define a correspondence over a compact set to which we apply the Kakutani fixed-point theorem to show the existence of an MFE. (Note that the set Π is not compact under the sup-norm.)

We begin by defining the set $\widehat{\Pi}$. Recall that $F(z,n) \to 0$ as $n \to \infty$ for all $z \in \mathbb{Z}$. Let K_0 be defined as

$$K_0 = \inf\left\{m : F(z,n) < \frac{(1-\gamma)^2}{2}\underline{V} \text{ for all } z \in \mathbb{Z} \text{ and } n \geq m\right\},$$

and let K_1 be defined as

$$K_1 = \inf\left\{n : \exp\left(-\frac{1}{8}\sqrt{n-1}\right) + \frac{2}{\sqrt{\log(n-1)}} + \gamma^{\lfloor\sqrt{\log(n-1)}\rfloor}(1-\gamma) < \frac{(1-\gamma)^2\underline{V}}{4\|F\|_\infty}\right\}$$

Let $K_{\max} = \max\{4K_0^2 + 1, K_1\}$. Define the set $\widehat{\Pi} \in \Pi$ as follows:

$$\widehat{\Pi} = \{\xi \in \Pi : \xi(z,n) = 0 \text{ for all } z \in \mathbb{Z} \text{ and } n \geq K_{\max}\}.$$

In other words, under any strategy $\xi \in \widehat{\Pi}$, each agent chooses to switch locations if the number of agents at the agent's location is greater than $K_{\max}$ irrespective of the resource level. It is straightforward to show that $\widehat{\Pi}$ is compact by noting that it is isomorphic to $[0,1]^{K_{\max}}$ under the Euclidean topology.

The following lemma states that if $\xi \in \widehat{\Pi}$ and $V_{\mathsf{sw}} \geq \underline{V}$, then the optimal action for an agent at the state (z,n) is to switch if $n \geq K_{\max}$.

Lemma F.1. *For $\xi \in \widehat{\Pi}$ and $V_{\mathsf{sw}} \geq \underline{V}$, we have $\mathscr{V}_{\mathsf{st}}(z,n;\xi,V_{\mathsf{sw}}) < V_{\mathsf{sw}}$ for all $z \in \mathbb{Z}$ and for all $n \geq K_{\max}$.*

Proof of Lemma F.1. Consider an agent i in location k facing the decision problem $\mathsf{DEC}(\xi, \kappa(\xi), V_{\mathsf{sw}})$ for a given $\xi \in \widehat{\Pi}$ and $V_{\mathsf{sw}} \geq \underline{V}$. Let $\tau^\ell > 0$ denote the time of the ℓ^{th}-decision epoch of the agent for $\ell = 1, 2, \cdots$. Let (Z_t, N_t) denote the state of the location at time t, and for brevity, we let (Z_ℓ, N_ℓ) denote $(Z_{\tau^\ell}, N_{\tau^\ell})$ for each $\ell = 1, 2, \cdots$.

Suppose that $(Z_0, N_0) = (z,n)$ for $z \in \mathbb{Z}$ and $n \geq K_{\max}$. Fix a strategy $\phi \in \Pi$ for the agent, and let τ^ϕ denote the first time at which the agent chooses to switch under ϕ. Let $V_{\mathsf{st}}^\phi(z,n)$ denote agent i's continuation payoffs under the strategy ϕ subsequent to the agent making the decision to stay and not leave the system given the state of the location (z,n). We have the following expression for $V_{\mathsf{st}}^\phi(z,n)$:

$$V_{\mathsf{st}}^\phi(z,n) = \mathbf{E}\left[\sum_{\ell=1}^\infty \gamma^{\ell-1} F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\} + \gamma^\ell \mathbf{I}\{\tau^\ell = \tau^\phi\} V_{\mathsf{sw}}\right]. \tag{F.1}$$

The first term inside the expectation denotes the total expected payoff until the agent chooses to switch; the second term denotes the payoff on switching. Here the expectation $\mathbf{E}$ is conditioned on $(Z_0, N_0) = (z,n)$ and on the fact that agent i follows strategy ϕ and all other agents follow strategy ξ. (We drop this explicit dependence from the notation for $\mathbf{E}$ for brevity.) From this, we obtain

$$\begin{aligned} V_{\mathsf{st}}^\phi(z,n) &= \sum_{\ell=1}^\infty \gamma^{\ell-1}\mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\}\right] + \sum_{\ell=1}^\infty \gamma^\ell V_{\mathsf{sw}} \mathbf{P}(\tau^\ell = \tau^\phi) \\ &\leq \sum_{\ell=1}^\infty \gamma^{\ell-1}\mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\}\right] + \gamma V_{\mathsf{sw}}. \end{aligned} \tag{F.2}$$

Let $\widehat{n} = \lfloor\sqrt{n-1}/2 + 1\rfloor$. For each $\ell = 1, 2, \cdots$, we have

$$\begin{aligned} \mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\}\right] &= \mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\} \middle| N_\ell \geq \widehat{n}\right]\mathbf{P}(N_\ell \geq \widehat{n}) \\ &\quad + \mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\} \middle| N_\ell < \widehat{n}\right]\mathbf{P}(N_\ell < \widehat{n}) \\ &\leq \frac{1}{2}(1-\gamma)^2\underline{V} + \|F\|_\infty \mathbf{P}(N_\ell < \widehat{n}). \end{aligned} \tag{F.3}$$

Here, in the inequality, the first term follows from the fact that because $n > K_{\max}$, we have $\widehat{n} > K_0$, and hence, $F(Z_\ell, N_\ell) < (1-\gamma)^2\underline{V}/2$, if $N_\ell \geq \widehat{n}$. In the second term, we have used the fact that $F(Z_\ell, N_\ell) \leq \|F\|_\infty$.

To bound $\mathbf{P}(N_\ell < \widehat{n})$, consider an auxiliary system with n agents at $t = 0$, where each agent other than agent i stays in the system for a time that is independently and identically distributed as an exponential distribution with rate λ. (We assume that agent i never leaves the auxiliary system.) Furthermore, there are no arrivals to this auxiliary system. Let $\tilde{N}_t$ denote the number of agents in this auxiliary system. It is straightforward to show that $\tilde{N}_t$ is first-order stochastically dominated by N_t via a coupling argument, and we omit the details here. This implies that $\mathbf{P}(N_\ell \leq \widehat{n}) \leq \mathbf{P}(\tilde{N}_\ell \leq \widehat{n})$, where we write $\tilde{N}_\ell$ to denote $\tilde{N}_{\tau^\ell}$. Thus, we obtain

$$\mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\}\right] \leq \frac{1}{2}(1-\gamma)^2\underline{V} + \|F\|_\infty \mathbf{P}(\tilde{N}_\ell < \widehat{n}).$$

Let $\widehat{\ell} = \lfloor\sqrt{\log(n-1)}\rfloor$ and $\widehat{t} = \log(n-1)/2\lambda$. For each $\ell \le \widehat{\ell}$, we have $\tilde{N}_\ell \ge \tilde{N}_{\widehat{\ell}}$, and hence,

$$\begin{aligned}\mathbf{P}(\tilde{N}_\ell < \widehat{n}) &\le \mathbf{P}(\tilde{N}_{\widehat{\ell}} < \widehat{n})\\ &= \mathbf{P}(\tilde{N}_{\widehat{\ell}} < \widehat{n} | \tau^{\widehat{\ell}} < \widehat{t})\mathbf{P}(\tau^{\widehat{\ell}} < \widehat{t}) + \mathbf{P}(\tilde{N}_{\widehat{\ell}} < \widehat{n} | \tau^{\widehat{\ell}} \ge \widehat{t})\mathbf{P}(\tau^{\widehat{\ell}} \ge \widehat{t})\\ &\le \mathbf{P}(\tilde{N}_{\widehat{\ell}} < \widehat{n} | \tau^{\widehat{\ell}} < \widehat{t})\mathbf{P}(\tau^{\widehat{\ell}} < \widehat{t}) + \mathbf{P}(\tau^{\widehat{\ell}} \ge \widehat{t})\\ &\le \mathbf{P}(\tilde{N}_{\widehat{t}} < \widehat{n} | \tau^{\widehat{\ell}} < \widehat{t})\mathbf{P}(\tau^{\widehat{\ell}} < \widehat{t}) + \mathbf{P}(\tau^{\widehat{\ell}} \ge \widehat{t})\\ &\le \mathbf{P}(\tilde{N}_{\widehat{t}} < \widehat{n}) + \mathbf{P}(\tau^{\widehat{\ell}} \ge \widehat{t}),\end{aligned} \tag{F.4}$$

where the third inequality follows from the fact that on $\tau^{\widehat{\ell}} < \widehat{t}$ we have $\tilde{N}_{\widehat{t}} \le \tilde{N}_{\widehat{\ell}}$, and the fourth inequality follows from the independence of $\tau^{\widehat{\ell}}$ and $\tilde{N}_t$.

Now observe that because each agent $j \ne i$ stays in the auxiliary system for a time distributed independently and exponentially with rate λ, the probability that the agent $j \ne i$ is still in the auxiliary system by time $\widehat{t}$ is equal to $\exp(-\lambda\widehat{t}) = 1/\sqrt{n-1}$. Thus, the number of agents $\tilde{N}_{\widehat{t}}$ in the auxiliary system at time $\widehat{t}$ is distributed as $1 + \text{Bin}(n-1, 1/\sqrt{n-1})$, where $\text{Bin}(\cdot,\cdot)$ denotes the binomial distribution. (Recall that, in the auxiliary system, agent i never leaves.) Now note that $\mathbf{E}[\text{Bin}(n-1, 1/\sqrt{n-1})] = \sqrt{n-1} > \widehat{n} - 1$. From this, we obtain

$$\begin{aligned}\mathbf{P}(\tilde{N}_{\widehat{t}} < \widehat{n}) &= \mathbf{P}\left(\text{Bin}\left(n-1, \frac{1}{\sqrt{n-1}}\right) < \widehat{n} - 1\right)\\ &\le \mathbf{P}\left(\text{Bin}\left(n-1, \frac{1}{\sqrt{n-1}}\right) < \frac{1}{2}\sqrt{n-1}\right)\\ &\le \exp\left(-\frac{1}{8}\sqrt{n-1}\right),\end{aligned} \tag{F.5}$$

where we have used the Chernoff bound (Mitzenmacher and Upfal 2005) for the lower tail of the binomial distribution in the last inequality.

Next, note that $\tau^\ell \sim \text{Gamma}(\ell, \lambda)$ because τ^ℓ is the sum of ℓ independently and exponentially distributed time intervals. Hence, from Markov's inequality, we obtain

$$\mathbf{P}(\tau^{\widehat{\ell}} > \widehat{t}) \le \frac{\mathbf{E}[\tau^{\widehat{\ell}}]}{\widehat{t}} = \frac{\widehat{\ell}}{\lambda\widehat{t}} \le \frac{2}{\sqrt{\log(n-1)}}. \tag{F.6}$$

Thus, combining (F.3)–(F.6), we obtain, for all $\ell \le \widehat{\ell}$,

$$\mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \le \tau^\phi\}\right] \le \frac{1}{2}(1-\gamma)^2\underline{V} + \|F\|_\infty\left(\exp(-\frac{1}{8}\sqrt{n-1}) + \frac{2}{\sqrt{\log(n-1)}}\right).$$

Thus, using (F.2), we have

$$\begin{aligned}V_{\mathsf{st}}^\phi(z,n) &= \sum_{\ell=1}^{\widehat{\ell}} \gamma^{\ell-1}\mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \le \tau^\phi\}\right] + \sum_{\ell=\widehat{\ell}+1}^{\infty} \gamma^{\ell-1}\mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \le \tau^\phi\}\right]\\ &\quad + \gamma V_{\mathsf{sw}}\\ &\le \frac{1}{1-\gamma}\left(\frac{1}{2}(1-\gamma)^2\underline{V} + \|F\|_\infty\left(\exp(-\frac{1}{8}\sqrt{n-1}) + \frac{2}{\sqrt{\log(n-1)}}\right)\right)\\ &\quad + \gamma^{\widehat{\ell}}\|F\|_\infty + \gamma V_{\mathsf{sw}},\end{aligned}$$

where, in the inequality, we use the fact that $F(Z_\ell, N_\ell) \le \|F\|_\infty$ for all $\ell > \widehat{\ell}$. Thus, we obtain

$$V_{\mathsf{st}}^\phi(z,n) \le \frac{1}{1-\gamma}\left(\frac{1}{2}(1-\gamma)^2\underline{V} + \|F\|_\infty\left(\exp(-\frac{1}{8}\sqrt{n-1}) + \frac{2}{\sqrt{\log(n-1)}} + \gamma^{\widehat{\ell}}(1-\gamma)\right)\right) + \gamma V_{\mathsf{sw}}.$$

Now note that because $n \ge K_{\max} \ge K_1$, we have

$$\|F\|_\infty\left(\exp(-\frac{1}{8}\sqrt{n-1}) + \frac{2}{\sqrt{\log(n-1)}} + \gamma^{\widehat{\ell}}(1-\gamma)\right) < \frac{(1-\gamma)^2\underline{V}}{4}.$$

Thus, we obtain

$$V_{\mathsf{st}}^\phi(z,n) \le \frac{1}{1-\gamma}\left(\frac{(1-\gamma)^2\underline{V}}{2} + \frac{(1-\gamma)^2\underline{V}}{4}\right) + \gamma V_{\mathsf{sw}} = \frac{3(1-\gamma)}{4}\underline{V} + \gamma V_{\mathsf{sw}}.$$

Because this inequality holds for all strategies $\phi \in \Pi$, and because $V_{\mathsf{sw}} \ge \underline{V}$, we obtain $\mathcal{V}_{\mathsf{st}}(z, n; \xi, V_{\mathsf{sw}}) < V_{\mathsf{sw}}$ for all $z \in \mathbb{Z}$ and all $n \ge K_{\max}$.

Let $\Upsilon = \widehat{\Pi} \times [\underline{V}, \overline{V}]$. The preceding lemma implies that for any $\zeta \in \mathcal{X}(\xi, V_{\mathsf{sw}})$ with $(\xi, V_{\mathsf{sw}}) \in \Upsilon$, it must be the case that $\zeta(z, n) = 0$ for all $z \in \mathbb{Z}$ and all $n \ge K_{\max}$. From the definition of $\widehat{\Pi}$, this implies that $\mathcal{X}(\xi, V_{\mathsf{sw}}) \subseteq \widehat{\Pi}$ for all $(\xi, V_{\mathsf{sw}}) \in \Upsilon$. Thus, we can view the map $(\xi, V_{\mathsf{sw}}) \to \mathcal{X}(\xi, V_{\mathsf{sw}})$ as defining a correspondence $\mathcal{X} : \Upsilon \rightrightarrows \widehat{\Pi}$. □

Appendix G. Upper Hemicontinuity of $\mathcal{R}$

For $(\xi, V_{\mathsf{sw}}) \in \Upsilon$, define the map $\mathcal{R}$ as $\mathcal{R}(\xi, V_{\mathsf{sw}}) = \mathcal{X}(\xi, V_{\mathsf{sw}}) \times \{\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})\}$. Note that from Lemmas E.2–F.1, we obtain that $\mathcal{R}(\xi, V_{\mathsf{sw}}) \subseteq \Upsilon$ for any $(\xi, V_{\mathsf{sw}}) \in \Upsilon$. This implies that we can view the map $\mathcal{R}$ as a correspondence $\mathcal{R} : \Upsilon \rightrightarrows \Upsilon$. In this appendix, we seek to show that this correspondence is upper hemicontinuous. This result is directly used in the proof of Theorem 1.

To prove this, we first show that the value functions $\mathcal{V}(\xi, V_{\mathsf{sw}})$ and $\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})$ are jointly continuous in $(\xi, V_{\mathsf{sw}}) \in \Upsilon$. In the following, we use the following notation: for $U \in \mathcal{C}_b(\mathbb{S})$, let $\|U\|_* \triangleq \max_{z \in \mathbb{Z}, n < K_{\max}} |U(z, n)|$. Note that $\|U\|_* \le \|U\|_\infty$.

Lemma G.1. *The map $(\xi, V_{\mathsf{sw}}) \to \mathcal{V}(\xi, V_{\mathsf{sw}})$ is (jointly) continuous in $(\xi, V_{\mathsf{sw}}) \in \Upsilon$.*

Proof of Lemma G.1. For $(\xi^i, V^i_{\mathsf{sw}}) \in \Upsilon$ for $i = 1, 2$, let $W_i(z, n) = \mathcal{V}(z, n; \xi^i, V^i_{\mathsf{sw}})$. Using the definition of $\mathbf{T}$ and Lemma F.1, we obtain $W_i(z, n) = F(z, n) + \gamma V^i_{\mathsf{sw}}$ for all $z \in \mathbb{Z}$ and $n \ge K_{\max}$. This implies that

$$|W_1(z, n) - W_2(z, n)| \le \gamma |V^1_{\mathsf{sw}} - V^2_{\mathsf{sw}}|, \quad \text{for } z \in \mathbb{Z} \text{ and } n \ge K_{\max}.$$

This implies that $\|W_1 - W_2\|_\infty \le \max\{\|W_1 - W_2\|_*, \gamma |V^1_{\mathsf{sw}} - V^2_{\mathsf{sw}}|\}$.

Next, we have

$$\begin{aligned}
\|W_1 - W_2\|_* &= \|\mathbf{T}(\xi^1, V^1_{\mathsf{sw}}, W_1) - \mathbf{T}(\xi^2, V^2_{\mathsf{sw}}, W_2)\|_* \\
&\le \|\mathbf{T}(\xi^1, V^1_{\mathsf{sw}}, W_1) - \mathbf{T}(\xi^2, V^2_{\mathsf{sw}}, W_1)\|_* \\
&\quad + \|\mathbf{T}(\xi^2, V^2_{\mathsf{sw}}, W_1) - \mathbf{T}(\xi^2, V^2_{\mathsf{sw}}, W_2)\|_* \\
&\le \|\mathbf{T}(\xi^1, V^1_{\mathsf{sw}}, W_1) - \mathbf{T}(\xi^2, V^2_{\mathsf{sw}}, W_1)\|_* \\
&\quad + \|\mathbf{T}(\xi^2, V^2_{\mathsf{sw}}, W_1) - \mathbf{T}(\xi^2, V^2_{\mathsf{sw}}, W_2)\|_\infty \\
&\le \|\mathbf{T}(\xi^1, V^1_{\mathsf{sw}}, W_1) - \mathbf{T}(\xi^2, V^2_{\mathsf{sw}}, W_1)\|_* + \gamma \|W_1 - W_2\|_\infty,
\end{aligned}$$

where we have used Lemma E.1 in the last inequality. Using the fact that $\|W_1 - W_2\|_* \le \|W_1 - W_2\|_\infty \le \max\{\|W_1 - W_2\|_*, \gamma |V^1_{\mathsf{sw}} - V^2_{\mathsf{sw}}|\} \le \|W_1 - W_2\|_* + \gamma |V^1_{\mathsf{sw}} - V^2_{\mathsf{sw}}|$ and after some straightforward algebra, we obtain

$$\|W_1 - W_2\|_\infty \le \frac{1}{1 - \gamma} \left(\|\mathbf{T}(\xi^1, V^1_{\mathsf{sw}}, W_1) - \mathbf{T}(\xi^2, V^2_{\mathsf{sw}}, W_1)\|_* + \gamma |V^1_{\mathsf{sw}} - V^2_{\mathsf{sw}}| \right).$$

From Lemma G.2, we obtain that the first term in the parentheses can be made arbitrarily small by setting $\|\xi^1 - \xi^2\|_\infty$ and $|V^1_{\mathsf{sw}} - V^2_{\mathsf{sw}}|$ correspondingly small enough. Thus, we conclude that $\mathcal{V}(\xi, V_{\mathsf{sw}})$ is jointly continuous in $(\xi, V_{\mathsf{sw}}) \in \Upsilon$.

The following auxiliary lemma is used in the proof of Lemma G.1. □

Lemma G.2. *Let $(\xi^m, V^m_{\mathsf{sw}}) \in \Upsilon$ with $(\xi^m, V^m_{\mathsf{sw}}) \to (\xi, V_{\mathsf{sw}}) \in \Upsilon$ as $m \to \infty$. For any $U \in \mathcal{C}_b(\mathbb{S})$, we have $\|\mathbf{T}(\xi^m, V^m_{\mathsf{sw}}, U) - \mathbf{T}(\xi, V_{\mathsf{sw}}, U)\|_* \to 0$ as $m \to \infty$.*

Proof of Lemma G.2. Let $(\xi^m, V^m_{\mathsf{sw}})$ be as in the statement of the lemma, and let $W_m = \mathbf{T}(\xi^m, V^m_{\mathsf{sw}}, U)$ and $W = \mathbf{T}(\xi, V_{\mathsf{sw}}, U)$. By definition of $\mathbf{T}$, we have

$$\begin{aligned}
|W_m(z, n) - W(z, n)| &= \gamma \,|\max\{\mathbf{E}^m[U(Z_\tau, N_\tau) \,|\, z, n], V^m_{\mathsf{sw}}\} \\
&\qquad - \max\{\mathbf{E}^\xi[U(Z_\tau, N_\tau) \,|\, z, n], V_{\mathsf{sw}}\}| \\
&\le \gamma \max\{|\mathbf{E}^m[U(Z_\tau, N_\tau) \,|\, z, n] \\
&\qquad - \mathbf{E}^\xi[U(Z_\tau, N_\tau) \,|\, z, n]|, |V^m_{\mathsf{sw}} - V_{\mathsf{sw}}|\},
\end{aligned}$$

where we let $\mathbf{E}^m = \mathbf{E}^{\xi^m}$. Thus, it suffices to show that the first term inside the maximization converges to zero as $m \to \infty$ for all $z \in \mathbb{Z}$ and $n < K_{\max}$. Observe that because $U \in \mathcal{C}_b(\mathbb{S})$ and τ is exponentially distributed with parameter λ, we have

$$\begin{aligned}
\mathbf{E}^\xi[U(Z_\tau, N_\tau) \,|\, z, n] &= \int_0^\infty \lambda \exp(-\lambda t) \mathbf{E}^\xi[U(Z_t, N_t) \,|\, z, n, \tau = t] dt \\
&= \int_0^T \lambda \exp(-\lambda t) \mathbf{E}^\xi[U(Z_t, N_t) \,|\, z, n, \tau = t] dt \\
&\quad + \int_T^\infty \lambda \exp(-\lambda t) \mathbf{E}^\xi[U(Z_t, N_t) \,|\, z, n, \tau = t] dt
\end{aligned}$$

with similar expressions for ξ^m in place of ξ. For a large enough value of $T > 0$, the second term in the last equation can be made arbitrarily small (uniformly for ξ and all ξ^m). Thus, again, it suffices to show that the first term in the last equation is continuous in (ξ, V_{sw}) for all $z \in \mathbb{Z}$ and $n < K_{\max}$ and for large enough T.

Now, using the definition (1) of the transition-rate matrix Q^ξ of the chain $\mathsf{MC}(\xi, \kappa(\xi))$ (and similarly $\mathsf{Q}^m = \mathsf{Q}^{\xi^m}$ of the chain $\mathsf{MC}(\xi^m, \kappa(\xi^m))$), we obtain that $\mathsf{Q}^m((u,k) \to (v,\ell)) \longrightarrow \mathsf{Q}((u,k) \to (v,\ell))$ as $m \to \infty$ for all $(u,k),(v,\ell) \in \mathbb{S}$. Then, from Xia (1994, example 1.1) or Ethier and Kurtz (1986, problem 8), we obtain that the measure $\mathbf{P}^m(\cdot \mid z, n, \tau = t)$ converges weakly to $\mathbf{P}^\xi(\cdot \mid z, n, \tau = t)$. From this, we conclude that $\int_0^T \lambda \exp(-\lambda t)\mathbf{E}^m[U(Z_t, N_t) \mid z, n, \tau = t]dt$ converges to $\int_0^T \lambda \exp(-\lambda t)\mathbf{E}^\xi[U(Z_t, N_t) \mid z, n, \tau = t]dt$ as $m \to \infty$. This completes the proof. □

The continuity of $\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})$ is then obtained as a corollary of Lemma G.1.

Lemma G.3. *The value function $\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})$ is jointly continuous in $(\xi, V_{\mathsf{sw}}) \in \Upsilon$.*

Proof of Lemma G.3. Recall the definition (23) of $\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})$:

$$\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}}) = \sum_{(z,n)\in\mathbb{S}} \pi(z,n)\mathcal{V}_{\mathsf{sw}}(z, n+1; \xi, V_{\mathsf{sw}}),$$

where $\pi = \pi(\xi, \kappa(\xi))$ is the invariant distribution of $\mathsf{MC}(\xi, \kappa(\xi))$. From Lemma E.2, we have $\|\mathcal{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}})\|_\infty \le \overline{\mathsf{V}}$. Also, note that Lemmas C.1 and D.4 imply that the invariant distribution $\pi(\xi, \kappa(\xi))$ is continuous. Moreover, from Lemma C.2, we obtain that the set of invariant distributions Γ is tight. These results together imply that it suffices to show that $\mathcal{V}_{\mathsf{st}}(z, n; \xi, V_{\mathsf{sw}})$ is uniformly continuous in $(\xi, V_{\mathsf{sw}}) \in \Upsilon$ for all $z \in \mathbb{Z}$ and all $n < M$ for some large enough M.

Let $(\xi^m, V_{\mathsf{sw}}^m) \in \Upsilon$ with $(\xi^m, V_{\mathsf{sw}}^m) \to (\xi, V_{\mathsf{sw}}) \in \Upsilon$ as $m \to \infty$. We have

$$\begin{aligned}
&|\mathcal{V}_{\mathsf{st}}(z, n; \xi^m, V_{\mathsf{sw}}^m) - \mathcal{V}_{\mathsf{st}}(z, n; \xi, V_{\mathsf{sw}})| \\
&\le |\mathbf{E}^{\xi^m}[\mathcal{V}(Z_\tau, N_\tau; \xi^m, V_{\mathsf{sw}}^m) \mid z, n] - \mathbf{E}^\xi[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}}) \mid z, n]| \\
&\le |\mathbf{E}^{\xi^m}[\mathcal{V}(Z_\tau, N_\tau; \xi^m, V_{\mathsf{sw}}^m) \mid z, n] - \mathbf{E}^{\xi^m}[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}}) \mid z, n]| \\
&\quad + |\mathbf{E}^{\xi^m}[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}}) \mid z, n] - \mathbf{E}^\xi[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}}) \mid z, n]| \\
&\le |\mathcal{V}(\xi^m, V_{\mathsf{sw}}^m) - \mathcal{V}(\xi, V_{\mathsf{sw}})\|_\infty + |\mathbf{E}^{\xi^m}[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}}) \mid z, n] \\
&\quad - \mathbf{E}^\xi[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}}) \mid z, n]|.
\end{aligned}$$

From Lemma G.1, we obtain that as $m \to \infty$, the first term converges to zero. Moreover, from the same argument as in the proof of Lemma G.1, we obtain that $\mathbf{E}^{\xi^m}[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}}) \mid z, n] \to \mathbf{E}^\xi[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}}) \mid z, n]$ as $m \to \infty$ for each $(z,n) \in \mathbb{S}$. From this, we conclude that $\mathcal{V}_{\mathsf{st}}(z, n; \xi, V_{\mathsf{sw}})$ is uniformly continuous in $(\xi, V_{\mathsf{sw}}) \in \Upsilon$ for all $z \in \mathbb{Z}$ and all $n < M$ for large enough M. □

We are now ready to show that the correspondence $\mathcal{R}$ is upper hemicontinuous.

Lemma G.4. *The correspondence $\mathcal{R} : \Upsilon \rightrightarrows \Upsilon$ is upper hemicontinuous.*

Proof of Lemma G.4. By definition, $\mathcal{R}(\xi, V_{\mathsf{sw}}) = \mathcal{X}(\xi, V_{\mathsf{sw}}) \times \{\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})\}$ for $(\xi, V_{\mathsf{sw}}) \in \Upsilon$. From Lemma G.3, we obtain that $\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})$ is jointly continuous in $(\xi, V_{\mathsf{sw}}) \in \Upsilon$. Thus, it suffices to show that the correspondence $\mathcal{X} : \Upsilon \rightrightarrows \widehat{\Pi}$ is upper hemicontinuous.

Consider a sequence $(\xi^n, V_{\mathsf{sw}}^n, \zeta^n) \to (\xi, V_{\mathsf{sw}}, \zeta)$ as $n \to \infty$ such that $\zeta^n \in \mathcal{X}(\xi^n, V_{\mathsf{sw}}^n)$ for each $n \ge 0$. By continuity of $\mathcal{V}_{\mathsf{st}}(\cdot)$, we obtain that if $\mathcal{V}_{\mathsf{st}}(z, n, \xi, V_{\mathsf{sw}}) > V_{\mathsf{sw}}$ for some $(z,n) \in \mathbb{S}$, then, for all large enough m, we must have $\mathcal{V}_{\mathsf{st}}(z, n, \xi^m, V_{\mathsf{sw}}^m) > V_{\mathsf{sw}}^m$ and, hence, $\zeta^m(z,n) = 1$. Similarly, if $\mathcal{V}_{\mathsf{st}}(z, n, \xi, V_{\mathsf{sw}}) < V_{\mathsf{sw}}$, then $\zeta^m(z,n) = 0$ for all large enough m. Because $\zeta^m \to \zeta$, this implies that $\zeta(z,n) = 1$ if $\mathcal{V}_{\mathsf{st}}(z, n, \xi, V_{\mathsf{sw}}) > V_{\mathsf{sw}}$ and $\zeta(z,n) = 0$ if $\mathcal{V}_{\mathsf{st}}(z, n, \xi, V_{\mathsf{sw}}) < V_{\mathsf{sw}}$. Thus, we obtain that $\zeta(z,m) \in \mathcal{X}(\xi, V_{\mathsf{sw}})$. □

Appendix H. Existence of an Optimal Threshold Strategy

In this appendix, we provide the proof of Theorem 2. We prove this result in two steps. First, we prove Lemma 1, which states that the value function $V_{\mathsf{st}} : \mathbb{S} \to \mathbb{R}$ is nonincreasing in the number of agents n at the location for any fixed resource level $z \in \mathbb{Z}$. Second, we show in Lemma F.1 that $\lim_{n\to\infty} V_{\mathsf{st}}(z,n) \le \gamma V_{\mathsf{sw}}$ for all $z \in \mathbb{Z}$. Therefore, there always exists a threshold strategy in the set of best responses $\mathsf{OPT}(\xi, \kappa, V_{\mathsf{sw}})$.

Proof of Lemma 1. We define a partial order $\preccurlyeq_p$ on the state space $\mathbb{S}$ of $\mathsf{MC}(\xi, \kappa)$ as follows: for $(z_1, n_1), (z_2, n_2) \in \mathbb{S}$, $(z_1, n_1) \preccurlyeq_p (z_2, n_2)$ if and only if $z_1 = z_2$ and $n_1 \le n_2$. For any function $f : \mathbb{S} \to \mathbb{R}$, we say that f is decreasing with respect to $\preccurlyeq_p$ if, for all $(z_1, n_1), (z_2, n_2) \in \mathbb{S}$ such that $(z_1, n_1) \preccurlyeq_p (z_2, n_2)$, we have $f(z_1, n_1) \ge f(z_2, n_2)$.

Thus, our goal is to show that the value function V_{st} of $\mathsf{DEC}(\xi, \kappa, V_{\mathsf{sw}})$ is decreasing with respect to $\preccurlyeq_p$. We note that the property "decreasing with respect to $\preccurlyeq_p$" is a closed convex cone property for functions on $\mathbb{S}$, as defined in Smith and McCardle (2002). Thus, using their proposition 5, we can conclude that V_{st} has this property if the following two conditions hold:

(1) The resource-sharing function F is decreasing with respect to $\preccurlyeq_p$.

(2) Let $(Z_t, N_t) \sim \mathsf{MC}(\xi, \kappa)$. For any $(z,n) \in \mathbb{S}$, let $\nu_{(z,n)}$ be the probability distribution of $(Z_\tau, N_\tau) \sim \nu_{(z,n)}$ conditioning on $(Z_0, N_0) = (z,n)$, where τ is distributed independently as an exponential with rate λ, denoting the first decision epoch of a fixed

agent. Then, for any $f:\mathbb{S}\to\mathbb{R}$ that is decreasing with respect to $\preccurlyeq_p$, it must hold that $\mathbf{E}[f(Z,N)\,|\,(Z,N)\sim\nu_{(z_1,n_1)}]\geq\mathbf{E}[f(Z,N)\,|\,(Z,N)\sim\nu_{(z_2,n_2)}]$ for all $(z_1,n_1)\preccurlyeq_p(z_2,n_2)$.

Because $F(z,n)$ is decreasing in n for each $z\in\mathbb{Z}$, we immediately obtain the first condition. We now show that the second condition also holds using a coupling argument.

Suppose that $(z_1,n_1)\preccurlyeq_p(z_2,n_2)$. By using an argument that is the same as that in the proof of Lemma I.2, we obtain that there exists a coupling of the two processes $(Z_t^{(i)},N_t^{(i)})\sim\mathsf{MC}(\xi,\kappa)$ with $(Z_0^{(i)},N_0^{(i)})=(z_i,n_i)$ for $i=1,2$ such that for all $t\geq 0$, $(Z_t^{(1)},N_t^{(1)})\preccurlyeq_p(Z_t^{(2)},N_t^{(2)})$. Thus, for any f that is decreasing with respect to $\preccurlyeq_p$, we have $f(Z_t^{(1)},N_t^{(1)})\geq f(Z_t^{(2)},N_t^{(2)})$ for all $t\geq 0$, and therefore, $\mathbf{E}[f(Z_\tau^{(1)},N_\tau^{(1)})]\geq\mathbf{E}[f(Z_\tau^{(2)},N_\tau^{(2)})]$, where τ is distributed independently as an exponential with rate λ. Because $(Z_\tau^{(i)},N_\tau^{(i)})\sim\nu_{(z_i,n_i)}$ for $i=1,2$, we obtain the result. □

Appendix I. Coupling Results

In this appendix, we obtain structural properties of the Markov chain $\mathsf{MC}(\xi,\kappa)$ by coupling the chain with an $M/M/\infty$ queue. Let $\mathsf{M/M/\infty}(\lambda,\kappa)$ denote an (independent) $M/M/\infty$ queue with arrival rate κ and service rate λ. We begin with the following simple result that states that a queue with a higher arrival rate and/or lower service rate is more likely to have more agents in the queue. The proof is straightforward and omitted.

Lemma I.1. *Let $N_t^{(i)}$ for $i=1,2$ denote the number of agents at time t in an (independent) $M/M/\infty$ queue with arrival rate κ_i and service rate λ_i. Suppose that $N_0^{(1)}=N_0^{(2)}$ and that one of the following two conditions holds: (1) $\lambda_1=\lambda_2$ and $\kappa_1\leq\kappa_2$ or (2) $\lambda_1\geq\lambda_2$ and $\kappa_1=\kappa_2$. Then, for all $t\geq 0$, $N_t^{(1)}$ is stochastically dominated by $N_t^{(2)}$; that is, for all $n\in\mathbb{N}_0$, we have $\mathbf{P}(N_t^{(1)}\geq n)\leq\mathbf{P}(N_t^{(2)}\geq n)$.*

In the proof of Theorem 1, we frequently compare the $\mathsf{MC}(\xi,\kappa)$ process for two (or more) different values of (ξ,κ) to show the monotonicity of various quantities. Our next result justifies these stochastic comparisons. Before we state the lemma, we make the following definition of stochastic dominance. Let $(Z_t^{(i)},N_t^{(i)})\sim\mathsf{MC}(\xi^i,\kappa^i)$ with $(Z_0^{(i)},N_0^{(i)})=(z_i,n_i)$ for $i=1,2$. We say the process $(Z_t^{(1)},N_t^{(1)})$ is stochastically dominated by the process $(Z_t^{(2)},N_t^{(2)})$ if

$$\mathbf{P}(Z_t^{(1)}=z,N_t^{(1)}\geq n)\leq\mathbf{P}(Z_t^{(2)}=z,N_t^{(2)}\geq n),\quad\text{for all }(z,n)\in\mathbb{S}.$$

In this case, we denote as $(Z_t^{(1)},N_t^{(1)})\preccurlyeq_{\mathsf{sd}}(Z_t^{(2)},N_t^{(2)})$. Note that this also implies that $N_t^{(1)}$ is stochastically dominated by $N_t^{(2)}$ under the usual sense of stochastic dominance.

Lemma I.2. *Let $\xi\in\Pi$ and $\kappa>0$. Let $(Z_t,N_t)\sim\mathsf{MC}(\xi,\kappa)$.*

(1) *Let $\kappa_0\geq\kappa$, and let $\xi_0\in\Pi$ be such that $\xi_0(z,n)\geq\xi(z,n)$ for all $(z,n)\in\mathbb{S}$. Then we have $(Z_t,N_t)\preccurlyeq_{\mathsf{sd}}(Z_t^{(0)},N_t^{(0)})$ for all $t\geq 0$, where $(Z_t^{(0)},N_t^{(0)})\sim\mathsf{MC}(\xi_0,\kappa_0)$ with $Z_0^{(0)}=Z_0$ and $N_0^{(0)}\geq N_0$.*

(2) *Let $X_t^i\sim\mathsf{M/M/\infty}(\lambda_i,\kappa)$ for $i=1,2$ be two independent processes with $X_0^1=X_0^2=N_0$, where $\lambda_1=\lambda$ and $\lambda_2=(1-\gamma)\lambda$. Then we have, for all $t\geq 0$, $(Z_t,X_t^1)\preccurlyeq_{\mathsf{sd}}(Z_t,N_t)\preccurlyeq_{\mathsf{sd}}(Z_t,X_t^2)$.*

Proof of Lemma I.2. First note that the second statement in the lemma is implied by the first. In particular, let $\xi_1(z,n)=0$ and $\xi_2(z,n)=1$ for all $(z,n)\in\mathbb{S}$. Then, using the first statement in the lemma, we obtain $(Z_t^{(1)},N_t^{(1)})\preccurlyeq_{\mathsf{sd}}(Z_t,N_t)\preccurlyeq_{\mathsf{sd}}(Z_t^{(2)},N_t^{(2)})$, where $(Z_t^{(i)},N_t^{(i)})\sim\mathsf{MC}(\xi_i,\kappa)$ with $(Z_0^{(i)},N_0^{(i)})=(Z_0,N_0)$. The second statement then follows directly by the fact that under ξ_i, the process $(Z_t^{(i)},N_t^{(i)})$ has the same distribution as (Z_t,X_t^i) for $i=1,2$.

To prove the first statement in the lemma, we use a coupling argument. We construct two chains as follows. Let $t_0=0$, $(Z_0,N_0)=(z_0,n_0)$, and $(Z_0^{(0)},N_0^{(0)})=(u_0,v_0)$ with $u_0=z_0$ and $v_0\geq n_0$. For $k=1,2\ldots$, define the following recursively:

(1) Let $\tau_k\sim\mathsf{Exp}(\Delta_k)$, where $\Delta_k\triangleq\sum_{y\neq u_{k-1}}\mu_{u_{k-1},y}+\kappa_0+\lambda v_{k-1}$. Let $t_k=t_{k-1}+\tau_k$.

(2) Let $(Z_t^{(0)},N_t^{(0)})=(u_{k-1},v_{k-1})$ and $(Z_t,N_t)=(z_{k-1},n_{k-1})$ for $t\in[t_{k-1},t_k)$.

(3) For $t=t_k$, let $(Z_t^{(0)},N_t^{(0)})=(u_k,v_k)$, where

$$(u_k,v_k)=\begin{cases}(y,v_{k-1}) \\ \qquad\text{with probability } \mu_{u_{k-1},y}/\Delta_k,\text{ for each } y\in\mathbb{Z}\text{ with } y\neq u_{k-1};\\ (u_{k-1},v_{k-1}+1)\\ \qquad\text{with probability }\kappa_0/\Delta_k;\\ (u_{k-1},v_{k-1}-1)\\ \qquad\text{with probability }\lambda v_{k-1}(1-\gamma\xi_0(u_{k-1},v_{k-1}))/\Delta_k;\\ (u_{k-1},v_{k-1})\\ \qquad\text{with probability }\lambda v_{k-1}\gamma\xi_0(u_{k-1},v_{k-1})/\Delta_k.\end{cases}$$

(4) Define $\zeta_k\triangleq n_{k-1}(1-\gamma\xi(z_{k-1},n_{k-1}))/(v_{k-1}(1-\gamma\xi_0(u_{k-1},v_{k-1})))$ and $\eta_k\triangleq(1-\gamma\xi_0(u_{k-1},v_{k-1})/(\gamma\xi_0(u_{k-1},v_{k-1})))\max(\zeta_k-1,0)$. It is straightforward to verify that $\eta_k\in[0,1]$.

(5) Let $(Z_t, N_t) = (z_k, n_k)$ for $t = t_k$, where

$$(z_k, n_k) = \begin{cases} (u_k, n_{k-1}) & \text{if } u_k \neq u_{k-1}; \\ \begin{cases} (z_{k-1}, n_{k-1}+1) \\ \quad \text{with probability } \frac{\kappa}{\kappa_0}; \\ (z_{k-1}, n_{k-1}) \\ \quad \text{with probability } 1-\frac{\kappa}{\kappa_0}, \end{cases} & \text{if } (u_k, v_k) = (u_{k-1}, v_{k-1}+1); \\ \begin{cases} (z_{k-1}, n_{k-1}-1) \\ \quad \text{with probability} \\ \quad \min(\zeta_k, 1); \\ (z_{k-1}, n_{k-1}) \\ \quad \text{with probability} \\ \quad \max(1-\zeta_k, 0), \end{cases} & \text{if } (u_k, v_k) = (u_{k-1}, v_{k-1}-1); \\ \begin{cases} (z_{k-1}, n_{k-1}-1) \\ \quad \text{with probability } \eta_k; \\ (z_{k-1}, n_{k-1}) \\ \quad \text{with probability } 1-\eta_k, \end{cases} & \text{if } (u_k, v_k) = (u_{k-1}, v_{k-1}). \end{cases}$$

It is straightforward to verify that under this construction, we have $(Z_t^{(0)}, N_t^{(0)}) \sim \mathsf{MC}(\xi_0, \kappa_0)$ and $(Z_t, N_t) \sim \mathsf{MC}(\xi, \kappa)$ with $Z_0^{(0)} = Z_0$ and $N_0 \le N_0^{(0)}$. Furthermore, by construction, we have $z_k = u_k$ for all k and, hence, $Z_t^{(0)} = Z_t$ for all $t \ge 0$.

To show that $N_t \le N_0^{(0)}$ for all $t \ge 0$, we perform induction on k in this construction. Note that $n_0 \le v_0$. Suppose that for some k we have $n_{k-1} \le v_{k-1}$. Then, from the definition, we obtain that $n_k \le v_k$ for all the cases except possibly when $(u_k, v_k) = (u_{k-1}, v_{k-1}-1)$ and $(z_k, n_k) = (z_{k-1}, n_{k-1})$. In this case, if $n_{k-1} < v_{k-1}$, then again we have $n_k \le v_k$. By contrast, if $n_{k-1} = v_{k-1}$, then together with the fact that $z_{k-1} = u_{k-1}$, we obtain $\xi(z_{k-1}, n_{k-1}) \le \xi_0(u_{k-1}, v_{k-1})$, implying that $\zeta_k \ge 1$. However, note that if $\zeta_k \ge 1$, then the event in which $(u_k, v_k) = (u_{k-1}, v_{k-1}-1)$ and $(z_k, n_k) = (z_{k-1}, n_{k-1})$ occurs with zero probability. Thus, we obtain that in all cases, $n_k \le v_k$. This completes the induction step and, hence, the proof. □

Endnotes

[1] Proving this statement rigorously is an interesting direction for future work.

[2] We conjecture that the equilibrium is unique when the resource-sharing function is decreasing and the resource level is binary, the setting we study for comparative statics in Section 5.2. An extensive numerical investigation supports this conjecture, but we do not have a formal proof.

References

Adlakha S, Johari R, Weintraub GY (2015) Equilibria of dynamic games with many players: Existence, approximation, and market structure. *J. Econom. Theory* 156(1):269–316.

Aliprantis CD, Border K (2006) *Infinite Dimensional Analysis: A Hitchhiker's Guide* (Springer Science & Business Media, Berlin).

Arnosti N, Johari R, Kanoria Y (2014) Managing congestion in decentralized matching markets. Babaioff M, Conitzer V, Easley D, eds. *Proc. 15th ACM Conf. Econom. Comput.* (ACM, New York), 451.

Arthur WB (1994) Inductive reasoning and bounded rationality. *Amer. Econom. Rev.* 84(2):406–411.

Balseiro SR, Besbes O, Weintraub GY (2015) Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Sci.* 61(4):864–884.

Banerjee S, Freund D, Lykouris T (2016) Pricing and optimization in shared vehicle systems. Working paper, Cornell University, Ithaca, NY.

Banerjee S, Johari R, Riquelme C (2015) Pricing in ride-sharing platforms: A queueing-theoretic approach. Roughgarden T, Feldman M, Schwarz M, eds. *Proc. 16th ACM Conf. Econom. Comput.* (ACM, New York), 639.

Berge C (1963) *Topological Spaces: Including a Treatment of Multi-Valued Functions, Vector Spaces, and Convexity*. Translated by EM Patterson (Dover, Mineola, NY).

Billingsley P (2013) *Convergence of Probability Measures* (Wiley, Hoboken, NJ).

Bimpikis K, Candogan O, Daniela S (2016) Spatial pricing in ride-sharing networks. Working paper, Cornell University, Ithaca, NY.

Braverman A, Dai J, Liu X, Ying L (2016) Empty-car routing in ridesharing systems. Working paper, Northwestern University, Evanston, IL.

Castillo JC, Knoepfle D, Weyl G (2017) Surge pricing solves the wild goose chase. Daskalakis C, Babaioff M, Moulin H, eds. *Proc. 2017 ACM Conf. Econom. Comput.* (ACM, New York), 241–242.

Chakrabarti AS, Chakrabarti BK, Chatterjee A, Mitra M (2009) The Kolkata Paise restaurant problem and resource utilization. *Physica A: Statist. Mech. Appl.* 388(12):2420–2426.

Chakrabarti BK (2007) Kolkata restaurant problem as a generalised El Farol bar problem. Chatterjee A, Chakrabarti BK, eds. *Econophysics of Markets and Business Networks* (Springer, Berlin), 239–246.

Chen MK (2016) Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform. Bergemann D, Chen Y, Conitzer V, eds. *Proc. 2016 ACM Conf. Econom. Comput.* (ACM, New York), 455.

Durrett R, Levin S (1994a) The importance of being discrete (and spatial). *Theoret. Population Biol.* 46(3):363–394.

Durrett R, Levin SA (1994b) Stochastic spatial models: A user's guide to ecological applications. *Philos. Trans. Roy. Soc. London Ser. B: Biol. Sci.* 343(1305):329–350.

Ethier SN, Kurtz TG (1986) *Markov Processes: Characterization and Convergence* (Wiley, Hoboken, NJ).

Fudenberg D, Tirole J (1991a) *Game Theory* (MIT Press, Cambridge, MA).

Fudenberg D, Tirole J (1991b) Perfect Bayesian equilibrium and sequential equilibrium. *J. Econom. Theory* 53(2):236–260.

Ghosh A, Chatterjee A, Mitra M, Chakrabarti BK (2010) Statistics of the Kolkata Paise restaurant problem. *New J. Phys.* 12(7):075033.

Hopenhayn HA (1992) Entry, exit and firm dynamics in long run equilibrium. *Econometrica* 60(5):1127–1150.

Huang M, Caines PE, Malhamé RP (2007) Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralizedϵ-Nash equilibria. *IEEE Trans. Automatic Control* 52(9):1560–1571.

Iyer K, Johari R, Sundararajan M (2014) Mean field equilibria of dynamic auctions with learning. *Management Sci.* 60(12):2949–2970.

Jovanovic B, Rosenthal RW (1988) Anonymous sequential games. *J. Math. Econom.* 17(1):77–87.

Keymer JE, Marquet PA, Velasco-Hernández JX, Levin SA (2000) Extinction thresholds and metapopulation persistence in dynamic landscapes. *Amer. Naturalist* 156(5):478–494.

Lachapelle A, Wolfram MT (2011) On a mean field game approach modeling congestion and aversion in pedestrian crowds. *Transportation Res. Part B: Methodological* 45(10):1572–1589.

Lasry JM, Lions PL (2007) Mean field games. *Japanese J. Math.* 2(1):229–260.

Le Van C, Stachurski J (2007) Parametric continuity of stationary distributions. *Econom. Theory* 33(2):333–348.

Levin R (1970) Extinction. Gerstenhaber M, ed. *Some Mathematical Problems in Biology* (American Mathematical Society, Providence, RI), 77–107.

Li J, Bhattacharyya R, Paul S, Shakkottai S, Subramanian V (2017) Incentivizing sharing in realtime D2D streaming networks: A mean field game perspective. *IEEE/ACM Trans. Networking* 25(1):3–17.

Manjrekar M, Ramaswamy V, Shakkottai S (2014) A mean field game approach to scheduling in cellular systems. Leon-Garcia A, Pack S, McNair J, eds. *IEEE INFOCOM 2014-IEEE Conf. Comput. Comm.* (IEEE, Piscataway, NJ), 1554–1562.

Mitzenmacher M, Upfal E (2005) *Probability and Computing: Randomized Algorithms and Probabilistic Analysis* (Cambridge University Press, Cambridge, UK).

Molofsky J (1994) Population dynamics and pattern formation in theoretical populations. *Ecology* 75(1):30–39.

Nelder JA, Mead R (1965) A simplex method for function minimization. *Comput. J.* 7(4):308–313.

New York City Taxi and Limousine Commission (2017) New York City Taxi & Limousine Commission trip record data. Accessed March 20, 2018, http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml.

Nisan N, Roughgarden T, Tardos E, Vazirani VV (2007) *Algorithmic Game Theory* (Cambridge University Press, Cambridge, UK).

Rosenthal RW (1973) A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory* 2(1):65–67.

Smith JE, McCardle KF (2002) Structural properties of stochastic dynamic programs. *Oper. Res.* 50(5):796–809.

Weather Underground (2017) Weather history and data archive. Accessed March 20, 2018, https://www.wunderground.com/history/.

Weintraub GY, Benkard CL, Van Roy B (2008) Markov perfect industry dynamics with many firms. *Econometrica* 76(6):1375–1411.

Weintraub GY, Benkard CL, Van Roy B (2011) Industry dynamics: Foundations for models with an infinite number of firms. *J. Econom. Theory* 146(5):1965–1994.

Xia A (1994) Weak convergence of Markov processes with extended generators. *Ann. Probab.* 22(4):2183–2202.

Xu J, Hajek B (2013) The supermarket game. *Stochastic Systems* 3(2):405–441.