SPATIAL RESOURCE COMPETITION GAMES

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Pu Yang

August 2019

SPATIAL RESOURCE COMPETITION GAMES

Pu Yang, Ph.D.

Cornell University 2019

This dissertation studies spatial resource competition settings where nomadic agents migrate across different locations, competing for time-varying and location-specific resources. Such setting arises in crowd-sourced transportation services, online communities, and traditional location-based economic activities. In these settings, many factors influence the agents' behavior: the resource dynamics, the way resource is shared among agents at different locations, the information available to the agents, etc. Understanding agents' behavior in equilibrium and how their decisions depend on these factors can help system operators design better mechanisms to improve social welfare of systems.

Analyzing these settings systematically is challenging, since agents' decisions influence each other spatially and temporally in a complicated nested way. This dissertation aims at building models that capture the essentials of spatial resource competitions, and are analytically tractable, to help understand the nature of agents' interactions in these settings, from a game theoretical point of view.

We first provide a general model for spatial resource competition settings. Using the methodology of mean field approximation, we analyze the dynamics and the game between the agents at a single location, in the limit where there are infinitely many locations. We characterize an equilibrium for agents in the mean field model where agents' equilibrium strategies have a simple Markovian struc-

ture.

We then provide a method to approximately compute the equilibrium for a common case of resource competition where the amount of resource each agent gets decreases as the number of agents competing with her increases. We study numerically how different factors affect agents' equilibrium behavior. We also extend our model and analysis to more general settings where locations are non-homogeneous and there is a two-sided market at each location.

Finally, we study information design problem in spatial resource competition scenarios. That is, how should a system operator communicate her extra information about the system to the agents in order to better position them and increase their welfare? We study both private and public signaling mechanisms. For private signaling, we provide a method to obtain the optimal mechanism in polynomial time. For public signaling, we show the sender preferred equilibrium has a simple threshold structure and characterize the structure of the optimal public mechanism under the sender preferred equilibrium. We show via numerical computations that the optimal private and public signaling mechanisms achieve substantially higher social welfare compared with no information sharing or full information sharing in many settings.

BIOGRAPHICAL SKETCH

Pu Yang grew up in Guiyang, a beautiful town in southwestern China. Before coming to Cornell University, he got his Bachelor's degree in Tsinghua University in Beijing, China, majoring in Electronic Engineering.

At Cornell, Pu's academic interest led him to explore many fields before he finally chose to focus on game theory. He got a lot of enjoyment from his research towards this dissertation.

During semester breaks, he travelled widely across United States and abroad, collecting many cherished memories. In his spare time, he enjoyed hiking, dancing, swimming and skiing in nearby areas. He also remembers the fun times playing board games and solving puzzles with his friends, especially those on windy and snowy nights of Ithaca, this weird yet a little bit adorable small town.

To my grandma, for teaching me how to play gwat-pai, giving me her calligraphy of poems, and drawing an adorable whale for me.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisors, Krishnamurthy Iyer and Peter Frazier, for their invaluable academic and personal support. My sincere gratitude also goes to my dissertation committee members, Siddhartha Banerjee and Robert Kleinberg, for their insightful comments about my work. Especially, I would like to thank Bobby for the idea that led to simplification of the private persuasion linear program in Chapter 4. I thank Cornell University and the School of Operations Research and Information Engineering, for supporting me financially and providing me a comfortable working space. Thanks to my parents, for their love and support. Especially my mother, who always supports my decisions in her unique way as a friend. Thanks to all the friends, teachers and co-workers I met in my life. The influence many of them had on me is subtle, hard-to-notice yet impressive and unforgettable, and contributes to this dissertation in its own way. Finally, thanks to time and life itself, for letting me gradually find my desires and passions, allowing me to know myself better and become a person I appreciate more during the past years.

Part of the work in this dissertation is supported by NSF CMMI-1462592, NSF CMMI-1254298 and AFOSR FA9550-15-1-0038.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

Consider a setting where a group of nomadic agents explore and compete for resources that are time-varying, stochastic and location-specific. Many real-world scenarios fall into this setting. For example, in the sharing economy, such as the crowd-sourced transportation services like Uber and Lyft, or the crowdsourced food delivery services like GrubHub and DoorDash, drivers choose neighborhoods to locate themselves and provide service. The payoffs they earn depend on the number of customers requesting service within the neighborhood they choose (the location-specific resource), and the number of other drivers working there. The overall resource level varies stochastically as demand rises and falls, and the payoff of a driver decreases as more drivers provide service in her neighborhood.

Examples of such setting also arises in more traditional scenarios. For example, mobile food vendors need to decide where to locate their trucks; pastoralists need to decide where to graze their livestock; and fishermen need to decide where to fish. In these examples, the utility each of agent at their location (whether profit from hungry passers by, or food for livestock provided by the range-land, or profit from the catch) depends both on the number of other agents at the location, and on the location's stochastically varying resource level (the amount of grazeable land for pastoralists and amount of fish for fishermen).

Online communities like Reddit and Twitch also fall into this setting, in which participants choose sub-communities or channels and then derive enjoyment depending both on some underlying but transitory societal interest in the sub-community's topic of focus (the overall resource) and the number of other participants in the sub-community. 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.

Finally, another example of this setting is the research community, where academics 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 in 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 scenarios, 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 co-operation among agents at the same location, and the distribution of agents and resources across locations.

Many interesting questions regarding this setting arise. How do agents make decisions to position themselves and when should they choose to leave their current location? How do the resource dynamics or the nature of interaction among the agents at one location influence agents' decisions? How does the social welfare of all agents depend on those factors? When resource dynamics and the nature of sharing are different across locations, how does an agent choose her destination? How should a system operator communicate her information about the state of different locations to the agents in order to better position them and increase their total welfare?

Answers to these questions enable a deeper understanding of such complicated multi-agent systems. Moreover, in many spatial resource competition scenarios, there is a system operator, or a policy maker that can affect the resource dynamics, the way resource is shared or the information that is available to each agent. Examples include Uber or Lyft in the ride-sharing market, city regulators in the market of local food trucks and government officials in the labor markets. From a practical point of view, answering these questions systematically can help system operators come up with better policies and mechanisms to improve the social welfare.

However, these questions are not easy to answer. For each agent, her utility depends on the choices of other agents in the system. Meanwhile, her choices also affect the payoffs and decisions of agents at the same location, as well as agents in other locations. On the other hand, the decision an agent makes at any point in time not only affects her current payoff but also her future payoffs. Furthermore, due to the stochasticity of resources and the dynamics at other locations, it is hard for an agent to have a complete information of the state of other locations when she decides whether to explore a new location. These issues together make it challenging to characterize agents' behavior in equilibrium in such a complicated system, or design mechanisms to alter the agents' behaviors in the way a policy maker anticipated.

This dissertation aims at providing insights to understand such spatial resource competition settings and answer the forementioned questions, by building models that capture the essentials of the underlying settings, and are tractable to analysis. Below, we describe the structure of the rest of this dissertation and provide an outline of the contents of each chapter.

In Chapter 2, we provide a general model that describes spatial resource competition settings. This model comprises a group of agents and a number of locations. Each location is endowed with an independent stochastic resource process. Each agent periodically 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 this location or move. Upon moving, an agent arrives at a different location whose dynamics are independent and identical to the original location. We study the game among the agents in this model using mean field approximation [1, 52, 10, 46], which is a useful tool for analyzing games in complicated systems by approximating the effect of all other agents on any given agent with a single averaged effect, conditioning on the system is in the limit where there is a large number of agents. With mean field approximation, we provide a simpler model with a single location that can be viewed as one representative location among infinitely many locations. We study the equilibrium behavior of the agents in the mean field model, as a function of the dynamics of the stochastic resource process and the nature of the competition among co-located agents. We show that an equilibrium exists, where each agent decides whether to switch locations based only on their current location's resource level and the number of other agents there. Additionally, we show that in the common case of resource sharing scenarios where an agent's payoff is decreasing in the number of other agents at her location, equilibrium strategies obey a simple threshold structure. We exploit this structure and provide a method to approximately compute the equilibria numerically, and use these numerical techniques to study how system structure affects agents' collective ability to explore their domain to find and effectively utilize resource-rich areas. We conclude this chapter by applying our model to a case study on how ride-sharing platforms could improve their drivers' total welfare by adopting different commission rates. The major contents of this chapter appears in [78, 77].

Our model in Chapter 2 provides a way to analyze general resource competition games. However, several important issues are ignored for convenience of analysis. In Chapter 3, we extend our models to incorporate these issues and provide corresponding equilibrium analysis. Specifically, we consider three major extensions. First, we consider non-homogeneous locations and allow agents choose their destinations strategically when switching. Second, we consider scenarios where the resource at each location is not exogenous, but comes endogenously from another set of agents. In these scenarios, there is a two-sided market at each location, with buyers and sellers. The buyers are strategic and their dynamics may depend on the number of sellers at each location through price and quality of the service. This extended model aligns closer to many resource competition settings that involve matching markets. Finally, we explicitly model the service time for the two-sided market extension, as such a model would better describe many real world scenarios such as the ride-sharing setting.

Chapter 4 studies information design in spatial resource competition games. In many spatial settings, agents do not observe the state of other locations. A principal with such information would usually like to design an information sharing mechanism to communicate her extra information to the agents, in order to better position them and increase the social welfare. In this chapter, we consider a model with two locations. The agents, gathering at the first location, do not observe the resource state of the second location, while a principal has this information. Each agent needs to decide whether to move to the other location. We adopt the Bayesian persuasion framework [48, 64] and study the optimal signaling mechanism design problem for the principal. We study both private signaling mechanisms, which gives the principal more flexibility and may achieve a higher social welfare, as well as public signaling mechanisms, which is not susceptible to "information leakage": agents share their private information, leading to unanticipated behaviors. For private signaling, we show the optimal mechanism can be computed in polynomial time with respect to the number of agents. Obtaining the optimal private mechanism involves two steps: first, solve a linear program to get the marginal probability each agent should be recommended to move; second, sample the moving agents satisfying the marginal probabilities with a sequential sampling procedure. We also provide conditions on model parameters under which recommending the agents to take the social optimal strategy profile is persuasive. For public signaling, we show the sender preferred equilibrium has a simple threshold structure and the optimal public mechanism with respect to the sender preferred equilibrium can be computed in polynomial time. We support our analytical results with numerical computations that show the optimal private and public signaling mechanisms achieve substantially higher social welfare compared with no information or full information benchmarks in many settings.

CHAPTER 2

MEAN FIELD EQUILIBRIA FOR SPATIAL RESOURCE COMPETITION GAMES

2.1 Introduction

In this chapter, we develop a formal model to analyze the spatial-temporal competition among agents and their equilibrium behavior. 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. Upon 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, we study the equilibrium behavior of the agents in this system as a function of the dynamics of the spatial-temporal 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 where the resource-sharing function is non-increasing 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 2.5 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.

2.1.1 Related Work

This work contributes to the literature on mean field equilibrium [1, 45, 47, 52, 71], 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 system, has been used to study settings including industry dynamics and oligopoly models [43, 71, 72], repeated dynamics auctions [10, 46], online labor markets [5], queueing [58, 76], content sharing [55], and pedestrian motion [51], 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.

This work also contributes to the literature on spatial models of ride-sharing and crowd-sourced transportation [11, 12, 19]. In this literature, the paper most closely related to ours is [18], 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 which location to switch to.

Our model is also related to congestion games [63, 65], in which agents choose paths on which to travel, and then incur costs that depend on the number of other agents that 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 [33, 34, 54, 61]. [49] consider a set of habitats, arranged on a lattice, each containing a subpopulation of a species, and where 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.

This work can be seen as an extension of the Kolkata Paise Restaurant Problem [24], a generalization of the El Farol bar problem [6, 25]. In this game, each agent chooses (simultaneously) a restaurant to visit, and earns a reward that depends both on 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 both in the one-shot and repeated settings, with results on the limiting behavior of myopic [24] and other strategies [40], 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.

2.2 Model

The model we would study in this chapter and Chapter 3 has only a single location, representing one among infinitely many locations among which the agents move. Before we formally present this model, we first describe a model that has a finite number of locations which captures the characteristics of general resource competition scenario. We then present the single location model and describe the connection between these two models, and illustrate why we choose to study the latter.

2.2.1 The Finite System

Consider a system with a finite set of locations and a set of N agents. We denote the set of locations as $\mathcal{K}$. 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 the resource process $\{Z_t^k : t \geq 0\}$ is a finite state continuous time Markov chain, and further assume 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 for resources and less competition, as we detail below. 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 her ℓ^{th} decision epoch respectively.

We let N_t^k denote the number of agents at the location k at time t. At each decision epoch $t = \tau_i^\ell$, the 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 refer to the function F as the *resource-sharing function*. We assume that the resource-sharing function is non-negative, i.e., $F(z, n) \geq 0$ 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 $z \in \mathbb{Z}$.

Subsequent to receiving the payoff, the agent i makes the decision whether to continue staying at her 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 *which* location to move to. Nevertheless, we make this assumption as, even under this restrictive assumption, the analysis of the agent's decision problem turns out to be challenging. In Chapter 3, we provide extension and modification of this model that relaxes this assumption and align closer to practical settings.

We assume agents in the system are short-lived: after each decision epoch τ_i^ℓ, subsequent to making her decision regarding whether to stay in her current location or move to a different location, the agent i departs the system independently with probability $1 - \gamma$, never to return, and we denote as τ_i the time she leaves the system. We also assume 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 her current location k, observes the resource level Z_t^k and the number of agents N_t^k. On the other hand, the agent cannot observe the resource level and the number of agents at any other location. We assume the agents have perfect recall, and hence, at any decision epoch τ_i^ℓ, agent i bases her decision to stay or move on the entire history (namely the resource levels and the number of agents at each location she has visited) she 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 her 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. Since 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 maximize her total expected discounted payoff.

Since 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 her 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. Since 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 $\mu^i(h^i_t)$ after any history h^i_t over all aspects of the system that she is uncertain of and that influence her 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^i_t, given their belief system and given all other agents' strategies; and (2) each agent i's beliefs $\mu^i(h^i_t)$ are updated via Bayes' rule whenever possible (see [38, 37] for more details).

A PBE supposes a complex model of agent behavior. Each agent keeps track of her entire history, and maintains complex beliefs about the rest of the system. While 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, since the agent would typically only visit a small fraction of the locations. In such settings, it is more plausible that each agent would base her decision to stay or switch solely on the current state of the location she 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.

Below, we seek to uncover this intuitive behavioral model as an equilibrium in large systems by letting the number of agents and the number of location 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 location 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 her 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.

2.2.2 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 $t \geq 0$ and assume the resource process $\{Z_t^k : t \geq 0\}$ is a finite state continuous time Markov chain, with the same assumptions as in the finite system.

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 below. 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 below it is required to satisfy consistency conditions in equilibrium.

Same as in the finite system, associated with each agent i at location k is a sequence of decision epochs $\{\tau_i^l\}_{l=1}^{\infty}$ that are separated by independent and identically distributed exponential time elapses. At each decision epoch τ_i^ℓ, the 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. The same assumptions on F are made as in the finite system.

We make the same modeling assumption on agents' finite lifetimes: 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. We refer to $\gamma \in (0,1)$ as the *survival probability*, and each agent i can exist in the system for at most a random time interval distributed exponentially with rate $\lambda(1-\gamma)$.

If the agent's lifetime does not expire, then the agent i decides whether to stay at her 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 above 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 below 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 is distributed according to the stationary distribution induced by equilibrium agent behavior.

2.2.3 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 [37], where in equilibrium, each agent finds it optimal to base her decision only on the current state of the location at her decision epoch, and not on her past (although she 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 her 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 towards 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 $\mathsf{Q}^{\xi,\kappa}$:

$$\begin{aligned} \mathsf{Q}^{\xi,\kappa}\left((z,n) \to (x,m)\right) &= \mathbf{I}\{x \neq z, m = n\}\mu_{z,x} + \mathbf{I}\{x = z, m = n+1\}\kappa \\ &+ \mathbf{I}\{x = z, m = n-1\}\lambda n\left(1 - \gamma\xi(z,n)\right) \\ &- \mathbf{I}\{x = z, m = n\}\left(\sum_{y \neq z}\mu_{z,y} + \kappa + \lambda n\left(1 - \gamma\xi(z,n)\right)\right), \end{aligned} \tag{2.1}$$

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 the location k at rate κ. The third term on the right-hand side represents the departure of one of the n agents from the 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 the location k. Since there are n agents at the location, each of whose decision epoch occur 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, where 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, she receives a payoff $F(Z_i^\ell, N_i^\ell)$, and must choose whether to "stay" in location k or to "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 her 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 her decision epoch, prior to her making a decision or receiving payoffs, given resource level $z \in \mathbb{Z}$ and the 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 her decision epoch, subsequent to her making the decision to stay and conditional on her not leaving the system, given resource level z and the 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:

$$\begin{aligned} V(z, n) &= F(z, n) + \gamma \max\{V_{\mathsf{st}}(z, n), V_{\mathsf{sw}}\} \\ V_{\mathsf{st}}(z, n) &= \mathbf{E}^{\xi}[V(Z_\tau, N_\tau)|z, n], \end{aligned} \tag{2.2}$$

where $\mathbf{E}^{\xi}[\cdot|z, n]$ denotes the expectation with respect to the process defined by (2.1), 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. Upon choosing to switch and surviving, the agent receives a continuation payoff equal to V_{sw}. On the other hand, upon 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 (2.2), any optimal strategy ξ_i for 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 (2.2). 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 $V_{\mathsf{st}}(z,n) = V_{\mathsf{sw}}$.

2.2.4 Mean Field Equilibrium

With the description of the model in place, we are now ready to formally define the notion of equilibrium we focus on.

First, for any arrival rate κ and the switching payoff V_{sw}, we require the agents play a Markov perfect equilibrium at the 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, the agent i maximizes her payoff (across all possibly history-dependent strategies) by following the strategy ξ. This leads us to the following condition:

$$\xi \in \mathsf{OPT}(\xi, \kappa, V_{\mathsf{sw}}). \tag{2.3}$$

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 the 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

$$\sum_{z \in \mathbb{Z}} \sum_{n \in \mathbb{N}_0} \pi_{z,n}(\xi, \kappa) \mathsf{Q}^{\xi,\kappa}((z, n) \to (x, m)) = 0, \quad \text{for all } x \in \mathbb{Z}, m \in \mathbb{N}_0. \tag{2.4}$$

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 her 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 upon 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, where 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. Since 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:

$$V_{\mathsf{sw}} = \sum_{(z,n)\in\mathbb{S}} \pi_{z,n}(\xi,\kappa) V_{\mathsf{st}}(z,n+1). \tag{2.5}$$

The final condition we impose on the equilibrium is a requirement on the 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:

$$\sum_{(z,n)\in\mathbb{S}} n\,\pi_{z,n}(\xi,\kappa) = \beta. \tag{2.6}$$

Given these three conditions, we are now ready to define a mean-field equilibrium:

Definition 1 (Mean field equilibrium)**.** *A mean field equilibrium (MFE) consists of a strategy ξ, an arrival rate κ and a switching payoff V_{sw} satisfying* (2.3)*,* (2.5)*, and* (2.6)*.*

Note that, in comparison to a PBE, a mean field equilibrium adopts a fairly natural and a vastly simpler model of agent behavior. In a PBE of a finite model, an agent's strategy depends on the state of her current location, her history, as well as her belief about the state of all other locations. Moreover, the agent constantly updates this belief based on her observations of the arrival process at her current location. For example, if an agent sees a high volume of arrivals at her current location, her updated belief would attribute lower resource levels at other locations, thereby lowering her expected payoff for switching. Such complex considerations do not arise in an MFE, where 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, as 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 a mean field equilibrium.

2.3 Existence of a mean field equilibrium

Below, we state the main result of the paper, proving the existence of an MFE for general resource-sharing functions. Subsequently, in Section 2.4, we analyze the structure and properties of a mean field equilibrium under specific assumptions on the resource-sharing function. We have the following main theorem.

Theorem 2.3.1. *For any $\lambda > 0$, $\beta > 0$ and $\{\mu_{z,y} > 0 : z, y \in \mathbb{Z}\}$, there exists a mean field equilibrium $(\xi, \kappa, V_{\mathsf{sw}})$, 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 a mean field equilibrium is equivalent to the existence of a fixed point of $\mathcal{R}$. The latter is obtained by an application of Fan-Glicksberg fixed point theorem [2]. Here, we first sketch the steps involved,

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

and highlight the technical challenges in each of those steps. Using these intermediate results, we then provide the proof of Theorem 2.3.1. (The complete proof is provided in Appendices A.1.1-A.1.6.)

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 (2.4). 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 A.1.1.

   Denote the (unique) invariant distribution of $\mathsf{MC}(\xi, \kappa)$ by $\pi(\xi, \kappa)$. In Appendix A.1.2, by applying Berge's maximum theorem [14], 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 (2.6). 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 A.1.3.

   Let $\kappa(\xi)$ denote the unique value of the arrival rate κ for which $\pi(\xi, \kappa)$ satisfies (2.6). 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 (2.6).

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 (2.2), and let $\mathcal{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}})$ denote the corresponding continuation payoff function. Finally, we let $\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})$ denote the right-hand-side of (2.5):

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

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

   In Appendix A.1.4, 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 $\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}}) \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$. 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 A.1.5, 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 $\Upsilon = \widehat{\Pi} \times [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$.

5. 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 Fig. 2.1. In Appendix A.1.6, 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 [35].

We then obtain the following proof for the existence of a mean field equilibrium.

Proof. The steps outlined above show that $\mathcal{R}$ is an upper-hemicontinuous correspondence on a convex, compact subset Υ of a metric space, with values that are non-empty and convex. From an application of the Fan-Glicksberg fixed point theorem [2], we obtain that $\mathcal{R}$ has a fixed point, i.e., 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 (2.3) 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 (2.6). Finally, from $V_{\mathsf{sw}} = \mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})$, we obtain that (2.5) holds. From this, we conclude that $(\xi, \kappa(\xi), V_{\mathsf{sw}})$ constitutes a mean-field equilibrium.

$\square$

$$
\begin{tikzcd}
\xi \arrow[r] \arrow[rr, bend left] \arrow[dr] & \kappa(\xi) \arrow[d] & \pi(\xi, \kappa(\xi)) \arrow[d] \\
V_{\mathsf{sw}} \arrow[r] \arrow[dr, Rightarrow] & \mathcal{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}}) \arrow[r] \arrow[d, Rightarrow] & \mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}}) \\
 & \mathcal{X}(\xi, V_{\mathsf{sw}}) &
\end{tikzcd}
$$

Figure 2.1: Illustration of the correspondence $\mathcal{R}(\xi, V_{\mathsf{sw}}) = \mathcal{X}(\xi, V_{\mathsf{sw}}) \times \{\mathcal{V}_{\mathsf{sw}}(\xi, V_{\mathsf{sw}})\}$. (Here single arrows denote functions, and double arrows denote correspondences.)

2.4 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, where the resource-sharing function is non-increasing 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 2.5 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* $F(z, n+1) \leq$

$F(z, n)$ 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 2.5 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 where 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 $\lambda n F(z, n) = \lambda f(z)$, a quantity independent of n. We refer to the quantity $W(z, n) \triangleq \lambda n F(z, n)$ as *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. While 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, where 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 $n \geq 0$. In particular, under strategy $\xi^{\mathbf{x}}$, an agent, at her decision epoch, will stay at her 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 2.4.1. *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 the theorem makes essential use of the following lemma, which states that with decreasing resource-sharing functions, the continuation values are non-increasing.

Lemma 2.4.1. *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 non-increasing in n.*

The proof of the lemma, provided in Appendix A.1.7, shows that the decision problem $\mathsf{DEC}(\xi, \kappa, V_{\mathsf{sw}})$ has a dynamic program that satisfies closed convex cone properties defined in [66]. With the lemma in place, the proof of Theorem 2.4.1 follows from minor modifications of the argument in the proof of Theorem 2.3.1, and is omitted.

2.5 Computation of MFE and Numerical Equilibrium Analysis

The implications of Theorem 2.4.1 are of substantial practical importance: when the resource-sharing function is decreasing, the equilibrium behavior of the agents can be fully described by $|\mathbb{Z}|$ non-negative 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 below.

2.5.1 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 $\mathcal{R}(\mathbf{x}, V_{\mathsf{sw}}) = \mathcal{X}(\mathbf{x}, V_{\mathsf{sw}}) \times \mathcal{V}_{\mathsf{arr}}(\mathbf{x}, V_{\mathsf{sw}})$. For any $(\mathbf{x}, V_{\mathsf{sw}})$, we define the distance metric $\mathsf{dist}_{\mathcal{R}}$ as follows:

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

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 below.

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 (2.4).

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
$$\left| \sum_{z \in \mathbb{Z}} \sum_{n=0}^{L} n\pi_{z,n} - \beta \right| \leq \epsilon_1,$$
where $\pi = \pi_L(\mathbf{x}, \kappa_L(\mathbf{x}))$ and $\epsilon_1 > 0$ denotes the tolerance level within which we seek to satisfy (2.6).

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 per-

form value iteration to compute approximate value functions $\mathcal{V}^{\epsilon}_{\mathsf{st}}(\mathbf{x}, V_{\mathsf{sw}})$ and $\mathcal{V}^{\epsilon}_{\mathsf{arr}}(\mathbf{x}, V_{\mathsf{sw}})$, 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 $\mathcal{V}_{\mathsf{arr}}$ and $\mathcal{X}$ in the definition of $\mathsf{dist}_{\mathcal{R}}$ with $\mathcal{V}^{\epsilon}_{\mathsf{arr}}$ 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 [62] 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 0 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 to find an approximate equilibrium on terminating, in all our computations in section 2.5.2, 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 non-uniqueness may arise, for example, when the resource-sharing function is multimodal, as 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.

2.5.2 Comparative Statics

In this section, we present the results of our numerical investigations of the agents' behavior in a mean field equilibrium using the computational approach described in the preceding section. We study the setting where $\mathbb{Z} = \{0, 1\}$, with transitions rates $\mu_{0,1} = \mu_{1,0} = \mu$. As our model is invariant to proportional scaling of the transition rate μ and the agents' inter-epoch 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 $z = 1$) while others do not ($z = 0$), and the single-location welfare function is increasing for $\alpha = 0.5$, constant for $\alpha = 1$, and decreasing for

[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 2.5.2. An extensive numerical investigation supports this conjecture, but we do not have a formal proof.

$\alpha = 1.5$ 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 $W(z, n) = \lambda n F(z, n)$. Since 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. Since 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 2.2 shows how the equilibrium thresholds and the welfare per location vary as the resource process changes more frequently, i.e, 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. Since the resource level changes slowly, an agent in a location with resource is willing to suffer significant competition (in the form of other agents) before choosing to switch her location. 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. On the other hand, 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 2.2 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, i.e., as μ increases. In 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 level. On the other hand, 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, since 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. On the other hand, when $\alpha > 1$, the former distribution incurs lower welfare per location due to severe competition among the agents at the location with resource. (When $\alpha = 1$, the distribution of the agents between locations

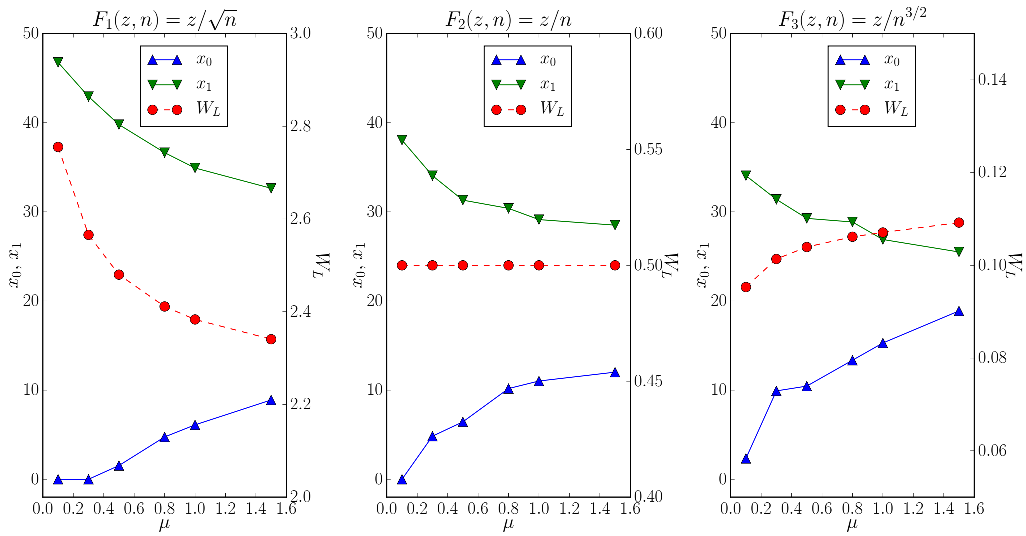

Figure 2.2: Equilibrium thresholds and welfare under different resource transition rates μ, with agent density fixed at $\beta = 20$.

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 2.3 shows equilibrium properties as a function of the 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, since there are more agents to compete against. Thus, when the current location has resource, 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 2.2, 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, since 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 2.2 demonstrates this by showing that agents' strategies change as the resource transition rate μ changes. In contrast, since 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 $\lambda n F(z, n)$ depends non-trivially 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.

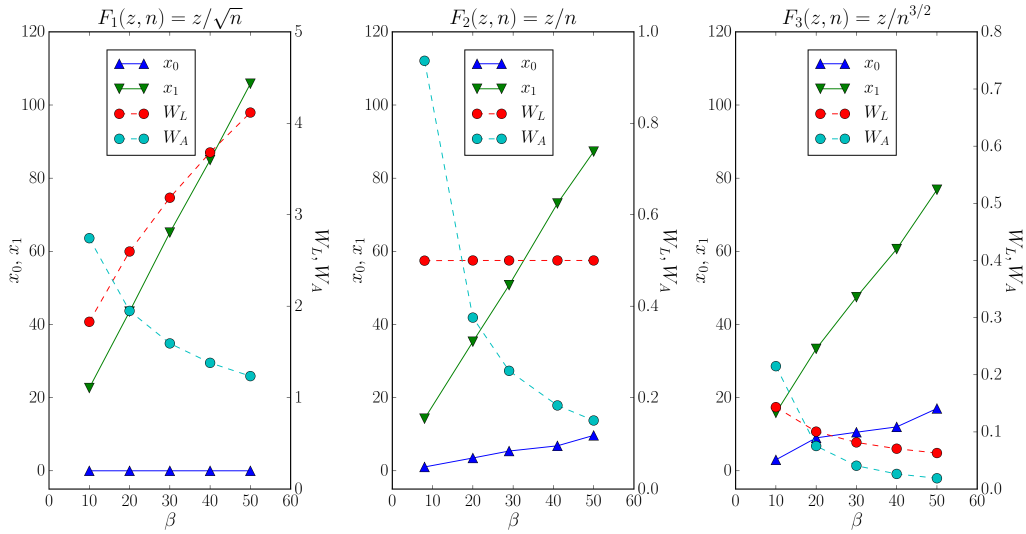

Figure 2.3: Equilibrium thresholds and welfare under different agent densities β. W_A is multiplied by 15 for all values. Note the resource transition rate is given by $\mu = 0.25$.

2.5.3 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 a driver makes this decision based on the trip demand in his current neighborhood as well as his estimate of the number of competing drivers in the same neighborhood.

Below 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 dataset [29] 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 4800 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 0 denoting the typical resource state, and 1 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 non-rainy days using historical weather data from Manhattan [42]. 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 4 days, for a duration of about 2 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 [23, 27], as 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 non-rainy days in Manhattan from [29], 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, since our data comes 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 $\lambda = 1$ 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 parameter values described above), we numerically compute the resulting mean field equilibrium in our model, and the driver and platform revenues in the computed equilibrium. We share these results in Table 2.1. 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 upon switching, which influences her switching decisions. Our model enables us to include this endogenous effect of the driver distribution on the drivers' switch-

$c(0)$	$c(1)$	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%

Table 2.1: 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.

ing decisions in evaluating different commission rates. Without the dynamics (and the tractable equilibrium concept of a mean field equilibrium), such effects would be hard to incorporate in a static analysis, rendering it incomplete.

2.6 Conclusion

In this chapter, our analysis establishes 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 spatial-temporal models under simple yet optimal models of agent behavior.

The spatial resource competition model in this chapter and analysis raise many topics for future research, which we will discuss in details in Chapter 5. On the modeling front, we have assumed that each location is homogeneous. In particular, we assume the resource process is distributed independently and identically across different locations. One consequence of this homogeneity is agents do not choose their destination when they switch. One natural extension of our model is to incorporate location heterogeneity and to let agents choose where to go strategically when switching. Such a model would better align the settings we study. In the next chapter, we will provide an extended model and corresponding equilibrium analysis for system with non-homogeneous locations. On the other hand, we assumed the resource process at a location is exogenous specified, and in next chapter we will consider a variation our model that allows the "resource" being a group of strategic customers whose dynamics may endogenously depend on the number of agents at the location.

On the other hand, our work also sets the stage for analyzing engineering interventions and their economic impact in settings we study. In Section 2.5.3, we provide an example of one such intervention that 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. More extensive studies could be undertaken that utilizes this model to help policy makers design mechanisms.

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. In Chapter 4, we would study the information mechanism design problem for spatial resource competition settings.

CHAPTER 3

HETEROGENEOUS LOCATIONS AND TWO-SIDED MARKETS

In Chapter 2, we studied a model aimed to describe spatial resource competition settings. We assumed the resource process at each location is independent and locations are homogeneous hence agents' decision about which location to move to was ignored. In many real-world spatial resource competition scenarios, the resource dynamics across different locations can be significantly different, and the agents usually take such difference into consideration when choosing where to locate themselves. For example, in the ride-sharing setting, residential neighborhoods and business districts have different demand patterns: the business districts are usually much more busier in peak hours and have a larger demand variations across different times of day. Drivers usually take the different demand pattern of each location into consideration when they choose where to provide service. Furthermore, the way resource is shared among agents in one location may also be different across different locations.

On the other hand, the resource is an exogenous Markov process in our model. In many real-world scenarios, however, the resource may be provided endogenously by a different set of agents. For example, in ride-sharing platforms such as Uber and Lyft, the resource is the demand of rides provided by riders; in crowd-sourced labor markets such as TaskRabbit, the resource is the demand for task fulfillments provided by consumers. In these scenarios, each location has a two-sided market with buyers and sellers. Applying our model to these settings, the exogenous resources assumption essentially implies that each buyer's dynamic is as simple as arriving at the market, making a purchase and then leaving the market.

This independent and identical Markovian resource process assumption ignores two important factors. First, in a two sided market, the buyers are usually strategic and a buyer's purchase decision depends on the price and the quality of the service. The price and the amount of purchase in turn jointly affect the payoff of sellers. An effective model for those scenarios should allow different pricing mechanisms to affect the resource level and the resource sharing dynamics. Secondly, when we consider the social welfare of a two-sided market, we should take both sellers' and buyers' welfare into consideration. Our current model only included sellers' welfare and it is not clear how to characterize buyers' welfare in this model. To better align to two-sided market settings, a model should allow us to conveniently characterize buyers' welfare as well as sellers'.

Finally, in our model we assumed the resource is derived instantly. However, in scenarios such as ride-sharing, each service would take a period of time to complete. In this chapter, we extend our model to take those important factors into consideration.

3.1 A Model with Heterogeneous Locations and Equilibrium Analysis

In this section, we present a mean field model where each location is endowed with a type drawn from a finite set of types, and the type determines the resource process and the sharing function at this location, thus incorporating location heterogeneity into the model.

3.1.1 Model

We begin with a description of the finite system, which has a set of K locations and N agents. Each location k now has a type θ_k coming from a finite set of types Θ. We assume the type is drawn according to a distribution over types, and let p_θ denote the probability of type θ. The resource process $\{Z_t^k : t \geq 0\}$ at location k evolves as a continuous time Markov chain whose state space $\mathbb{Z}^{\theta_k}$ and transition $\{\mu_{zy}^{\theta_k} : z, y \in \mathbb{Z}^{\theta_k}\}$ depends on the type of the location (without losing of generality, we let $\mathbb{Z} = \cup_{\theta \in \Theta} Z^{\theta}$). The resource-sharing function F_{θ_k} also depends on the type of the location, and for each $\theta \in \Theta$, F_θ follows the same assumptions of resource-sharing functions we made in Chapter 2.

The assumptions on agents' decision epochs and exponential lifetimes remain the same as in Chapter 2. However, unlike the homogeneous locations model, at the end of a decision epoch of an agent i, if the agent chooses to switch to a new location, she first chooses a type $\theta \in \Theta$, and then uniformly selects a location among all type-θ locations.

Similar to Chapter 2, we let the number of locations and number of agents both increase to infinity, while the agent density $\beta = N/K$ remain fixed, obtaining a limiting infinite system. We also assume the type of each location is still drawn from the distribution $\{p_\theta\}_{\theta\in\Theta}$ as the system grows to infinity. We assume the dynamics at each location decouple in the limiting system, and study a mean field model where there is a single location for each type $\theta \in \Theta$, representing all type-θ locations in the limiting system.

In the mean field model, for each type θ, agents arrive at the type-θ location following a Poisson process with rate κ_θ, and this arrival flow models arrivals from other locations as well as outside in the limiting infinite system. For each agent, at each of her decision epoch, she observes a payoff V^θ_{arr} for arriving at the type-θ location for each type θ. She then decides whether to switch to another location. On choosing to switch, she obtains a one time payoff $V_{\mathsf{sw}} = \max_{\theta\in\Theta} V^\theta_{\mathsf{arr}}$ and exits the system. V_{sw} models the agent's expected total future payoff on choosing a location with type $\theta \in \arg\max_{\theta\in\Theta} V^\theta_{\mathsf{arr}}$ and continues to stay in the system.

3.1.2 Equilibrium Analysis

We aim to characterize a symmetric equilibrium where all agents follow the same strategy. We first characterize, for each type $\theta \in \Theta$, the Markov perfect equilibrium of a type-θ location with given arrival rate κ_θ and switching payoff V_{sw}, and then provide consistency conditions for κ_θ and V_{sw} and define the mean field equilibrium.

Let Z_t be the resource level and N_t be the number of agents at the location. Consider the dynamics of Z_t and N_t at the type-θ location where all agents adopt the same Markovian strategy $\xi_\theta : \mathbb{Z}\times\mathbb{N}_0 \to \Delta([0,1])$, where $\xi_\theta(z,n)$ is the probability an agent chooses to stay at the location when the resource level is z and number of agents is n. Along with the arrival rate κ_θ, the process $\{(Z_t, N_t) : t \geq 0\}$ evolves as a continuous time Markov chain, denoted as $\mathsf{MC}(\xi_\theta, \kappa_\theta)$, with transition kernel $\mathsf{Q}_\theta^{\xi_\theta,\kappa_\theta}$ given as:

$$
\begin{aligned}
\mathsf{Q}_\theta^{\xi_\theta,\kappa_\theta}((z,n) \to (x,m)) = &\mathbf{I}\{x \neq z, m = n\}\mu^\theta_{z,x} + \mathbf{I}\{x = z, m = n+1\}\kappa_\theta \\
&+ \mathbf{I}\{x = z, m = n-1\}\lambda n\,(1 - \gamma\xi_\theta(z,n)) \\
&- \mathbf{I}\{x = z, m = n\}\left(\sum_{y\neq z}\mu^\theta_{z,y} + \kappa_\theta + \lambda n\,(1 - \gamma\xi_\theta(z,n))\right).
\end{aligned}
\tag{3.1}
$$

The decision problem faced by a particular agent i when all other agents follow strategy ξ_θ is an optimal stopping problem and we denote it as $\tilde{\mathsf{DEC}}(\xi_\theta, \kappa_\theta, V_{\mathsf{sw}})$.

Let $V^\theta(z,n)$ denote the value function of agent i at her decision epoch prior to making a decision and receiving payoffs, and let $V^\theta_{\mathsf{st}}(z,n)$ denote the continuation payoff on choosing to stay, given resource level z and number of other agents n. The problem $\tilde{\mathsf{DEC}}(\xi_\theta, \kappa_\theta, V_{\mathsf{sw}})$ is characterized by the following Bellman Equation:

$$
\begin{aligned}
V^\theta(z,n) &= F_\theta(z,n) + \gamma \max\{V^\theta_{\mathsf{st}}(z,n), V_{\mathsf{sw}}\} \\
V^\theta_{\mathsf{st}}(z,n) &= \mathbf{E}[V^\theta(Z_\tau, N_\tau) | (Z_0, N_0) = (z,n); \xi_\theta, \kappa_\theta],
\end{aligned}
\tag{3.2}
$$

where γ and τ has the same definition as in Chapter 2. We let $\tilde{\mathsf{OPT}}(\xi_\theta, \kappa_\theta, V_{\mathsf{sw}})$ be the set of best responses of the decision problem $\tilde{\mathsf{DEC}}(\xi_\theta, \kappa_\theta, V_{\mathsf{sw}})$. The Markov perfect equilibrium at the location given κ_θ and V_{sw} is then a symmetric Markov strategy played by all agents that satisfies

$$\xi_\theta \in \tilde{\mathsf{OPT}}(\xi_\theta, \kappa_\theta, V_{\mathsf{sw}}). \tag{3.3}$$

For a given κ_θ and $\{V^\theta_{\mathsf{arr}}\}_{\theta\in\Theta}$, assume a Markov perfect equilibrium strategy ξ_θ is being played. We denote the invariant distribution of the corresponding Markov process $\mathsf{MC}(\xi_\theta, \kappa_\theta)$ as $\pi^\theta(\xi_\theta, \kappa_\theta)$, satisfying

$$\sum_{z\in\mathbb{Z}}\sum_{n\in\mathbb{N}_0} \pi^\theta_{z,n}(\xi_\theta, \kappa_\theta)\mathsf{Q}^{\xi_\theta,\kappa_\theta}_\theta((z,n)\to(x,m)) = 0, \quad \text{for all } x\in\mathbb{Z}, m\in\mathbb{N}_0. \tag{3.4}$$

Let β_θ be the number of agents in steady state of a type-k location (agent density), which is given as

$$\sum_{z\in\mathbb{Z},n\in\mathbb{N}_0} n\pi^\theta_{z,n}(\xi_\theta, \kappa_\theta) = \beta_\theta. \tag{3.5}$$

We assumed the agent density remained fixed as β, and locations' types still follow the distribution $\{p_\theta\}_{\theta\in\Theta}$ as we obtaining the limiting system. Therefore, we require the overall agent density equals the expected agent density with respect to the location distribution:

$$\sum_{\theta\in\Theta}\beta_\theta p_\theta = \beta. \tag{3.6}$$

For each type $\theta \in \Theta$, the expected payoff of arriving at the type-θ location V^θ_{arr} should be equal to the expected continuation payoff of staying at that location under the steady state:

$$V_{\mathsf{arr}}^{\theta} = \sum_{z \in \mathbb{Z}, n \in \mathbb{N}_0} \pi_{z,n}^{\theta}(\kappa_\theta, \xi_\theta) V_{\mathsf{st}}^{\theta}(z, n+1). \tag{3.7}$$

Since locations are not homogeneous, the arriving payoff for each type of location may be different. Let $\bar{\Theta} \triangleq \arg\max_{\theta \in \Theta} V_{\mathsf{arr}}^{\theta}$ be the set of types whose arriving payoff is maximal among all types. For any type $\theta \in \Theta \backslash \bar{\Theta}$, agents in a system with a large but finite number of locations would not choose type-θ locations as destination on choosing to switch, and we call such locations as “abandoned locations”. Arrivals to an abandoned location comprises of only the new agents from outside the system that lands at this location. In the finite system, the rate of agents arriving from outside the system is equal to the rate of agents leaving the system, and incoming agents choose their destination uniformly, hence the arrival rate of new agents to a particular location is given as $\beta\lambda(1-\gamma)$. We then require for all $\theta \in \Theta \backslash \bar{\Theta}$, the arrival rate must satisfy $\kappa_\theta = \beta\lambda(1-\gamma)$. On the other hand, for locations that are not abandoned, the arriving payoff should equal the one-time switching payoff V_{sw}, since V_{sw} models the total continuation payoff of switching to a non-abandoned location. Summarizing the above argument, we require

$$\begin{aligned} (V_{\mathsf{arr}}^{\theta} - V_{\mathsf{sw}})(\kappa_\theta - \beta\lambda(1-\gamma)) &= 0, \quad \forall \theta \in \Theta, \\ \kappa_\theta \geq \beta\lambda(1-\gamma), &\quad \forall \theta \in \Theta, \end{aligned} \tag{3.8}$$

and

$$V_{\mathsf{sw}} = \max_{\theta \in \Theta} V_{\mathsf{arr}}^{\theta}. \tag{3.9}$$

We then define the mean field equilibrium of a heterogeneous locations model as follows:

Definition 3. *A mean field equilibrium is characterized by* $(\{\xi_\theta : \theta \in \Theta\}, \{\kappa_\theta : \theta \in \Theta\}, \{\beta_\theta : \theta \in \Theta\}, \{V^\theta_{\mathsf{arr}} : \theta \in \Theta\}, V_{\mathsf{sw}})$ *that satisfies (3.9), (3.3), (3.5), (3.6), (3.7) and (3.8).*

Our major results of the homogeneous location model carries through to this extended model.

Proposition 1. *If for all type* $\theta \in \Theta$*, the resource sharing function* $F_\theta(z, n)$ *is non-increasing in* n *for all* $z \in \mathbb{Z}$*, then for any type* $\theta \in \Theta$*, let* ξ_θ *be a Markovian strategy,* $\kappa_\theta > 0$ *and* $V_{\mathsf{sw}} > 0$*, the continuation payoff* $V^\theta_{\mathsf{st}}(z, n)$ *for the decision problem* $\tilde{\mathsf{DEC}}(\xi_\theta, \kappa_\theta, V_{\mathsf{sw}})$ *is non-increasing in* n *for all* $z \in \mathbb{Z}$*.*

The proof of Proposition 1 is similar to that of Lemma 2.4.1 and we omit it. Following this proposition, the threshold structure of the equilibrium strategy is retained.

Theorem 3.1.1. *If for all type* $\theta \in \Theta$*,* $F_\theta(z, n)$ *is non-increasing in* n *for all* $z \in \mathbb{Z}$*, given any* $\lambda > 0, \beta > 0, \{\mu^\theta_{z,y} > 0 : z, y \in \mathbb{Z}\}$ *for all* $\theta \in \Theta$ *and type distribution* $\{p_\theta : \theta \in \Theta\} \in \Delta(\Theta)$*, there exists a mean field equilibrium.*

The proof of this theorem adopts the same general approach as that of the homogeneous locations model, which constructs a correspondence whose fixed point characterizes an equilibrium, and show the existence of a fixed point using Fan-Glicksberg fixed point theorem. The construction of the correspondence is different from that of the previous chapter, so we provide a proof sketch in Appendix B, emphasizing these differences.

3.1.3 Discussion

This model and the corresponding equilibrium analysis can be further extended easily to allow agents from outside the system choose the type of their destination to join not uniformly but according to a distribution over location types. We omit the details of this extension.

Similar to the previous chapter, approximate equilibria of this model can be computed via finding a fixed point of the correpondence using neighborhood search methods. One could study numerically how different resource dynamics and resource sharing functions at each location type affect the agent density and the social welfare at each type of location. Such study would help, for example, ridesharing platforms such as Uber and Lyft better design their spatial pricing schemes.

3.2 A Model for Two-Sided Markets and Equilibrium Analysis

In this section, we consider spatial resource sharing settings where there is a two-sided market between buyers and sellers at each location. In this scenario, the term "resource" is from the viewpoint of the sellers. The same service or goods is provided at all locations. At each location, buyers come to the market, observe the price of the service and make a one-time purchase decision based on the price and the quality of the service, and then leave the market. The price is determined endogenously by supply and demand, or set by a system operator. If a buyer decides to purchase, one of the sellers is chosen uniformly to provide the service and is paid the price. Motivated by the ride-sharing setting, we assume the service is not completed instantly, and the seller may be relocated to another location after the service (and these assumptions can be simplified to allow instantaneous service without relocation, and we omit the details). The sellers can move across different locations to provide their service and they would always prefer locations with higher volume of buyers and lower competition from other sellers.

We provide a formal model for the above setting in Section 3.2.1 and the corresponding equilibrium analysis in Section 3.2.2.

3.2.1 Model

Consider a model with K locations and N sellers. At each location there is a two-sided market with buyers and sellers. Buyers arrive at the location following a Poisson process with rate ϕ. Each buyer i's valuation of the service U_i is drawn independently from a known distribution F. Upon arriving at the location, the buyer is quoted a price $p(n, m) \geq 0$ for the service, where n is the number of sellers at the location ready to provide service, and m is the number of sellers at the location who are currently serving other buyers, at the time the buyer enters the market. The price can be interpreted as the market clearing price at this location. We assume for any $m \in \mathbb{N}_0$, $\lim_{n\to\infty} p(n, m) = 0$ and for any $n \in \mathbb{N}_0$, $\lim_{m\to\infty} p(n, m) =$

0.

The total utility of a particular buyer i depends on her valuation $U_i \sim F$ and a disutility $f(n) \geq 0$, which models the case where the service quality depends on the number of available sellers n. The buyer would purchase the service if her total utility is higher than the price. Let $q(n, m)$ denote the probability of purchasing. $q(n, m)$ is given as

$$q(n, m) = \mathbf{P}(U - f(n) \geq p(n, m)) = 1 - F(p(n, m) + f(n)). \tag{3.10}$$

We require $f(0) = +\infty$ so when no seller is available at the location, the buyer would not purchase. We also assume the price p and disutility f satisfies that there exists $n_0, m_0 \in \mathbb{N}_0$ such that $q(n_0, m_0) > 0$ to avoid triviality.

If the buyer decides to purchase, among the n available sellers at the location, each one is selected uniformly at random to serve the buyer. The selected seller would get paid the price $p(n, m)$. The service time is exponentially distributed with rate $\tau > 0$. Upon completing the service, the selected seller would stay in the current location with probability $\eta \in [0, 1]$ and be relocated uniformly at random to another location with probability $1 - \eta$. Each seller also decides periodically whether she should switch to another location to provide service. Specifically, each seller has an associated sequence of decision epochs, with $\exp(\lambda)$ intervals. At each decision epoch, she exits the system permanently with probability γ, and if not exiting, she decides whether to stay in the current location, or switch to another location chosen uniformly at random.

The mean field model has a single representative location. Sellers are idle and available to provide service when not assigned to serve a buyer, and if gets selected to serve, the seller becomes busy for an $\exp(\tau)$ period. Upon service completion, with probability η the agent becomes idle and stays in the location; and with probability $1-\eta$, she switches to other locations, receiving a one-time payoff V_{sw}, accounting for her total future payoff of being relocated to another location. On the other hand, if an agent decides to switch in a certain decision epoch, she also receives V_{sw} and switches.

New sellers arrive at the location following a Poisson process, accounting for sellers completing service and joining this location, switching sellers from other locations, as well as sellers from outside landing at this location. We assume the average number of sellers in steady state is still exogenously given as β.

3.2.2 Equilibrium Analysis

We first identify the Markov perfect equilibrium for the game between sellers at the location, given an arrival rate of sellers κ and switching payoff V_{sw}. In a Markovian equilibrium, each seller's decision on whether to switch only depends on the number of busy and idle sellers at the location at the decision time, and a Markovian strategy is a function $\xi : \mathbb{N}_0^2 \to [0, 1]$ where $\xi(n, m)$ denotes the probability to switch when there are n idle and m busy agents. Let N_t and M_t be the number of idle and busy sellers respectively at time t. If all sellers follow a Markovian strategy, (N_t, M_t) is Markov process, and we denote it as $\mathsf{MC}(\xi, \kappa)$, with the transition kernel $Q^{\kappa,\xi}$ given as

$$Q^{\xi,\kappa}((n,m) \to (k,l)) = \begin{cases} n\lambda(1-\gamma\xi(n,m)), & k = n-1, l = m; \\ \kappa, & k = n+1, l = m; \\ \tau(1-\eta), & k = n, l = m-1; \\ \tau\eta, & k = n+1, l = m-1; \\ \phi q(n,m), & k = n-1, l = m+1; \\ -n\lambda(1-\gamma\xi(n,m)) - \kappa - \tau - \phi q(n), & k = n, l = m. \end{cases} \tag{3.11}$$

Here the first case corresponds to the event a seller's lifetime expires and exits the system or she decides to switch in a decision epoch; the second case corresponds to arrival of an agent to this location, and upon arrival this agent becomes idle and waits for service assignment; the third case corresponds to a seller completes service and switches away from the location; the fourth case corresponds to a service completion and the underlying seller stays in the location and becomes idle again; the fifth case corresponds to an arrival of a buyer and she decides to request a service. We let $\pi(\xi, \kappa)$ be the steady state distribution of $\mathsf{MC}(\xi, \kappa)$.

On the other hand, for each seller, at each of her decision epoch, assuming all other sellers adopting a Markovian strategy ξ, the agent faces a decision problem $\hat{\mathsf{DEC}}(\xi, \kappa, V_{\mathsf{sw}})$. Let $V(n,m)$ denote her value function of this decision problem, and $V_{\mathsf{st}}(n,m)$ denote her expected continuation payoff of staying at the location, when there are n idle sellers and m busy sellers at the location at this decision epoch. $V(n,m)$ is characterized by the Bellman equation:

$$\begin{aligned} V(n,m) =& \gamma \max\{V_{\mathsf{st}}(n,m), V_{\mathsf{sw}}\}, \\ V_{\mathsf{st}}(n,m) =& \mathbf{E}[V(N_\delta, M_\delta) | (N_0, M_0) = (n,m); \delta \sim \exp(\lambda); \kappa, \xi], \end{aligned} \tag{3.12}$$

where given $(N_0, M_0) = (n,m)$, the distribution of (N_δ, M_δ) is determined by $Q^{\xi,\kappa}$.

Let $\hat{\mathsf{OPT}}(\xi, \kappa, V_{\mathsf{sw}})$ be the set of best responses of $\hat{\mathsf{DEC}}(\xi, \kappa, V_{\mathsf{sw}})$. The Markov perfect equilibrium strategy is a Markovian strategy ξ such that

$$\xi \in \hat{\mathsf{OPT}}(\xi, \kappa, V_{\mathsf{sw}}). \tag{3.13}$$

For consistency in the equilibrium, we require in steady state the expected number of agents at the location equals β:

$$\sum_{n,m \in \mathbb{N}_0} (n+m)\pi(n,m;\xi,\kappa) = \beta. \tag{3.14}$$

We also require the expected continuation payoff for staying at the location, which should equal the payoff of arriving at this location in a finite system with large number of locations, to be consistent with the switching payoff:

$$V_{\mathsf{sw}} = \sum_{n,m \in \mathbb{N}_0} V_{\mathsf{st}}(n+1,m)\pi(n,m;\xi,\kappa). \tag{3.15}$$

A mean field equilibrium for this model is a Markovian strategy ξ, an arrival rate κ and a switching payoff V_{sw} that satisfies (3.12), (3.13) and (3.15), and our existence result carries through:

Theorem 3.2.1. *For any valuation distribution F, $\lambda > 0$, $\phi > 0$, $\tau \geq 0$, $\beta > 0$, there exists a mean field equilibrium.*

The proof of this result is similar to the proof of the existence of mean field equilibria for the original model, by constructing a correspondence whose fixed point corresponds to an equilibrium, and show the existence of a fixed point with Fan-Glicksberg theorem. We omit the details.

On the other hand, in many two-sided markets, the market clearing price decreases as supply increases. If we assume for any $m \in \mathbb{N}_0$, $p(n, m)$ is non-increasing in n and for any $n \in \mathbb{N}_0$, $p(n, m)$ is non-increasing in m, with a similar argument as that in the proof of Lemma 2.4.1, it can be shown the agents' equilibrium strategy has a threshold structure for both n and m. Specifically, the following proposition holds

Proposition 2. *If for for any $m \in \mathbb{N}_0$, $p(n, m)$ is non-increasing in n and for any $n \in \mathbb{N}_0$, $p(n, m)$ is non-increasing in m, then the continuation payoff $V^{\theta}_{\mathsf{st}}(n, m)$ for the decision problem $\hat{\mathsf{DEC}}(\xi, \kappa, V_{\mathsf{sw}})$ is non-increasing in n for all $m \in \mathbb{N}_0$ and non-increasing in m for all $n \in \mathbb{N}_0$.*

3.2.3 Discussion

With Proposition 2, the dimension of the equilibrium strategy is reduced and we can use similar methodology as in Section 2.5 to approximately compute a mean field equilibrium for given model parameters. We can further investigate numerically, for example, how different pricing schemes affect the equilibrium of the market, and how the equilibrium strategies change under different market dynamics: e.g., the buyer arrival rate, the service time and the relocation probability. Such analysis enables this model to be applied in practical scenarios.

An important practical application of this model is to study how platforms or system operators could set the prices $p(n, m)$ in order to improve the social welfare or other metrics of the system. In contrast to our previous model, this two-sided model allows us to characterize both the sellers' and the buyers' welfare naturally. When there are n idle sellers and m busy sellers at the location, upon a service request, the buyer's welfare is given as $U - p(n, m) - f(n)$, and the seller's utility is the price $p(n, m)$. Therefore, let $W(n, m)$ be the social welfare including the welfare of both the sellers and the buyers, when there are n idle sellers and m busy sellers, we have

$$W(n, m) = \phi q(n, m)\Big(\mathbf{E}_F[U|U - f(n) \geq p(n, m)] - f(n)\Big),$$

and the social welfare per location in steady state is given as:

$$W_L = \mathbf{E}_\pi[W(N, M)] = \sum_{n,m} \phi q(n, m)\Big(\mathbf{E}_F[U|U - f(n) \geq p(n, m)] - f(n)\Big)\pi(n, m).$$

This model can be extended to align with more general settings. For example, in many real-world scenarios, the utility of the buyer and buyers' arrival rate ϕ may be different when exogenous environment changes, e.g., in the ride-sharing setting, the ride demands and riders' valuation of a ride may be different under different weather conditions. This model can be extended to incorporate this scenario. We may assume at each location, at time t, a state of nature Z_t determines the buyers' arrival rate ϕ or buyers' valuation distribution F, and Z_t is a Markovian process. In this case, the price of the service may also depend on the state of nature z.

We can also include location heterogeneity in this model by letting the arrival rate ϕ, the price function p and the buyers' valuation distribution F be different for different type of locations, and obtain a mean field model similar to what we did in Section 3.1.1.

CHAPTER 4

INFORMATION DESIGN IN SPATIAL RESOURCE COMPETITION GAMES

4.1 Introduction

We study information design in the spatial resource competition settings, where a group of agents migrate across a network of locations, competing for stochastic time-varying resources at each location. This setting characterizes many real world scenarios: on crowd-sourced transportation platforms, drivers migrate across different neighborhoods of a city, competing for ride demands; in unskilled labor markets, workers migrate across different cities for more job opportunities; in nomadic animal husbandry, pastoralists migrate across different range lands, competing for water and pastures, etc.

In these scenarios, information about other locations largely affects an agent's decision about whether to leave her current location, and where to explore. However, such information is usually limited or even unavailable to the agents. Meanwhile, in many cases, there is a principal who has access to more information than any individual agent, e.g., platforms such as Uber and Lyft in the ride sharing market, or the government agencies in labor markets, or non-profits in the case of pastoralists. Using this information, the principal can influence the decisions of the migrating agents, better locate them and improve the total social welfare. For example, Uber and Lyft show the demand trend at different neighborhoods as a heatmap to the drivers [69] in order to locate them to areas with higher demands, and government agencies provide information about employment and job openings in various sectors.

However, sharing information may not always bring higher welfare, if the information is not released in a careful way: if a sports match or a concert at a stadium is about to finish and Uber informs every available driver in the nearby area, too many drivers may flood to this area, with many of them failing to find a customer due to over-supply. In such cases, it might be wiser for the platform to only inform a particular group of drivers to better match supply and demand. How to choose such a group? More generally, spatial resource competition scenarios typically exhibit negative externalities. In presence of such negative externalities, how should a principal effectively communicate her information to the agents in order to better position them?

Answering this question for a large network of locations is challenging: the effect of a signal that attracts a group of agents to a particular location may percolate across the entire network through agents' subsequent migration. On the other hand, signals can take very complicated forms in large networks, depending on many factors, including the number of locations and agents, the dynamics at each location, and the agents' belief about the state of the system and other agents' strategies. With the goal of obtaining insights to this spatial information design problem while retaining a tractable analysis, we consider a two-location model, which serves as a foundation for more complicated analysis.

4.1.1 Overview of Model and Main Results

We consider a model with 2 locations, ℓ_0, ℓ_1, and N agents. Initially, all the agent are at ℓ_0, and must decide whether move to ℓ_1, which has a stochastic resource. The utility each agent receives upon moving to ℓ_1 depends on the stochastic resource level at ℓ_1, as well as how many other agents move there. The agents do not know the resource level at ℓ_1 while a principal can observe it. The principal would like to design a mechanism to share this information to the agents in order to attract a proper number of agents to move to ℓ_1.

We adopt the framework of Bayesian persuasion [48, 64] to study this information design problem. The principal's goal is to choose a signaling mechanism that maximizes the expected social welfare, i.e., the total expected utility of the agents. In this work, we consider both private signaling mechanisms, where the principal sends personalized signals to each agent privately, as well as public signaling mechanisms, where the principal sends the same information to all the agents.

The standard approach using a revelation-principle style argument [16] to find the optimal private signaling mechanism leads to a linear program in 2^N variables, rendering the computation challenging. Instead, our first main result in Section 4.3 characterizes a computationally efficient two-step approach to find the optimal private signaling mechanism. First, we perform a change-of-variables and instead of the signaling mechanism, we focus on the marginal probabilities p_{ik} that an agent i is recommended to move to ℓ_1 along with $k-1$ other agents, for each i and k. We show that the marginal probabilities $\{p_{ik}\}_{i,k}$ corresponding to the optimal private signaling mechanism can be found by solving a linear program in $O(N^2)$ variables. Then, we describe an efficient sampling procedure that samples sets of agents according to the the optimal marginal probabilities $\{p_{ik}\}_{i,k}$. The optimal private signaling mechanism then asks the sampled set of agents to move to ℓ_1 and the rest to stay at ℓ_0. Finally, we provide a condition on the model parameters under which the optimal signaling mechanism has a simple threshold structure and can be computed in $O(\log N)$ time.

Although private signaling mechanisms provide the principal more flexibility, a number of practical concerns often render private mechanisms infeasible [32, 57, 22]. First, private mechanisms make the strong assumption of no "information leakage" among the agents, i.e., the agents do not share their personalized information with each other. This assumption may easily fail in practice. Furthermore, fairness considerations may prevent a principal from sharing different information with different agents; a fair-minded principal may even seek to avoid the semblance of providing conflicting information to different agents.

Owing to these reasons, in Section 4.4, we analyze the problem of finding the optimal public signaling mechanism, where the principal shares the same information with all the agents. To do this, we first characterize the equilibria of the incomplete information game among the agents subsequent to receiving any public signal. While the equilibrium set is quite complex, we show that for any common posterior belief of the agents, the equilibrium that maximizes the social welfare has a simple threshold structure. Using this result, we show that the optimal public signaling mechanism can be found as a solution to a linear program with $O(N)$ variables and constraints. Furthermore, we show that the optimal mechanism only randomizes over two signals.

Finally, we numerically investigate the performance of the optimal private and public signaling mechanisms, and show that these achieve substantially higher social welfare than the no-information or full-information mechanisms.

The main point of departure of our work from past literature on information design is the modeling of negative externality among the agents. Past work on information design has focused mainly on settings with no externalities [32] or settings where there is positive externalities among the agents [22, 21]. In such settings, correlation among the agents' choice of action is beneficial, whereas the main challenge in our work is in de-correlating the agents' actions. This is especially challenging under public signaling, where public signals naturally tend to correlate the agents' actions.

To conclude this section, we note that the our approach for finding the optimal private signaling mechanism is not restricted to our model, but applies more broadly to settings where an agents' utility upon taking the action depends on the number of agents who take that action. Thus, our results on the computation of the optimal private signaling mechanism might be of independent interest to the research community.

4.1.2 Related Work

Our work focuses on information design in a resource competition scenario, adopting the Bayesian persuasion framework. We briefly survey here the related works on information design and Bayesian persuasion.

Information design problems on how a sender should persuade one or more receivers date back to [41] and [59]. The two mainstream frameworks studying this problem is the Bayesian persuasion framework, originating from [48, 64] and the "cheap talk" model [28, 36], where the main difference is the former assumes the signal sender has the power to commit to a particular information sharing mechanism.

[48] studies the basic setting with one sender and one receiver. [39] considers a setting where multiple senders wish to influence one receiver. For the more general "one sender, many receivers" setting, [8, 4, 32] considered the simplest scenario: each receiver's action imposes no externalities on other agents. [7, 8, 4] characterize polynomial time computable optimal mechanisms when the sender's utility is supermodular or anonymous submodular. [32] provide an $(1 - 1/e)$ optimal mechanism for general submodular sender utilities.

When agents' payoffs depend on other agents' actions, not too much about the structure and computation of the optimal signaling mechanism has been explored for general settings. However, a lot of work have studied various special cases. [67] characterizes the optimal mechanism in two-agent and binary action settings.

[4] provides polynomial time computable optimal mechanism for binary action settings when the sender's utility is supermodular, and similar result is given in [22, 21]. Both works point out the optimal policy correlates recommendations to take the positive action as much as possible. Our work considers resource competition settings where agents' actions have negative externalities, and to the best of our knowledge, is the first to study such settings.

Information design and Bayesian persuasion are studied in many other settings, including voting [70, 3, 13], ad auction [9], online retailing [57, 30], bilateral trade [15], advertising [26], security games [75], customer queues with delays [56], Stackelberg competition between firms [74] and team formation [44], etc. A more thorough review of this topic can be found in [31]. Beyond the one-time persuasion setting considered in most of the work we listed, several other works [20, 50] considered the problem of sequentially persuading a group of agents.

4.2 Model and Preliminaries

4.2.1 Model

We consider a model with N agents, a principal, and 2 locations, denoted by ℓ_0 and ℓ_1. Initially, all the agents are at location ℓ_0. There is a stochastic resource at ℓ_1, with the resource level denoted by θ. We primarily focus on the binary setting, with $\theta \in \Theta \triangleq \{0, 1\}$ capturing the presence or absence of a resource; our model can be extended to the case where Θ is a finite set. Without observing θ, each agent at ℓ_0 independently decides whether or not to move to ℓ_1, where she obtains a utility that depends on θ, as well as the number of other agents who also choose to move to ℓ_1. In addition, we assume that each agent incurs a moving cost if she moves to ℓ_1.

Formally, each agent $i \in [N]$ simultaneously chooses an action $a_i \in \{0, 1\}$, where $a_i = 0$ implies the agent chooses to stay at ℓ_0 and $a_i = 1$ implies she chooses to move to ℓ_1. Let $a = (a_i, a_{-i})$ denote the profile of actions chosen by all the agents, and $A \triangleq \{0, 1\}^N$ denote the set of action profiles. Note that for any $a \in A$, the number of agents that choose to move to ℓ_1 is given by $n(a) = \sum_{i=1}^N a_i$. Then, for any action profile a and any resource level θ, an agent i's utility, $U_i : \Theta \times A \to \mathbb{R}$, is given by

$$U_i(\theta, a) = \begin{cases} \theta \cdot F\left(\sum_{j=1}^N a_j\right) - r(i), & \text{if } a_i = 1; \\ 0, & \text{if } a_i = 0. \end{cases}$$

Here, $F : [N] \to \mathbb{R}_+$ is the resource sharing function that determines how the resource at ℓ_1 is shared among the agents at ℓ_1, and $r(i)$ denotes agent i's moving cost. In particular, an agent i who chooses to move to ℓ_1 receives an utility of $\theta F(n(a))$ from the resource, and incurs a cost $r(i)$ for moving. Furthermore, we have normalized the utility of staying at ℓ_0 to be zero. Without loss of generality, we assume that $r(i)$ is increasing in i, i.e., $r(1) \leq r(2) \leq \cdots r(N)$. For notational convenience, in the following, we let $r(0) = 0$ and $F(0) = F(1)$.

We make the following assumptions on the resource sharing function F:

Assumption 1. *The resource sharing function $F(n)$ is decreasing and convex in $n \geq 1$. Furthermore, the total utility from the resource $nF(n)$ is increasing and concave in $n \geq 0$.*

The first condition is meant to capture the fact that in most resource sharing settings, the amount of resource each agent receives decreases as competition increases, with the decrement diminishing with the level of competition due to market saturation. On the other hand, the second condition captures the fact that as the competition increases, the total level of resource available to the agents also increases, albeit also at a diminishing rate. This is especially common in platform markets, where the presence of more agents on side leads to better service quality for the other side of the market, leading to more conversion.

4.2.2 Information structure

We assume that the principal and the agents hold a common prior belief $\mu \in \Delta(\Theta)$ about the resource level θ, where $\mu(1)$ denotes the prior probability that $\theta = 1$, and $\mu(0) = 1 - \mu(1)$. While the resource level θ is unobserved by the agents, we assume that the principal observes θ prior to the agents' choice of actions. The principal's goal is to communicate this information about θ to the agents prior to their moving decision, in order to better position them at the two locations. (We describe the principal's objectives in more detail below.)

Following the methodology of Bayesian persuasion, the principal commits to a *signaling mechanism* as a means to share information with the agents. Formally, a signaling mechanism (Σ, ϕ) consists of a signal set Σ and a signaling scheme $\phi : \Theta \to \Delta(\Sigma^N)$. Given a signaling mechanism (Σ, ϕ) and upon observing the resource level θ, the principal first chooses a signal profile $s = (s_1, \cdots, s_N) \in \Sigma^N$ with probability $\phi(s|\theta) \in [0, 1]$. Then, the principal (privately) reveals the signal s_i to agent i. In particular, s_i is not revealed to an agent $j \neq i$.

Note that $\phi(s|\theta)$ denotes the conditional probability that the signal profile is s, given the resource level is θ. In particular, we have $\sum_{s \in \Sigma^N} \phi(s|\theta) = 1$ for each $\theta \in \Theta$. Analogously, we define $\phi(\theta, s) = \mu(\theta)\phi(\theta|s)$ to be the unconditional joint probability that the signal profile is s and the resource level is θ. Finally, we let $\phi(s)$ denote the probability of the principal selecting the signal profile $s \in \Sigma^N$, given by $\phi(s) = \phi(0, s) + \phi(1, s)$.

A special case of a signaling mechanism is a *public* signaling mechanism, where the principal publicly announces the information about θ to all the agents. In other words, the principal always shares the same information with all the agents. Such public signaling can be captured by a signaling mechanism (Σ, ϕ) where $\phi(s|\theta) = 0$ for any $s \in \Sigma^N$ with $s_i \neq s_j$ for some $i, j \in [N]$. Finally, we refer to any signaling mechanism that is not public as a private signaling mechanism.

4.2.3 Strategies and equilibrium

Since an agent i does not have access to the signals of the other agents, she maintains a belief over both the resource level θ and the signals s_{-i} of the other agents. Upon receiving her signal s_i from the principal, the agent updates her belief using Bayes' rule, before deciding whether to move. Let $q_i(\cdot|s_i)$ denote posterior belief of agent i about the resource level and the signals sent to other agents, given by

$$q_i(\theta, s_{-i}|s_i) = \frac{\mu(\theta)\phi(s_i, s_{-i}|\theta)}{\sum_{\theta', s'_{-i}} \mu(\theta')\phi(s_i, s'_{-i}|\theta')}. \tag{4.1}$$

A strategy $p_i : \Sigma \to [0, 1]$ for agent i specifies, for each possible signal s_i, the probability $p_i(s_i)$ with which she decides to move to location ℓ_1. Each agent, given her posterior belief and the strategies p_{-i} of the other agents, seeks to choose a strategy p_i that maximizes her expected utility. More precisely, given a signaling mechanism (Σ, ϕ) and the strategies p_{-i} of the other agents, the expected utility of an agent i for moving to location ℓ_1 upon receiving a signal s_i is given by

$$u_i(s_i, \textsf{move}, p_{-i}) = \mathbf{E}_{q_i}[U_i(\theta, a_i = 1, a_{-i})|a_{-i} \sim p_{-i}(s_{-i})],$$

Then, in a *Bayes-Nash equilibrium*, each agent i, upon receiving a signal s_i, decides to move to ℓ_1 if her expected utility $u_i(s_i, \textsf{move}, p_{-i})$ is positive. We have the following formal definition:

Definition 4. *Given a signaling mechanism (Σ, ϕ), a strategy profile $(p_1, \ldots, p_N)$ forms a Bayes-Nash equilibrium (BNE), if for each $i \in [N]$ and $s_i \in \Sigma$,*

$$p_i(s_i) = \begin{cases} 1 & \text{if } u_i(s_i, \textsf{move}, p_{-i}) > 0; \\ 0 & \text{if } u_i(s_i, \textsf{move}, p_{-i}) < 0. \end{cases}$$

(If $u_i(s_i, \textsf{move}, p_{-i}) = 0$, then $p_i(s_i)$ in$[0, 1]$.)

As mentioned earlier, the principal's goal is to choose a signaling mechanism to maximize the *expected social welfare*. Formally, the social welfare $W : \Theta \times A \to \mathbb{R}$ is defined as

$$W(\theta, a) \triangleq \theta \cdot n(a) \cdot F(n(a)) - \sum_{j=1}^{N} a_j r(j).$$

Here the first term denotes the total utility obtained by the $n(a) = \sum_{j=1}^{N} a_j$ agents from the resource at ℓ_1, and the second term denotes the total moving costs incurred by the agents. We assume that the principal knows the agents' moving costs. Then, the principal's decision problem is to choose a signaling mechanism (Σ, ϕ) such that in a resulting Bayes-Nash equilibrium $(p_1, \cdots, p_N)$ among the agents, the expected social welfare, given by $\mathbf{E}[W(\theta, a) | a_i \sim p_i(s_i), (\theta, s) \sim \phi]$, is maximized. In the next section, we study the problem of computation of the optimal signaling mechanism and characterize its structure. Subsequently, in Section 4.4, we analyze the related problem of optimal public signaling, where the principal is restricted to sharing information via public signaling mechanisms.

4.3 Private Signaling Mechanism

From a standard revelation-principle style argument [48, 16], there exists a *straightforward* and *persuasive* signaling mechanism that optimizes the expected social welfare. In a straightforward mechanism, the principal makes an action recommendation to each agent, and if it is optimal for each agent to follow the recommendation (assuming all others do so), then the mechanism is said to be persuasive. Thus, to obtain an optimal private mechanism, it is sufficient to restrict our attention to persuasive straightforward mechanisms. This implies that it suffices for the principal to determine the subset of agents to recommend to move for each resource level.

For $\theta = 0, 1$ and $S \subseteq [N]$, let $\phi(S|\theta)$ be the probability of recommending the set S of agents to move to ℓ_1, given resource level θ. We abuse the notation to let

$$W(\theta, S) \triangleq \theta|S|F(|S|) - \sum_{i \in S} r(i)$$

be the social welfare when resource level is θ and agents in S move to ℓ_1. The optimal signaling scheme ϕ is then obtained as solution to the following linear program:

$$\text{(LP.1)} \quad \begin{aligned} \max_{\phi} \quad & \sum_{\theta=0,1} \mu(\theta) \sum_{S \subseteq [N]} \phi(S|\theta) W(\theta, S) && (4.2) \\ s.t. \quad & \sum_{\theta=0,1} \mu(\theta) \sum_{S: i \in S} \phi(S|\theta)(\theta F(|S|) - r(i)) \geq 0, \quad i \in [N], && (4.3) \\ & \sum_{\theta=0,1} \mu(\theta) \sum_{S: i \notin S} \phi(S|\theta)(\theta F(|S|+1) - r(i)) \leq 0, \quad i \in [N], && (4.4) \\ & \sum_{S \subseteq [N]} \phi(S|\theta) = 1, \quad \theta = 0, 1, \\ & \phi(S|\theta) \geq 0, \quad \theta = 0, 1; S \subseteq [N]. \end{aligned}$$

The first two sets of constraints ensure ϕ is persuasive: the first constraint states that any agent i who is recommended to move to ℓ_1 must have non-negative utility for moving, whereas the second constraint states that any agent i who is recommended to stay must have non-positive utility for moving. The other constraints ensure that $\phi(\cdot|\theta)$ is a valid probability distribution for both $\theta = 0, 1$. Note that this linear program has $O(2^N)$ variables, and is computationally challenging. As a first step towards simplifying the problem, we note the following lemma, which allows us to only consider mechanisms that recommend all agent to stay at ℓ_0 when $\theta = 0$.

Lemma 4.3.1. *For any persuasive mechanism, there exists another persuasive mechanism that recommends every agent to stay at ℓ_0 when $\theta = 0$, and achieves a higher social welfare than the original mechanism.*

We provide proofs of all results in Appendix C.1. Although Lemma 4.3.1 reduces the size of (LP.1) by half, this linear program is still computationally challenging. However, taking a closer look at each agent's utility function, we notice that each agent i's payoff for moving when $\theta = 1$ depends only on how many other agents are moving. With this observation, we now consider an alternative formulation of (LP.1).

Given a persuasive signaling scheme ϕ satisfying Lemma 4.3.1, i.e., $\phi(\varnothing|0) = 1$, we define p_{ik} to be the joint probability, given $\theta = 1$, that this signaling scheme recommends k agents to move and agent i is among them, i.e.,

$$p_{ik} = \sum_{S} \phi(S|1) \cdot \mathbf{I}\{i \in S, |S| = k\}. \tag{4.5}$$

The following lemma allows us to write the objective and persuasive constraints of (LP.1) in terms of $p = \{p_{ik} : i \in [N], k \in [N]\}$.

Lemma 4.3.2. *The objective* (4.2) *of (LP.1) can be written as*

$$\mu(1) \sum_{k=1}^{N} \sum_{i=1}^{N} p_{ik}(F(k) - r(i)),$$

and the persuasive constraints (4.3), (4.4) *can be written as*

$$\sum_{k=1}^{N} p_{ik}(F(k) - r(i)) \geq 0, \quad i \in [N], \tag{4.6}$$

$$\sum_{k=1}^{N-1} \left(\frac{1}{k} \sum_{j=1}^{N} p_{jk} - p_{ik} \right) (F(k+1) - r(i)) + \left(1 - \sum_{k=1}^{N} \frac{1}{k} \sum_{j=1}^{N} p_{jk} \right) (F(1) - r(i)) \leq \frac{\mu(0)}{\mu(1)} r(i), \quad i \in [N]. \tag{4.7}$$

Our main result in this section shows the converse is also true: given $p = \{p_{ik} > 0 : i \in [N], k \in [N]\}$ satisfying persuasive constraints (4.6) and (4.7), with a few more linear constraints ensuring the p_{ik}'s are valid joint probabilities, there exists a persuasive signaling scheme ϕ satisfying Lemma 4.3.1 such that (4.5) holds. Furthermore, there exists a polynomial time sequential sampling procedure that samples the set of agents to recommend to move as per the signaling scheme ϕ.

Lemma 4.3.3. *Assume $p = \{p_{ik} > 0 : i, k \in [N]\}$ satisfies the persuasive constraints* (4.6) *and* (4.7)*. If p further satisfies*

$$\sum_{k=1}^{N} \frac{1}{k} \sum_{i=1}^{N} p_{ik} \leq 1, \tag{4.8}$$

$$k p_{ik} \leq \sum_{j=1}^{N} p_{jk}, \quad k \in [N], i \in [N], \tag{4.9}$$

then there exists a persuasive signaling scheme ϕ *such that* $\phi(\varnothing|0) = 1$ *and* $\phi(\cdot|1)$ *satisfies* (4.5).

Note that if p_{ik} is the joint probability that k agents are recommended to move to ℓ_1 and agent i is among them, then $q_k \triangleq \frac{1}{k}\sum_{i=1}^{N} p_{ik}$ is the probability that k agents are recommended to move. Upon letting $q_0 = 1 - \sum_{k=1}^{N} q_k$, we note that (4.8) ensures $\{q_k\}_{k=0}^{N}$ is a valid probability distribution over $\{0, \ldots, N\}$. On the other hand, $q_{ik} \triangleq \frac{p_{ik}}{q_k} = \frac{kp_{ik}}{\sum_{i=1}^{N} p_{ik}}$ is the conditional probability that agent i is recommended to move given there are k agents asked to move, and (4.9) ensures q_{ik}'s are valid probabilities.

We briefly describe the sketch of the proof of Lemma 4.3.3, and omit the details due to space limit. The main idea is that there exists a sampling procedure that samples a set of agents such that for each $i, k \in [N]$, p_{ik} is the probability that k agents are sampled and agent i is among them. Specifically, this sampling procedure first samples the size k' of the output set according $\{q_k\}_{k=0}^{N}$. Following that, given any $k' > 0$, we adopt a sequential sampling subroutine presented by [68] to sample k' agents. This subroutine eliminates one agent from agents $1, \ldots, N$ at each step, and ensures the probability a particular agent i still remains in the pool after each step is either 1, or strictly less than 1 while proportional to $p_{ik'}$. If (4.9) is satisfied, when $N - k'$ agents are eliminated, the probability each agent i remains in the output set is ensured by this subroutine to be q_{ik}, and the joint probability that k' agents are sampled and agent i is included is then $p_{ik'}$. We briefly describe the main flow of this subroutine in Appendix C.1.3 and omit its details and proof of correctness, and refer interested readers to [68]. We then let $\phi(\cdot|1)$ be the probability distribution of the set sampled according to this procedure, and note that $\phi(\cdot|1)$ thereby satisfies (4.5) and the persuasive constraints.

Summarizing the preceding discussion, from Lemma 4.3.2 and Lemma 4.3.3, to obtain the optimal private signaling mechanism, we first solve the following linear program (LP.2) to obtain the optimal solution p^*.

$$
\begin{aligned}
\max_{p_{ik}:i,k\in[N]} \quad & \sum_{k=1}^{N}\sum_{i=1}^{N} p_{ik}(F(k)-r(i)) \\
s.t. \quad & \sum_{k=1}^{N} p_{ik}(F(k)-r(i)) \geq 0, \quad i \in [N], \\
& \sum_{k=1}^{N-1}\left(\frac{1}{k}\sum_{j=1}^{N} p_{jk} - p_{ik}\right)(F(k+1)-r(i)) \\
\text{(LP.2)} \qquad & +\left(1-\sum_{k=1}^{N}\frac{1}{k}\sum_{j=1}^{N} p_{jk}\right)(F(1)-r(i)) \leq \frac{\mu(0)}{\mu(1)}r(i), \quad i \in [N], \\
& \sum_{k=1}^{N}\frac{1}{k}\sum_{i=1}^{N} p_{ik} \leq 1, \\
& kp_{ik} \leq \sum_{j=1}^{N} p_{jk}, \quad k \in [N], i \in [N], \\
& p_{ik} \geq 0, \quad k \in [N], i \in [N],
\end{aligned}
$$

After obtaining p^*, we sample the set of agents to recommend to move according to the following procedure given in Algorithm 1. Let ϕ^* be the persuasive signaling mechanism corresponding to p^* and Algorithm 1, as given by Lemma 4.3.3. ϕ^* must be feasible for (LP.1). Furthermore, the objective corresponding to ϕ^* in (LP.1) is equal to that of p^* in (LP.2). Therefore ϕ^* must be the optimal solution to

(LP.1), and is the optimal private signaling mechanism.

Note that (LP.2) has N^2 variables and $N^2 + 2N + 1$ constraints, therefore can be solved efficiently in polynomial time. The sequential sampling subroutine in [68] takes at most N steps to sample the set of agents to move, and in each step it takes polynomial time to determine which agent to eliminate. Therefore, the entire process to obtain the set of agents to recommend to move under the optimal private signaling mechanism can be completed in polynomial time.

Algorithm 1: Sampling the set of Agents to Move.

```
Input: p = {p_ik : i ∈ [N], k ∈ [N]}.
Output: S.
for k ← 1 to N do
    q_k ← (Σ_{i=1}^N p_ik) / k,
q_0 ← 1 − Σ_{k=1}^N q_k
Sample k′ according to {q_k}_{k=0}^N.
if k′ > 0 then
    Sample k′ agents from all agents.
    Let S be the set of sampled agents.
else
    S ← ∅.
```

We summarize our main result in this section as the following theorem.

Theorem 4.3.1. *The optimal private signaling mechanism can be computed by first solving (LP.2), then sampling the set of agents to recommend to move according to Algorithm 1.*

To conclude this section, we provide an observation that under certain conditions of modeling parameters, recommending the agents to follow the social optimal strategy profile is persuasive.

Let $\tilde{W}(n) \triangleq nF(n) - \sum_{i=1}^{n} r(i)$ be the social welfare when $\theta = 1$ and the first n agents move to ℓ_1. Since $nF(n)$ is concave and $r(i)$ is increasing in i, $\tilde{W}(n)$ is also concave in n. Let $i^* \triangleq \max\{i : i \in \arg\max_{0 \le i' \le N} \tilde{W}(i')\}$ be the largest maximizer of $\tilde{W}$. In the following proposition, we show the strategy profile that all agents staying when $\theta = 0$, and the first i^* agents moving when $\theta = 1$, achieves the highest social welfare among all strategy profiles, and we give a sufficient condition under which recommending the agents to follow this strategy profile is persuasive.

Proposition 3. *If $\mu(1) \le r(i^* + 1)/F(i^* + 1)$, then recommending all agents to stay when $\theta = 0$ and recommending the first i^* agents to move when $\theta = 1$ is an optimal persuasive mechanism.*

Remark. Note that $r(i^* + 1)/F(i^* + 1)$ does not depend on μ. Thus, this proposition gives an upper bound on $\mu(1)$ for achieving the maximum possible social welfare using signaling mechanisms. In Appendix C.2 we compute this upper bound for $\mu(1)$ under several typical families of resource sharing functions and cost structures. Also, since $\tilde{W}$ is concave, computing for i^* takes only $O(\log N)$ time, reducing the computational time for the optimal mechanism substantially compared with the general method.

4.4 Public Signaling Mechanism

Having characterized the optimal private signals, we next consider the principal's problem under the restriction that the signaling mechanism be public. One reason for studying this restriction is that private signaling makes strong assumptions of no information leakage among the agents, an assumption that can easily fail in practice. Public signaling by construction avoids this information leakage concern. From a practical standpoint, restriction to public signaling can arise due to fairness requirements, where the principal seeks to avoid the semblance of providing conflicting information to different agents.

A main technical difficulty in analyzing public signaling is the failure of the revelation principle argument[48, 16]: it no longer suffices to optimize only over on straightforward and persuasive *public* mechanisms. To overcome this difficulty, we first note that in a public signaling mechanism, after any information transmission from the principal, all the agents have a common belief about the resource level θ, and participate in a Bayesian game under this common belief, where each agent simultaneously chooses whether to move to ℓ_1. Thus, we begin our analysis of public signaling mechanisms by first analyzing the structure of the equilibria of this Bayesian game under any common belief of the agents.

4.4.1 Equilibrium structure

Subsequent to receiving a public signal, let $q \in [0,1]$ denote the *common belief* of the agents that $\theta = 1$. The Bayes-Nash equilibrium of the subsequent game can be represented by a strategy profile $p = (p_1, \cdots, p_N)$, where p_i denotes the probability that agent i chooses to move to ℓ_1. Our goal is to identify, for any common belief $q \in [0,1]$, the equilibrium profile that achieves the highest expected social welfare, given by $W(q,p) = q\mathbf{E}[W(1,a)|a \sim p] = q\mathbf{E}[n(a)F(n(a))|a \sim p] - \sum_i p_i r_i$.

We first note that the set of equilibria of this game can be quite complex. First, there can be a multiplicity of equilbria for any $q \in [0,1]$, as the following example illustrates for an instance with 2 agents:

Example: Let $N = 2$, $F(1) = 1$, $F(2) = 0.6$, $r(1) = 0.5$ and $r(2) = 0.6$. For $q \in [0, 0.5)$, the only equilibrium is $p(q) = (0,0)$. For $q \in [0.5, 0.6)$, the only equilibrium is $p(q) = (1,0)$. For $q \in [0.6, 5/6]$, there are three equilibria: $p_1(q) = (1,0)$, $p_2(q) = (0,1)$ and $p_3(q) = ((5q-3)/2q, (10q-5)/4q)$. For $q \in (5/6, 1)$, there is a unique equilibrium $(1,0)$ and for $q = 1$, there are two equilibria, $(1,0)$ and $(1,1)$.

Secondly, as the following proposition shows, the equilibria themselves have very counterintuitive features, where if multiple agents randomize, then the one with larger moving costs must move with a higher probability:

Proposition 4. *For an equilibrium profile $(p_1, \ldots, p_N)$ under a common belief q, let $\mathcal{I}_{mix} \subseteq [N]$ be the set of agents who randomize: $\mathcal{I}_{mix} \triangleq \{i \in [N] : 0 < p_i < 1\}$. For any agent $i, j \in \mathcal{I}_{mix}$ where $i < j$, it must be that that $p_i \leq p_j$.*

While the preceding result goes counter to the intuition that agents with higher moving cost should be less likely to move, it is explained by the fact that in order for an agent with higher moving cost to be indifferent between moving and staying, it must be that she must find location ℓ_1 to be less competitive than the one with the lower moving cost, which implies that the one with the lower moving cost must move to ℓ_1 with lower probability.

Despite the complexity of the equilibrium set, one can nevertheless consider a simple class of equilibria, namely the *threshold* equilibria, defined as follows:

Definition 5. *For common belief $q \in [0,1]$, an equilibrium $p = (p_1, \cdots, p_N)$ is said to be a threshold equilibrium, if there exists a $t \in [0, N]$ such that*

$$p_i = \begin{cases} 1, & \text{if } i < \lceil t \rceil; \\ t + 1 - \lceil t \rceil, & \text{if } i = \lceil t \rceil; \\ 0, & \text{if } i > \lceil t \rceil, \end{cases}$$

where $\lceil t \rceil$ is the smallest integer that is greater than or equal to t. We denote such an equilibrium by (q, t).

In a threshold equilibrium (q, t), at most one agent randomizes between moving and staying. Furthermore, all agents $i < \lceil t \rceil$ move to ℓ_1, whereas all agents $i > \lceil t \rceil$ stay at ℓ_0. Thus, threshold equilibria capture the intuition that agents with higher moving costs should move with lower probability. Our first result shows that, indeed, for any $q \in [0,1]$, there exists a threshold equilibrium. Before we state our result, we define the following quantities for any common belief q:

$$\begin{aligned} \bar{i}(q) &\triangleq \max\{i \in \{0, \ldots, N\} : qF(i) - r(i) \geq 0\}, \\ \underline{i}(q) &\triangleq \max\{i \in \{0, \ldots, N\} : qF(i) - r(i) > 0\}. \end{aligned} \tag{4.10}$$

Note that $qF(i) - r(i)$ denotes the expected utility of agent i for moving to location ℓ_1, if all agents $j < i$ move to ℓ_1 and all agents $j > i$ stay in ℓ_0. Thus, $\underline{i}(q)$ denotes the agent with the largest moving cost who strictly prefers to move to ℓ_1, if all agents with smaller moving costs move, and those with larger moving costs stay. Similarly, $\bar{i}(q)$ denotes the agent with the largest moving cost who does not strictly prefer to stay, under similar choices of other agents. Moreover, since $F(i)$ is decreasing in i and $r(i)$ is increasing in i, we have $\bar{i}(q) \geq \underline{i}(q)$ for any belief q. Furthermore, since $qF(0) - r(0) \geq 0$, we have $\bar{i}(q) \geq 0$ and hence $[\underline{i}(q), \bar{i}(q)] \cap [0, N] \neq \emptyset$. We have the following result:

Lemma 4.4.1. *For any $q \in [0, 1]$, and any $t \in [0, N]$, (q, t) is a threshold equilibrium if and only if $t \in [\underline{i}(q), \bar{i}(q)]$.*

The preceding lemma guarantees the existence of threshold equilibria. The question then is how do threshold equilibria fare against other equilibria in terms of their expected social welfare. The following result, our main theorem of this section, establishes that for any common belief, the expected social welfare over all equilibria is attained at a threshold equilibrium.

Theorem 4.4.1. *For any common belief $q \in [0, 1]$, the expected social welfare $W(q, p)$ under any equilibrium (q, p) is no more than that under the threshold equilibrium $(q, \underline{i}(q))$:*

$$W(q, \underline{i}(q)) \geq W(q, p), \quad \textit{for any equilibrium } p.$$

The preceding theorem has the following implications: First, the expected social welfare is always maximized at a pure equilibrium $(q, \underline{i}(q))$: no agent strictly randomizes between moving and staying. Second, under the optimal *public* signaling mechanism (Σ, ϕ), for any signal s and the induced common belief $q \in [0, 1]$, it must be that the resulting equilibrium among the agents is $(q, \underline{i}(q))$. For otherwise, one can always (publicly) recommend the agents to move to this equilibrium and improve the expected social welfare. In the next section, we use these two facts to completely characterize the optimal public signaling mechanism as a solution to linear program.

4.4.2 Optimal Public Signaling Mechanism

The results in the preceding section imply the following structure for the optimal public signaling mechanism: For each θ, the principal (publicly) recommends the threshold i number of agents to move to location ℓ_1, with the constraint that, under the induced common belief q upon receiving the recommendation i, it must be the case that $\underline{i}(q) = i$. This follows from an argument similar to the revelation-principle style argument for the private signaling mechanism, with the modification where the condition $\underline{i}(q) = i$ plays the same role of the persuasive constraints. We omit the details due to its similarity to that of the private signaling case.

Thus, the public signaling mechanism can be described by choosing the signal set to be $\Sigma = \{0, \cdots, N\}$ and $\phi : \Theta \to \Delta(\Sigma)$, where $\phi(i|\theta)$ denotes the probability that the principal recommends the first i agents to move to ℓ_1 when the resource level is θ. The common belief upon sending a public signal $s = i$ is then given by $q_i = \mu(1)\phi(i|1)/(\mu(1)\phi(i|1) + \mu(0)\phi(i|0))$. The equilibrium constraint is thus $\underline{i}(q_i) = i$ for each $i \in \Sigma$. Note that the definition of $\underline{i}(q)$ then implies that

$$q_i F(i) - r(i) > 0 \tag{4.11}$$

$$q_i F(i+1) - r(i+1) \leq 0. \tag{4.12}$$

Finally, the expected social welfare is given by $\sum_{\theta=0,1} \mu(\theta) \sum_{i=0}^{N} \phi(i|\theta) W(\theta, i)$, where $W(\theta, i) = \theta i F(i) - \sum_{j \leq i} r_j$. Taken together, using the fact that $q_i = \mu(1)\phi(i|1)/(\mu(1)\phi(i|1) + \mu(0)\phi(i|0))$ and letting $\phi(\theta, i) = \mu(\theta)\phi(i|\theta)$, we obtain the following LP in ϕ:

$$\begin{aligned} \max_{\phi} \quad & \sum_{\theta=0,1} \sum_{i=0}^{N} W(\theta, i) \cdot \phi(\theta, i) \\ \text{s.t.} \quad & (F(i) - r(i)) \cdot \phi(1, i) - r(i) \cdot \phi(0, i) \geq 0, \qquad \text{for all } i \in \{1, \cdots, N\}; \\ & (F(i+1) - r(i+1)) \cdot \phi(1, i) - r(i+1) \cdot \phi(0, i) \leq 0, \quad \text{for all } i \in \{0, \cdots, N-1\}; \\ & \sum_{i=0,\dots,N} \phi(\theta, i) = \mu(\theta), \quad \phi(\theta, i) \geq 0, \quad \text{for } \theta = 0, 1, \text{ for all } i \in \{0, \cdots, N\}. \end{aligned} \tag{4.13}$$

In (4.13), the objective is the expected social welfare under ϕ. The first two sets of constraints encode (4.11), ensuring that each signal i indeed induces the corresponding threshold equilibrium. (Note that since $W(\theta, i)$ is decreasing in i, the strict inequality in (4.11) can be replaced by the weaker inequality without loss of optimality.) The last constraints ensure $\phi(\cdot|\theta)$ is a valid probability distribution for each θ. Note that the preceding LP has $2(N + 1)$ variables and $2(N + 1)$ constraints, and thus can be efficiently solved in polynomial time.

Remark. With a similar argument as that in [48], one can show the optimal public signaling mechanism sends at most two signals with positive probabilities. That is, the optimal public signaling mechanism randomizes over two thresholds, with different weights for $\theta = 0$ and $\theta = 1$. Our results in this section can be easily extended to allow for larger state space Θ, with cardinality any finite K. In this case, the optimal public signal is the solution to a linear program with $O(NK)$ number of variables and constraints, and the optimal public scheme sends at most K signals with positive probability.

4.5 Computational Results

In this section, we compare numerically the social welfare under the optimal private and public signaling mechanisms with three benchmarks: the no-information benchmark, the full-information benchmark, and the social optimal benchmark, where all agents choose the social optimal action, that is, all agents stay when $\theta = 0$, and the first i^* agents move when $\theta = 1$ and i^* is defined earlier as the largest maximizer of $\tilde{W}(n) = nF(n) - \sum_{i=1}^{n} r(i)$. Note the third benchmark may not be achievable through a signaling mechanism, but rather provides an upper bound for the social welfare any signaling mechanism can generate.

We consider three different resource sharing function $F(i) = 1/i^{\alpha}$ for $\alpha =$

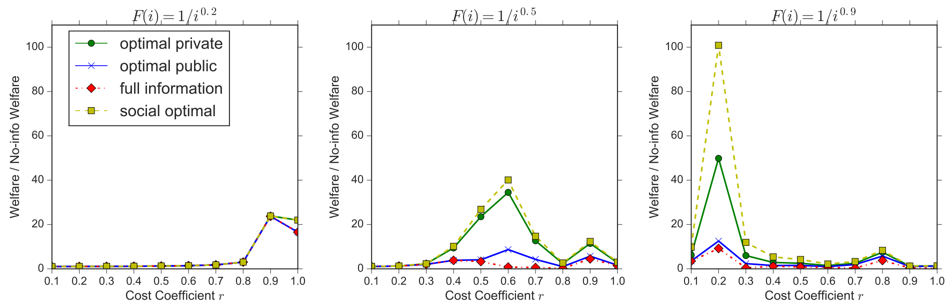

(a) Constant costs: $r(i) = 0.5r$.

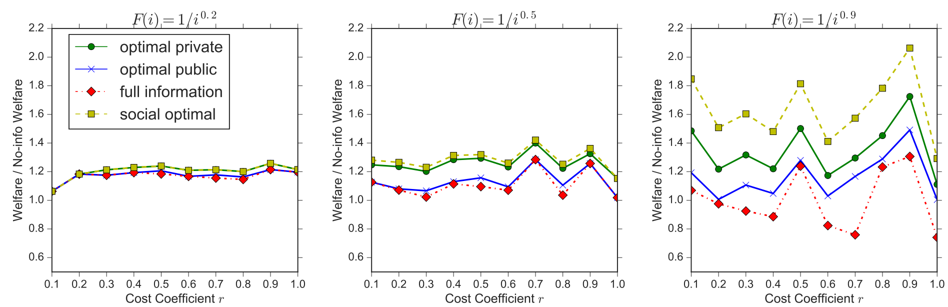

(b) Linear costs: $r(i) = 0.1ri$.

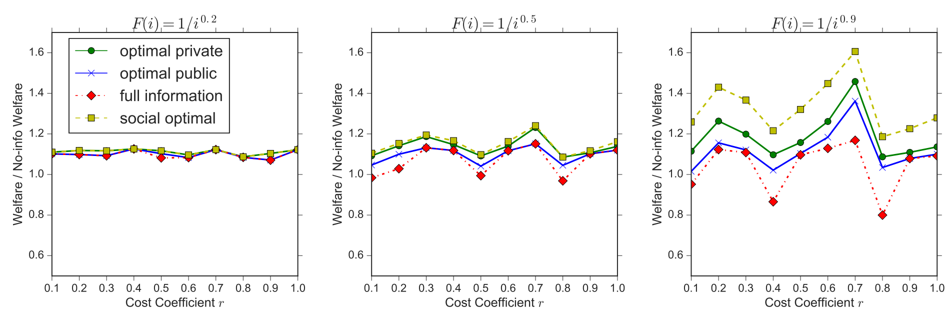

(c) Quadratic costs: $r(i) = 0.02ri^2$.

Figure 4.1: Social welfare of the optimal signaling mechanisms and the benchmarks. x-axis is the cost coefficient r, and y-axis is the social welfare of different mechanisms/benchmarks over social welfare under no information sharing. For all cases, the total number of agents $N = 20$ and the prior is $\mu(0) = 0.2$, $\mu(1) = 0.8$.

$0.2, 0.5, 0.9$. The parameter α controls the curvature of the resource sharing function, representing different resource sharing scenarios. We consider three different cost structures, constant costs where $r(i) = 0.5r$, linear costs where $r(i) = 0.1ri$ and quadratic costs where $r(i) = 0.02ri^2$, where r is the cost coefficient and we vary r to represent different cost functions. We compute the social welfare of the optimal mechanisms and the benchmarks for different resource sharing and cost functions and plot the results in Figure 4.1.

Numerical results in Figure 4.1 show that the optimal private signaling mechanism generates a social welfare close to the social optimal benchmarks in most cases, and the optimal public signaling mechanism also generates social welfare higher than the full information or no information benchmark in most scenarios. The benefit of private signaling is high when agents have similar costs, as illustrated by the subplots in the first column. [4] and [22] show that when agents' actions have positive externalities, the optimal private mechanism should correlate the recommendations to take the positive action for every agent as much as possible. Here in our setting where agents' actions have negative externalities, a different phenomenon is observed: private signaling de-correlates the agents' recommendations and generates more expected social welfare compared to public signaling where agents' actions are more correlated.

In Proposition 3, we give an upper bound for $\mu(1)$ under which the social optimal benchmark is achievable via private signaling. We also study numerically how much social welfare can be generated when this condition does not hold, by considering different priors. We presents these results in Appendix C.3.

4.6 Discussion

We consider information design in a two location resource competition model. For private signaling, we provide a method to compute the optimal mechanism in polynomial time, and also characterize a condition of model parameters under which the optimal mechanism has a simple threshold structure. For public signaling mechanisms, we characterize the structure of the socially optimal equilibrium and establish the form of the optimal public signaling mechanism. Numerical results show the optimal private and public signaling mechanisms increase the social welfare substantially compared with the no-information and full-information setting.

Readers may have noticed that our method for obtaining the optimal private signaling mechanism does not restrict to the resource sharing setting we are considering. In fact, it is applicable to all settings where the externality of each agent's action is anonymous (the utility of an agent depends only on how many other agents are taking the same action, but not on which agents are taking this action). Formally, for information design settings with binary action spaces where receiver i's utility has form $f_i\left(a_i, \sum_{j\neq i} a_j\right)$ and the sender's utility has form $f\left(\sum_i a_i\right)+\sum_i g_i(a_i)$, our method can be used to obtain the optimal private signaling mechanism in polynomial time. Such utility functions can be found in many settings, for example, in most voting settings where the sender does not differentiate across voter.

CHAPTER 5

CONCLUSION AND FUTURE DIRECTIONS

This dissertation studies spatial resource competition settings, with focus on building models to capture various characteristics of these settings and studying agents' equilibrium behaviors in these models; as well as studying the role of information sharing in these settings.

In Chapter 2, we present a model to describe general resource competition settings. We use mean field appproximation to consider a single location model whose dynamic represents that of any particular location in systems with large of number of locations and agents. Our analysis and results for the mean field model show that in equilibrium, the agents base their decision to explore solely on the state of the location they currently reside in, and on its steady state distribution, justifying analyzing spatio-temporal models under simple yet optimal models of agent behavior. In Chapter 3, we discussed several important extensions of the original model and provided corresponding equilibrium analysis for these extensions. Finally, to study information sharing in these settings, in Chapter 4, we present a two-location model and study the optimal mechanism design problem with this model. We provide polynomial time schemes to obtain the optimal private and public signaling mechanisms, and our method for computing the optimal private signaling mechanism can be applied to more general settings.

Our models and analysis raise many questions for future research, beyond the ones we discussed in each chapter.

First, in Chapter 2 and Chapter 3, we have used the methodology of mean field approximation to analyze a single location in isolation (or a single location for each type of location in Chapter 3), assuming the joint effects of the agents in other locations are approximated by the mean field distribution. A natural question is whether the resulting mean field equilibrium strategies constitute an approximate equilibrium in a system with a large but finite number of agents and locations. Such approximation results for mean field equilibrium have been obtained in other contexts (see, for example, [1, 10, 46]). In the finite system, a single agent visits multiple locations over her lifetime, inducing correlations among the states of those locations. The analytical challenge in obtaining an approximation result involves showing such correlations vanish as the system size increases, and the dynamics of a location in the finite system approaches the dynamics of the single-location in the mean field model in the limit. To overcome these challenges and rigorously show whether and how the mean field equilibrium approximates an equilibrium in a finite system provides theoretical guarantee of the effectiveness of our model and is itself an interesting research problem.

We provided two extensions of our original model in Chapter 3. Many important factors in various resource competition scenarios are still ignored and yet to be considered. For example, agents may be heterogeneous and have different valuation of resources. Agents may compete for more than one resources that have different dynamics. These settings are more complicated, and studying them may bring better insights to understant corresponding real world scenarios.

Going beyond our model in Chapter 4, in many spatial settings, there are more than two locations. The form of signals and the action and belief structure of the agents can be much more complicated. Fully solving the optimal signaling mechanism in such a system may be computationally intractable, while it is hopeful heuristic mechanisms with theoretical guarantee or good empirical performance on the social welfare generated may exist and is an interesting future research path. Meanwhile, in most settings, agents migrate among locations so signaling is not "one-shot" but rather a sequential and dynamic process, where signals sent at a time impact the actions and beliefs of agents ever since hence affect signals that should be sent later. The Bayesian persuasion framework does not extend naturally to these settings. Building models and analyze such dynamic signaling process is an important future path for better understanding real world spatial signaling problems.

APPENDIX A

APPENDIX OF CHAPTER 2

A.1 Proofs

A.1.1 Existence and uniqueness of invariant distribution of $\mathsf{MC}(\xi,\kappa)$

In this section, 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 A.1.1. *For any Markovian strategy ξ and arrival rate $\kappa \geq 0$, there exists a unique steady state distribution for* $\mathsf{MC}(\xi,\kappa)$ *satisfying* (2.4).

Proof. Fix a Markovian strategy ξ, and an arrival rate $\kappa > 0$. We prove the lemma statement 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 (2.1) 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 $T_z(\mathbb{S}_0)$ is the first return time of the chain to $\mathbb{S}_0$, given it starts at $(z,0)$:

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

given $(Z_0, N_0) = (z, 0)\}$.

In the following, we show that $T_z(\mathbb{S}_0)$ has 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 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 A.1.19. Define $\tilde{T}_z(0)$ to be the first return time to 0 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 $z \in \mathbb{Z}$. □

A.1.2 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 will play 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 [53], where 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 $C_b(\mathbb{S})$ denote the set of bounded function $h : \mathbb{S} \to \mathbb{R}$. (Note that since we impose the discrete topology on $\mathbb{S}$, any such h is also continuous.) We endow $C_b(\mathbb{S})$ with the sup-norm:

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

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

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

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

Let $\Gamma = \{\pi(\xi,\kappa) : \xi \in \Pi, \kappa \in [\beta\lambda(1-\gamma), \beta\lambda]\} \subseteq \mathcal{M}(\mathbb{S})$ 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:

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

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 A.1.2. *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. In Lemma A.1.3, we show that the set of distributions Γ is uniformly tight. Then, from Prohorov's theorem [17], we obtain that $\overline{\Gamma}$ is compact. Observe that

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

This follows from the fact that $\pi(\xi, \kappa)$ is the unique probability distribution over $\mathbb{S}$ for which (2.4) holds.

In Lemma A.1.4, we show that $\Lambda(\xi, \kappa, \nu)$ is jointly (and uniformly) continuous 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 [14] to (A.4).

□

The following two auxiliary lemmas are used in the proof of Lemma A.1.2.

Lemma A.1.3. *The set* Γ *of invariant distributions is tight.*

Proof. We prove the 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 Lemma A.1.18 and Lemma A.1.19, 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

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

Since k^ϵ is independent of the choice of (ξ, κ), we obtain that Γ is tight. $\square$

The following lemma proves the joint continuity of Λ.

Lemma A.1.4. *The function* Λ *as defined in* (A.3) *is jointly (and uniformly) continuous.*

Proof. 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

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

Now, note that

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

Now, using (2.1), we obtain that

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

and hence

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

Thus, from (A.6), we obtain

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

Next, observe that

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

Now, again from (2.1) and after some straightforward algebra, we obtain that

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

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

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

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

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

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

A.1.3 Existence of κ satisfying equilibrium condition

In this section we show for any Markovian strategy ξ, 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 the equation (2.6).

Towards 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 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 $\phi(\xi, \beta\lambda(1 - \beta)) \leq \beta$ and $\phi(\xi, \beta\lambda) \geq \beta$, which implies any κ such that $\phi(\xi, \kappa) = \beta$ must lie in $[\beta\lambda(1 - \gamma), \beta\lambda]$. The result then follows once we show $\phi(\xi, \kappa)$ is a continuous function of κ.

In the rest of this section, 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.

Strict monotonicity of $\phi(\cdot)$

Lemma A.1.5. *Given any Markovian strategy ξ, $\phi(\kappa)$ is a strictly increasing function of κ on $[\beta\lambda(1-\gamma), \beta\lambda]$.*

Proof. 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 A.1.19, where $Z_t^{(1)} = Z_t^{(2)}$ and $N_t^{(1)} \leq N_t^{(2)}$ for all $t \geq 0$.

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}. Since $N_t^{(1)} \leq N_t^{(2)}$ for all t, we have $\phi(\kappa_1) \leq \phi(\kappa_2)$.

Next, suppose for the sake of contradiction that $\phi(\kappa_1) = \phi(\kappa_2)$. Since $Z_t^{(1)} = Z_t^{(2)}$ and $N_t^{(1)} \leq N_t^{(2)}$ for all $t \geq 0$, we have

$$\mathbf{I}\{Z_t^{(1)} = z, N_t^{(1)} \geq n\} \leq \mathbf{I}\{Z_t^{(2)} = z, N_t^{(2)} \geq n\}, \tag{A.10}$$

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

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

$$\frac{1}{t}\int_0^t \mathbf{I}\{Z_s^{(i)} = z, N_s^{(i)} \geq n\}ds \to \sum_{n' \geq n} \pi_i(z, n') \tag{A.11}$$

almost surely as $t \to \infty$, for $i = 1, 2$. By (A.10) and (A.11) we have

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

for any (z, n), and

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

Since 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 $i = 1, 2$, the equation (2.4) 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

$$\begin{aligned} &\kappa_i(\pi_i(z, n) - \pi_i(z, n-1)) \\ =& \sum_{y \neq z} \mu_{y,z} \pi_i(y, n) + \pi_i(z, n+1)\lambda(n+1)(1 - \gamma \xi(z, n+1)) \\ &- \pi_i(z, n) \left(\sum_{y \neq z} \mu_{z,y} + \lambda n (1 - \gamma \xi(z, n)) \right). \end{aligned} \tag{A.12}$$

Since $\pi_1 = \pi_2$, the right hand side of (A.12) 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 $\pi_1 = \pi_2$ implies 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. $\square$

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 A.1.5, this implies that any κ for which $\phi(\kappa) = \beta$ must lie in the interval $[\beta\lambda(1-\gamma), \beta\lambda]$.

Lemma A.1.6. *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. 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 A.1.19, we have $N_t^{(1)} \preccurlyeq_{\mathsf{sd}} N_t \preccurlyeq_{\mathsf{sd}} N_t^{(2)}$ 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 Poisson distribution with mean κ/λ_i. □

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 A.1.7. *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. Given Markovian strategies ξ_1 and ξ_2, and arriving rates κ_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 $k \geq 0$,

$$\begin{aligned}
&|\phi(\xi_1, \kappa_1) - \phi(\xi_2, \kappa_2)| \\
&= \left| \sum_{z,n} n(\pi_1(z,n) - \pi_2(z,n)) \right| \\
&\leq \left| \sum_{z\in\mathbb{Z}} \sum_{n\leq k} n(\pi_1(z,n) - \pi_2(z,n)) \right| + \left| \sum_{z\in\mathbb{Z}} \sum_{n>k} n(\pi_1(z,n) - \pi_2(z,n)) \right| \\
&\leq \left| \sum_{z\in\mathbb{Z}} \sum_{n\leq k} n(\pi_1(z,n) - \pi_2(z,n)) \right| + \sum_{z\in\mathbb{Z}} \sum_{n>k} n\pi_1(z,n) + \sum_{z\in\mathbb{Z}, n>k} n\pi_2(z,n).
\end{aligned} \tag{A.13}$$

Now, bounding the first term, we obtain

$$\begin{aligned}
&\left| \sum_{z\in\mathbb{Z}} \sum_{n\leq k} n(\pi_1(z,n) - \pi_2(z,n)) \right| \\
\leq & k \sum_{z\in\mathbb{Z}} \sum_{n\leq k} |\pi_1(z,n) - \pi_2(z,n)| \leq k\|\pi_1 - \pi_2\|_1.
\end{aligned} \tag{A.14}$$

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 $\widehat{N}_0 = n$, denote 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 Lemma A.1.18 and Lemma A.1.19, 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

$$\sum_{z\in\mathbb{Z}, n>k} n\pi_i(z,n) \leq \sum_{n>k} n\widehat{\pi}(n), \quad i = 1, 2. \tag{A.15}$$

Thus, from(A.13), (A.15) and (A.14), 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 A.1.2, 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.

□

Continuity

For any $\xi \in \Pi$, let $\kappa(\xi)$ denote the unique value of κ for which $\pi(\xi, \kappa)$ satisfies (2.6). Below, we show that $\kappa(\xi)$ is a continuous function of ξ.

Lemma A.1.8. *The map $\xi \mapsto \kappa(\xi)$ is continuous.*

Proof. Define $W(\xi, \kappa) = -|\beta - \phi(\xi, \kappa)|$. Note that, from Lemma A.1.5, we obtain

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

From Lemma A.1.7, 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 [14]. □

A.1.4 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 $\mathbf{T} : \Pi \times \mathbb{R}_+ \times C_b(\mathbb{S}) \to C_b(\mathbb{S})$ 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 $U \in C_b(\mathbb{S})$, the function $W = \mathbf{T}(\xi, V_{\mathsf{sw}}, U)$ is defined as follows:

$$W(z, n) = F(z, n) + \gamma \max \left\{ \mathbf{E}^{\xi} \left[U(Z_\tau, N_\tau) | z, n \right], V_{\mathsf{sw}} \right\}, \quad \text{for all } (z, n) \in \mathbb{S}. \tag{A.16}$$

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 A.1.9. *For any $\xi \in \Pi$ and $V_{\mathsf{sw}} > 0$, we have $\mathbf{T}(\xi, V_{\mathsf{sw}}, U) \in C_b(\mathbb{S})$ for all $U \in C_b(\mathbb{S})$. Furthermore, the map $\mathbf{T}(\xi, V_{\mathsf{sw}}, \cdot) : C_b(\mathbb{S}) \to C_b(\mathbb{S})$ is a contraction (with contraction parameter γ) for any $\xi \in \Pi$ and $V_{\mathsf{sw}} > 0$.*

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

$$\begin{aligned} \mathcal{V}_{\mathsf{st}}(z, n; \xi, V_{\mathsf{sw}}) &= \mathbf{E}^{\xi} \left[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}}) | z, n \right] \\ \mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}}) &= \sum_{(z,n) \in \mathbb{S}} \pi(z, n) \mathcal{V}_{\mathsf{st}}(z, n+1; \xi, V_{\mathsf{sw}}), \end{aligned} \tag{A.17}$$

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

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

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

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 definition.

Lemma A.1.10. *For any $\xi \in \Pi$ and $V_{\mathsf{sw}} > 0$, the value functions satisfy $|\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})| \leq \|\mathcal{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}})\|_\infty \leq \|\mathcal{V}(\xi, V_{\mathsf{sw}})\|_\infty \leq \overline{\mathsf{V}} = \frac{\|F\|_\infty}{1-\gamma}$.*

Proof. Observe that from (A.17), we have $|\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})| \leq \|\mathcal{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}})\|_\infty \leq \|\mathcal{V}(\xi, V_{\mathsf{sw}})\|_\infty$. Also, from the fact that $\mathcal{V}(\xi, V_{\mathsf{sw}})$ is the fixed-point of $\mathbf{T}(\xi, V_{\mathsf{sw}}, \cdot)$, we obtain

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

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

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 A.1.11. *For any $\xi \in \Pi$ and $V_{\mathsf{sw}} > 0$, we have $\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}}) \geq \underline{\mathsf{V}}$.*

Proof. Observe that

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

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

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

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

$$\mathcal{V}_{\mathsf{st}}(z, n; \xi, V_{\mathsf{sw}}) \geq \frac{1}{n + \beta + \Psi} F(z, n) \geq \frac{1}{n(1 + \beta + \Psi)} F(z, n). \tag{A.18}$$

Next, observe that $\pi = \pi(\xi, \kappa(\xi))$ satisfies the steady-state equation (2.4):

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

where Q^{ξ} denote the transition kernel of the Markov chain $\mathsf{MC}(\xi, \kappa(\xi))$. Using the expression (2.1) for Q^{ξ}, we obtain

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

This implies that

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

Thus, we obtain

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

Now, from Lemma A.1.19, 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 $\pi(z, 0) \geq \pi_{\mathsf{res}}(z)\mathbf{P}(N_\infty^{(1)} = 0)$, where the steady state $N_\infty^{(1)}$ is given by a Poisson distribution with parameter $\beta/(1-\gamma)$, implying $\mathbf{P}(N_\infty^{(1)} = 0) = \exp(-\beta/(1-\gamma))$. Thus, we obtain

$$\pi(z,n) \geq \pi_{\mathsf{res}}(z)\exp\left(-\frac{\beta}{1-\gamma}\right)\frac{\beta^n(1-\gamma)^n}{(1+\beta+\Psi)^n n!}. \tag{A.19}$$

Finally, from (A.17), we have

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

where we use (A.18) in the first inequality and (A.19) in the second. □

A.1.5 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 non-empty and convex.

In this section, we provide characterization of a compact set $\widehat{\Pi} \subseteq \Pi$ of strategies such that if $\xi \in \widehat{\Pi}$ and $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$, then $\mathcal{X}(\xi, V_{\mathsf{sw}}) \subseteq \widehat{\Pi}$. 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

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

and let K_1 be defined as

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

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 her location, if the number of agents at her 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 $V_{\mathsf{sw}} \geq \underline{\mathsf{V}}$, then the optimal action for an agent at the state (z,n) is to switch if $n \geq K_{\max}$.

Lemma A.1.12. *For $\xi \in \widehat{\Pi}$ and $V_{\mathsf{sw}} \geq \underline{\mathsf{V}}$, we have $\mathcal{V}_{\mathsf{st}}(z,n;\xi,V_{\mathsf{sw}}) < V_{\mathsf{sw}}$ for all $z \in \mathbb{Z}$ and for all $n \geq K_{\max}$.*

Proof. 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 $V_{\mathsf{sw}} \geq \underline{\mathsf{V}}$. 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 $(Z_0, N_0) = (z, n)$ 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 $V^\phi_{\mathsf{st}}(z, n)$ denote agent i's continuation payoffs under the strategy ϕ, subsequent to her making the decision to stay and not leaving the system, given the state of the location (z, n). We have the following expression for the $V^\phi_{\mathsf{st}}(z, n)$:

$$V^\phi_{\mathsf{st}}(z, n) = \mathbf{E}\left[\sum_{\ell=1}^{\infty} \gamma^{\ell-1} F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\} + \gamma^\ell \mathbf{I}\{\tau^\ell = \tau^\phi\} V_{\mathsf{sw}}\right]. \tag{A.20}$$

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,

$$\begin{aligned} V^\phi_{\mathsf{st}}(z, n) &= \sum_{\ell=1}^{\infty} \gamma^{\ell-1} \mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\}\right] + \sum_{\ell=1}^{\infty} \gamma^\ell V_{\mathsf{sw}} \mathbf{P}\left(\tau^\ell = \tau^\phi\right) \\ &\leq \sum_{\ell=1}^{\infty} \gamma^{\ell-1} \mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\}\right] + \gamma V_{\mathsf{sw}}. \end{aligned} \tag{A.21}$$

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

$$\begin{aligned} \mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\}\right] &= \mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\} \middle| N_\ell \geq \widehat{n}\right] \mathbf{P}(N_\ell \geq \widehat{n}) \\ &\quad + \mathbf{E}\left[F(Z_\ell, N_\ell)\mathbf{I}\{\tau^\ell \leq \tau^\phi\} \middle| N_\ell < \widehat{n}\right] \mathbf{P}(N_\ell < \widehat{n}) \\ &\leq \frac{1}{2}(1-\gamma)^2 \underline{\mathsf{V}} + \|F\|_\infty \mathbf{P}(N_\ell < \widehat{n}). \end{aligned} \tag{A.22}$$

Here, in the inequality, the first term follows from the fact that since $n > K_{\max}$, we have $\widehat{n} > K_0$, and hence $F(Z_\ell, N_\ell) < \frac{(1-\gamma)^2\underline{V}}{2}$ 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 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 $\widehat{t} = \frac{\log(n-1)}{2\lambda}$. For each $\ell \leq \widehat{\ell}$, we have $\tilde{N}_\ell \geq \tilde{N}_{\widehat{\ell}}$, and hence,

$$\begin{aligned}
\mathbf{P}(\tilde{N}_\ell < \widehat{n}) &\leq \mathbf{P}(\tilde{N}_{\widehat{\ell}} < \widehat{n}) \\
&= \mathbf{P}(\tilde{N}_{\widehat{\ell}} < \widehat{n} | \tau^{\widehat{\ell}} < \widehat{t})\mathbf{P}(\tau^{\widehat{\ell}} < \widehat{t}) + \mathbf{P}(\tilde{N}_{\widehat{\ell}} < \widehat{n} | \tau^{\widehat{\ell}} \geq \widehat{t})\mathbf{P}(\tau^{\widehat{\ell}} \geq \widehat{t}) \\
&\leq \mathbf{P}(\tilde{N}_{\widehat{\ell}} < \widehat{n} | \tau^{\widehat{\ell}} < \widehat{t})\mathbf{P}(\tau^{\widehat{\ell}} < \widehat{t}) + \mathbf{P}(\tau^{\widehat{\ell}} \geq \widehat{t}) \qquad \text{(A.23)} \\
&\leq \mathbf{P}(\tilde{N}_{\widehat{t}} < \widehat{n} | \tau^{\widehat{\ell}} < \widehat{t})\mathbf{P}(\tau^{\widehat{\ell}} < \widehat{t}) + \mathbf{P}(\tau^{\widehat{\ell}} \geq \widehat{t}) \\
&\leq \mathbf{P}(\tilde{N}_{\widehat{t}} < \widehat{n}) + \mathbf{P}(\tau^{\widehat{\ell}} \geq \widehat{t}),
\end{aligned}$$

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

Now, observe that since each agent $j \neq i$ stays in the auxiliary system for a time distributed independently and exponentially with rate λ, the probability that the agent $j \neq i$ 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 $1 + \text{Bin}(n-1, \frac{1}{\sqrt{n-1}})$, where $\text{Bin}(\cdot,\cdot)$ denotes the binomial distribution. (Recall that in the auxiliary system, agent i never leaves.) Now, note that $\mathbf{E}[\text{Bin}(n-1, \frac{1}{\sqrt{n-1}})] = \sqrt{n-1} > \widehat{n} - 1$. From this, we obtain

$$\begin{aligned}\mathbf{P}(\tilde{N}_{\widehat{t}} < \widehat{n}) &= \mathbf{P}\left(\text{Bin}\left(n-1, \frac{1}{\sqrt{n-1}}\right) < \widehat{n} - 1\right) \\ &\leq \mathbf{P}\left(\text{Bin}\left(n-1, \frac{1}{\sqrt{n-1}}\right) < \frac{1}{2}\sqrt{n-1}\right) \\ &\leq \exp\left(-\frac{1}{8}\sqrt{n-1}\right),\end{aligned} \tag{A.24}$$

where we have used the Chernoff bound [60] for the lower tail of the binomial distribution in the last inequality.

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

$$\mathbf{P}(\tau^{\widehat{\ell}} > \widehat{t}) \leq \frac{\mathbf{E}[\tau^{\widehat{\ell}}]}{\widehat{t}} = \frac{\widehat{\ell}}{\lambda\widehat{t}} \leq \frac{2}{\sqrt{\log(n-1)}}. \tag{A.25}$$

Thus, combining (A.22), (A.23), (A.24) and (A.25), we obtain for all $\ell \leq \widehat{\ell}$,

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

Thus, using (A.21), we have

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

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

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

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

Now, note that since $n \geq K_{\max} \geq K_1$, we have

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

Thus we obtain $V_{\mathsf{st}}^{\phi}(z,n) \leq \frac{1}{1-\gamma}\left(\frac{(1-\gamma)^2\underline{\mathsf{V}}}{2} + \frac{(1-\gamma)^2\underline{\mathsf{V}}}{4}\right) + \gamma V_{\mathsf{sw}} = \frac{3(1-\gamma)}{4}\underline{\mathsf{V}} + \gamma V_{\mathsf{sw}}$. Since this inequality holds for all strategies $\phi \in \Pi$ and since $V_{\mathsf{sw}} \geq \underline{\mathsf{V}}$, we obtain $\mathcal{V}_{\mathsf{st}}(z,n;\xi,V_{\mathsf{sw}}) < V_{\mathsf{sw}}$ for all $z \in \mathbb{Z}$ and all $n \geq K_{\max}$. □

Let $\Upsilon = \widehat{\Pi} \times [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$. 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 $n \geq K_{\max}$. 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 $\mathcal{X} : \Upsilon \rightrightarrows \widehat{\Pi}$.

A.1.6 Upper-hemicontinuity of $\mathcal{R}$

For $(\xi, V_{\mathsf{sw}}) \in \Upsilon$, define the map $\mathcal{R}$ as $\mathcal{R}(\xi, V_{\mathsf{sw}}) = \mathcal{X}(\xi, V_{\mathsf{sw}}) \times \{\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})\}$. Note that from Lemma A.1.10, Lemma A.1.11 and Lemma A.1.12, 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 section, we seek to show that this correspondence is upper-hemicontinuous. This result directly used in proof for Theorem 2.3.1.

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

Lemma A.1.13. *The map $(\xi, V_{\mathsf{sw}}) \to \mathcal{V}(\xi, V_{\mathsf{sw}})$ is (jointly) continuous in $(\xi, V_{\mathsf{sw}}) \in \Upsilon$.*

Proof. 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 A.1.12, we obtain $W_i(z, n) = F(z, n) + \gamma V^i_{\mathsf{sw}}$ for all $z \in \mathbb{Z}$ and $n \geq K_{\max}$. This implies that

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

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

Next, we have

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

where we have used Lemma A.1.9 in the last inequality. Using the fact that $\|W_1 - W_2\|_* \leq \|W_1 - W_2\|_\infty \leq \max\{\|W_1 - W_2\|_*, \gamma|V^1_{\mathsf{sw}} - V^2_{\mathsf{sw}}|\} \leq \|W_1 - W_2\|_* + \gamma|V^1_{\mathsf{sw}} - V^2_{\mathsf{sw}}|$ and after some straightforward algebra, we obtain

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

From Lemma A.1.14, we obtain that the first term in the parenthesis 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 $(\xi, V_{\mathsf{sw}}) \in \Upsilon$. $\square$

The following auxiliary lemma is used in the proof of Lemma A.1.13.

Lemma A.1.14. *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 $U \in C_b(\mathbb{S})$, 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. 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

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

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, since $U \in C_b(\mathbb{S})$ and τ is exponentially distributed with parameter λ, we have

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

with similar expressions for ξ^m in place of ξ. For 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 (2.1) 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 [73, See pg. 2183, Example 1.1] or [35, pg. 262, problem 8], we obtain that the measure $\mathbf{P}^m(\cdot|z, n, \tau = t)$ converges weakly to $\mathbf{P}^{\xi}(\cdot|z, n, \tau = t)$. From this, we conclude that $\int_0^T \lambda \exp(-\lambda t)\mathbf{E}^m[U(Z_t, N_t)|z, n, \tau = t]dt$ converges to $\int_0^T \lambda \exp(-\lambda t)\mathbf{E}^{\xi}[U(Z_t, N_t)|z, n, \tau = t]dt$ as $m \to \infty$. This completes the proof. □

The continuity of $\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})$ is then obtained as a corollary of Lemma A.1.13.

Lemma A.1.15. *The value function $\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})$ is jointly continuous in $(\xi, V_{\mathsf{sw}}) \in \Upsilon$.*

Proof. Recall the definition (A.17) of $\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})$:

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

where $\pi = \pi(\xi, \kappa(\xi))$ is the invariant distribution of $\mathsf{MC}(\xi, \kappa(\xi))$. From Lemma A.1.10, we have $\|\mathcal{V}_{\mathsf{st}}(\xi, V_{\mathsf{sw}})\|_\infty \leq \overline{\mathsf{V}}$. Also, note that Lemma A.1.2 and Lemma A.1.8 imply that the invariant distribution $\pi(\xi, \kappa(\xi))$ is continuous. Moreover, from Lemma A.1.3, 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

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

From Lemma A.1.13, we obtain that as $m \to \infty$, the first term converges to zero. Moreover, from the same argument as in the proof of Lemma A.1.13, we obtain that $\mathbf{E}^{\xi^m}[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}})|z, n] \to \mathbf{E}^{\xi}[\mathcal{V}(Z_\tau, N_\tau; \xi, V_{\mathsf{sw}})|z, n]$ 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 A.1.16. *The correspondence $\mathcal{R} : \Upsilon \rightrightarrows \Upsilon$ is upper-hemicontinuous.*

Proof. By definition, $\mathcal{R}(\xi, V_{\mathsf{sw}}) = \mathcal{X}(\xi, V_{\mathsf{sw}}) \times \{\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})\}$ for $(\xi, V_{\mathsf{sw}}) \in \Upsilon$. From Lemma A.1.15, we obtain that $\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})$ 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 $(\xi^n, V^n_{\mathsf{sw}}, \zeta^n) \to (\xi, V_{\mathsf{sw}}, \zeta)$ as $n \to \infty$ such that $\zeta^n \in \mathcal{X}(\xi^n, V^n_{\mathsf{sw}})$ for each $n \geq 0$. 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 $\mathcal{V}_{\mathsf{st}}(z, n, \xi^m, V^m_{\mathsf{sw}}) > V^m_{\mathsf{sw}}$, 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. Since $\zeta^m \to \zeta$, 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 $\zeta(z, m) \in \mathcal{X}(\xi, V_{\mathsf{sw}})$. □

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

Lemma A.1.17. *The correspondence $\mathcal{R} : \Upsilon \rightrightarrows \Upsilon$ is upper-hemicontinuous.*

Proof. By definition, $\mathcal{R}(\xi, V_{\mathsf{sw}}) = \mathcal{X}(\xi, V_{\mathsf{sw}}) \times \{\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})\}$ for $(\xi, V_{\mathsf{sw}}) \in \Upsilon$. From Lemma A.1.15, we obtain that $\mathcal{V}_{\mathsf{arr}}(\xi, V_{\mathsf{sw}})$ 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 $(\xi^n, V^n_{\mathsf{sw}}, \zeta^n) \to (\xi, V_{\mathsf{sw}}, \zeta)$ as $n \to \infty$ such that $\zeta^n \in \mathcal{X}(\xi^n, V^n_{\mathsf{sw}})$ for each $n \geq 0$. 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. Since $\zeta^m \to \zeta$, 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 $\zeta(z, m) \in \mathcal{X}(\xi, V_{\mathsf{sw}})$. □

A.1.7 Existence of an optimal threshold strategy

In this section we provide the proof of Theorem 2.4.1. We prove this result in two steps: first, we prove Lemma 2.4.1, which states that the value function $V_{\mathsf{st}} : \mathbb{S} \to \mathbb{R}$ is non-increasing in the number of agents n at the location for any fixed resource level z $in\mathbb{Z}$. Second, we show in Lemma A.1.12 that $\lim_{n\to\infty} V_{\mathsf{st}}(z, n) \leq \gamma V_{\mathsf{sw}}$ 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 Theorem 2.4.1. 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 $n_1 \leq n_2$. For any function $f : \mathbb{S} \to \mathbb{R}$, we say 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 $f(z_1, n_1) \geq f(z_2, n_2)$.

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 [66]. 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 $(z_1, n_1) \preccurlyeq_p (z_2, n_2)$.

Since $F(z, n)$ 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 $(z_1, n_1) \preccurlyeq_p (z_2, n_2)$. By using an argument same as that in the proof of Lemma A.1.19, 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 a distributed independently as an exponential with rate λ. Since $(Z_\tau^{(i)}, N_\tau^{(i)}) \sim \nu_{(z_i,n_i)}$ for $i = 1, 2$, we obtain the result. □

A.1.8 Coupling results

In this section, 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 which states that a queue with 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 A.1.18. *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 $N_0^{(1)} = N_0^{(2)}$, and 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 $N_t^{(2)}$, i.e., 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 2.3.1, 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 that 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 A.1.19. *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. 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 $t_0 = 0$ and $(Z_0, N_0) = (z_0, n_0)$ 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 $\zeta_k \triangleq \frac{n_{k-1}(1-\gamma\xi(z_{k-1},n_{k-1}))}{v_{k-1}(1-\gamma\xi_0(u_{k-1},v_{k-1}))}$ and $\eta_k \triangleq \left(\frac{1-\gamma\xi_0(u_{k-1},v_{k-1})}{\gamma\xi_0(u_{k-1},v_{k-1})}\right)\max(\zeta_k - 1, 0)$. 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 $N_0 \leq N_0^{(0)}$. Furthermore, by construction, we have $z_k = u_k$ for all k, and hence $Z_t^{(0)} = Z_t$ for all $t \geq 0$.

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

APPENDIX B

APPENDIX OF CHAPTER 3

B.1 Proof of Theorem 3.1.1

The proof consists of the following two main steps.

STEP 1 Given $\lambda > 0, \beta > 0$ transition rates $\{\mu^{\theta}_{zy} > 0 : z, y \in \mathbb{Z}\}_{\theta\in\Theta}$ and type distribution $\{p_\theta : \theta \in \Theta\}$, we show for each Markovian strategy profile $\xi \triangleq \{\xi_\theta : \theta \in \Theta\}$, there exists $\{\kappa_\theta(\xi) : \theta \in \Theta\}$, $\{V^{\theta}_{\mathsf{arr}}(\xi) : \theta \in \Theta\}$, $\{\beta_\theta(\xi) : \theta \in \Theta\}$ and $V_{\mathsf{sw}}(\xi)$ that satisfies (3.1) ~ (3.2) and (3.4) ~ (3.8).

STEP 2 Given $\lambda > 0, \beta > 0$ transition rates $\{\mu^{\theta}_{zy} > 0 : z, y \in \mathbb{Z}\}_{\theta\in\Theta}$ and type distribution $\{p_\theta : \theta \in \Theta\}$, for any Markovian strategy profile ξ, let $\{\kappa_\theta(\xi) : \theta \in \Theta\}$, $\{V^{\theta}_{\mathsf{arr}}(\xi) : \theta \in \Theta\}$, $\{\beta_\theta(\xi) : \theta \in \Theta\}$ and $V_{\mathsf{sw}}(\xi)$ be given as in STEP 1. Consider the correspondence $\xi \mapsto \mathcal{X}(\xi)$, where $\mathcal{X}(\xi) \triangleq \times_{\theta\in\Theta}\mathcal{X}_\theta(\xi_\theta)$ and $\mathcal{X}_\theta(\xi_\theta)$ is the set of best responses for the decision problem $\mathsf{DEC}(\xi_\theta, \kappa_\theta(\xi), V_{\mathsf{sw}}(\xi))$. We show the correspondence has a fixed point ξ^* and $(\xi^*, \{\kappa_\theta(\xi^*) : \theta \in \Theta\}, \{V^{\theta}_{\mathsf{arr}}(\xi^*) : \theta \in \Theta\}, \{\beta_\theta(\xi)^* : \theta \in \Theta\}, V_{\mathsf{sw}}(\xi^*))$ gives the mean field equilibrium.

Below we provide a complete proof of STEP 1 and STEP 2. We omit details of the proof of lemmas and claims that can be shown straightforwardly using similar arguments as in the proof for existence of mean field equilibria of the homogeneous locations model.

We first prove STEP 1, and throughout the proof we assume we are given a Markovian strategy profile $\xi = \{\xi_\theta : \theta \in \Theta\}$, and we would sometimes omit ξ and ξ_θ in notations when no confusions will arise.

For each type $\theta \in \Theta$, for any $\kappa_\theta \geq 0$ and $V_{\mathsf{sw}} > 0$, from Lemma A.1.1 let $\pi_\theta^{\kappa_\theta}$ be the steady state distribution of the Markov chain $\mathsf{MC}(\xi_\theta, \kappa_\theta)$. Let $V^{\theta,\kappa_\theta,V_{\mathsf{sw}}}$ and $V_{\mathsf{st}}^{\theta,\kappa_\theta,V_{\mathsf{sw}}}$ be the value functions satisfying the following Bellman equation:

$$\begin{aligned} V^{\theta,\kappa_\theta,V_{\mathsf{sw}}}(z,n) &= F_\theta(z,n) + \gamma \max\{V_{\mathsf{st}}^{\theta,\kappa_\theta,V_{\mathsf{sw}}}(z,n), V_{\mathsf{sw}}\} \\ V_{\mathsf{st}}^{\theta,\kappa_\theta,V_{\mathsf{sw}}}(z,n) &= \mathbf{E}[V^{\theta,\kappa_\theta,V_{\mathsf{sw}}}(Z_\tau, N_\tau)|z,n;\xi_\theta,\kappa_\theta]. \end{aligned} \tag{B.1}$$

Let

$$\mathcal{V}_{\mathsf{arr}}^\theta(\kappa_\theta, V_{\mathsf{sw}}) \triangleq \sum_{z\in\mathbb{Z},n\in\mathbb{N}_0} V_{\mathsf{st}}^{\theta,\kappa_\theta,V_{\mathsf{sw}}}(z,n+1)\pi_\theta^{\kappa_\theta}(z,n)$$

and

$$\phi_\theta(\kappa_\theta, V_{\mathsf{sw}}) \triangleq \mathcal{V}_{\mathsf{arr}}^\theta(\kappa_\theta, V_{\mathsf{sw}}) - V_{\mathsf{sw}}.$$

Lemma B.1.1. *The function $\phi_\theta(\kappa_\theta, V_{\mathsf{sw}})$ is (i) jointly continuous in κ_θ and V_{sw}, (ii) for each κ_θ, $\phi_\theta(\kappa_\theta, V_{\mathsf{sw}})$ is strictly decreasing in V_{sw}; and (iii) for each V_{sw}, $\phi_\theta(\kappa_\theta, V_{\mathsf{sw}})$ is strictly decreasing in κ_θ.*

Proof. Continuity is straightforward using similar arguments as in the proof of Lemma A.1.7 and Lemma A.1.13 in Appendix A, and we omit the details here.

To show $\phi_\theta(\kappa_\theta, V_{\mathsf{sw}})$ is decreasing in V_{sw}, given any $\kappa_\theta \geq 0$, let $V^{(0)}, V'^{(0)} : \mathbb{Z}\times\mathbb{N}_0 \to \mathbb{R}^+$ be defined as

$$V^{(0)}(z,n) = V'^{(0)}(z,n) = 0, \quad \text{for all} z,n. \tag{B.2}$$

For any $\delta > 0$, define the two sequences $V^{(i)}, V'^{(i)} : \mathbb{Z} \times \mathbb{N}_0 \to \mathbb{R}^+$, $i = 1, 2, \ldots$ inductively as:

$$\begin{aligned} V^{(i)}(z,n) =& F_\theta(z,n) + \gamma \max\{\mathbf{E}[V^{(i-1)}(Z_\tau, N_\tau)|z,n], V_{\mathsf{sw}}\}, \\ V'^{(i)}(z,n) =& F_\theta(z,n) + \gamma \max\{\mathbf{E}[V'^{(i-1)}(Z_\tau, N_\tau)|z,n], V_{\mathsf{sw}} + \delta\}, \end{aligned}$$

for all $z \in \mathbb{Z}, n \in \mathbb{N}_0$.

Convergence of value iteration implies $\{V^{(i)}\}_{i=0}^\infty$ and $\{V'^{(i)}\}_{i=1}^\infty$ converge uniformly to $V, V' : \mathbb{Z} \times \mathbb{N}_0 \to \mathbb{R}^+$ respectively where V and V' satisfy

$$\begin{aligned} V(z,n) =& F_\theta(z,n) + \gamma \max\{\mathbf{E}[V(Z_\tau, N_\tau)|z,n], V_{\mathsf{sw}}\}, \\ V'(z,n) =& F_\theta(z,n) + \gamma \max\{\mathbf{E}[V'(Z_\tau, N_\tau)|z,n], V_{\mathsf{sw}} + \delta\}, \end{aligned}$$

for all $z \in \mathbb{Z}, n \in \mathbb{N}_0$.

It is straightforward to observe for all $i \geq 0$, for all $z \in \mathbb{Z}$ and $n \in \mathbb{N}_0$, $V'^{(i)}(z,n) \geq V^{(i)}(z,n)$. Next, we show inductively $\sup_{z\in\mathbb{Z}, n\in\mathbb{N}_0} V'^{(i)}(z,n) - V^{(i)}(z,n) \leq \gamma\delta$ for all $i \geq 0$. The claim holds trivially for $i = 0$. Given $\sup_{z\in\mathbb{Z}, n\in\mathbb{N}_0} V'^{(i)}(z,n) - V^{(i)}(z,n) \leq \gamma\delta$ for some $i \geq 0$, we have for any $z \in \mathbb{Z}$ and $n \in \mathbb{N}_0$,

$$\begin{aligned} V'^{(i+1)}(z,n) =& F_\theta(z,n) + \gamma \max\{\mathbf{E}[V'^{(i)}(Z_\tau, N_\tau)|z,n], V_{\mathsf{sw}}\} \\ \leq& F_\theta(z,n) + \gamma \max\{\mathbf{E}[V^{(i)}(Z_\tau, N_\tau) + \gamma\delta|z,n], V_{\mathsf{sw}}\} \\ \leq& F_\theta(z,n) + \gamma \max\{\mathbf{E}[V^{(i)}(Z_\tau, N_\tau)|z,n], V_{\mathsf{sw}}\} + \gamma\delta \\ =& V^{(i+1)}(z,n) + \gamma\delta, \end{aligned}$$

which completes the induction step. Along with the uniform convergence of $\{V^{(i)}\}_{i=0}^\infty$ and $\{V'^{(i)}\}_{i=1}^\infty$, we have $0 \leq V'(z,n) - V(z,n) \leq \gamma\delta, \quad \forall z \in \mathbb{Z}, n \in \mathbb{N}_0$.

Since $V_{\mathsf{st}}^{\theta,\kappa_\theta,V_{\mathsf{sw}}+\delta}(z,n) = \mathbf{E}[V'(Z_\tau, N_\tau)|z,n]$ and $V_{\mathsf{st}}^{\theta,\kappa_\theta,V_{\mathsf{sw}}}(z,n) = \mathbf{E}[V(Z_\tau, N_\tau)|z,n]$, we have

$$\sup_{z\in\mathbb{Z},n\in\mathbb{N}_0} V_{\mathsf{st}}^{\theta,\kappa_\theta,V_{\mathsf{sw}}+\delta}(z,n) - V_{\mathsf{st}}^{\theta,\kappa_\theta,V_{\mathsf{sw}}}(z,n) \le \gamma\delta,$$

and therefore

$$\mathcal{V}_{\mathsf{arr}}^{\theta}(\kappa_\theta, V_{\mathsf{sw}} + \delta) - \mathcal{V}_{\mathsf{arr}}^{\theta}(\kappa_\theta, V_{\mathsf{sw}}) \le \gamma\delta.$$

Finally, we have

$$\begin{aligned}\phi_\theta(\kappa_\theta, V_{\mathsf{sw}} + \delta) - \phi_\theta(\kappa_\theta, V_{\mathsf{sw}}) =& \mathcal{V}_{\mathsf{arr}}^{\theta}(\kappa_\theta, V_{\mathsf{sw}} + \delta) - \mathcal{V}_{\mathsf{arr}}^{\theta}(\kappa_\theta, V_{\mathsf{sw}}) - \delta \\ \le& \gamma\delta - \delta < 0,\end{aligned}$$

which completes the proof of monotonicity in V_{sw}.

Next, we show $\phi_\theta(\kappa_\theta, V_{\mathsf{sw}})$ is decreasing κ_θ given any V_{sw}. The argument is similar to that of showing monotonicity in V_{sw}, by induction on the value iteration procedure. Consider κ_θ and κ'_θ where $\kappa_\theta < \kappa'_\theta$. Define $V^{(0)}, V'^{(0)} : \mathbb{Z}\times\mathbb{N}_0 \to \mathbb{R}^+$ as in (B.2) and $V^{(i)}, V'^{(i)} : \mathbb{Z}\times\mathbb{N}_0 \to \mathbb{R}^+$, $i = 1, 2, \ldots$ inductively as:

$$\begin{aligned}V^{(i)}(z,n) =& F_\theta(z,n) + \gamma \max\{\mathbf{E}[V^{(i-1)}(Z_\tau, N_\tau)|z,n; (Z_t, N_t) \sim \mathsf{MC}(\xi_\theta, \kappa_\theta)], V_{\mathsf{sw}}\},\\ V'^{(i)}(z,n) =& F_\theta(z,n) + \gamma \max\{\mathbf{E}[V'^{(i-1)}(Z_\tau, N_\tau)|z,n; (Z_t, N_t) \sim \mathsf{MC}(\xi_\theta, \kappa'_\theta)], V_{\mathsf{sw}}\}.\end{aligned}$$

Since $F(z,\cdot)$ is decreasing and non-constant, it is straightforward to show for all i, for each z, $V^{(i)}(z,\cdot)$ and $V'^{(i)}(z,\cdot)$ is decreasing and non-constant. We show inductively for all i, for each z,n, $V'^{(i)}(z,n) \le V^{(i)}(z,n)$. The claim holds for $i = 0$ trivially. For the induction step, assume the claim holds for some $i \ge 0$. For each z,n, let $(Z_t^1, N_t^1) \sim \mathsf{MC}(\xi_\theta, \kappa_\theta)$ and $(Z_t^2, N_t^2) \sim \mathsf{MC}(\xi_\theta, \kappa'_\theta)$ and $(Z_0^j, N_0^j) = (z,n)$ for $j = 1, 2$, and we couple Z_1, Z_2. From Lemma A.1.18, for any $t \geq 0$, since $\kappa'_\theta > \kappa_\theta$, we have $Z_t^1 = Z_t^2$ and $N_t^1 \preccurlyeq_{\mathsf{sd}} N_t^2$, and hence $V'^{(i)}(Z_\tau^2, N_\tau^2) \preccurlyeq_{\mathsf{sd}} V^{(i)}(Z_\tau^1, n_\tau^1)$. Therefore we have $V^{(i+1)}(z,n) \leq V'^{(i+1)}(z,n)$.

It follows that $V'(z,n) \leq V(z,n)$ for each z, n where $V, V' : \mathbb{Z}\times\mathbb{N}_0 \to \mathbb{R}^+$ are defined as

$$V(z,n) = F_\theta(z,n) + \gamma \max\{\mathbf{E}[V(Z_\tau, N_\tau)|z,n;(Z_t,N_t) \sim \mathsf{MC}(\xi_\theta, \kappa_\theta)], V_{\mathsf{sw}}\},$$

$$V'(z,n) = F_\theta(z,n) + \gamma \max\{\mathbf{E}[V'(Z_\tau, N_\tau)|z,n;(Z_t,N_t) \sim \mathsf{MC}(\xi_\theta, \kappa'_\theta)], V_{\mathsf{sw}}\}.$$

It then follows that

$$\begin{aligned} V_{\mathsf{st}}^{\theta,\kappa'_\theta}(z,n) =& \mathbf{E}[V'(Z_\tau, N_\tau)|z,n;(Z_t,N_t) \sim \mathsf{MC}(\xi_\theta, \kappa'_\theta)] \\ \leq& \mathbf{E}[V(Z_\tau, N_\tau)|z,n;(Z_t,N_t) \sim \mathsf{MC}(\xi_\theta, \kappa'_\theta)] \\ \leq& \mathbf{E}[V(Z_\tau, N_\tau)|z,n;(Z_t,N_t) \sim \mathsf{MC}(\xi_\theta, \kappa_\theta)] \\ =& V_{\mathsf{st}}^{\theta,\kappa_\theta}(z,n). \end{aligned}$$

Therefore we have $\mathcal{V}_{\mathsf{arr}}^\theta(\kappa'_\theta, V_{\mathsf{sw}}) \leq \mathcal{V}_{\mathsf{arr}}^\theta(\kappa_\theta, V_{\mathsf{sw}})$.

To show strict monotonicity of $\mathcal{V}_{\mathsf{arr}}^\theta(,\cdot, V_{\mathsf{sw}})$, we can prove by contradiction using a similar approach as in the proof of Lemma A.1.5, and we omit the detail here. Since $\mathcal{V}_{\mathsf{arr}}^\theta(,\cdot, V_{\mathsf{sw}})$ is strictly decreasing, $\phi_\theta(\cdot, V_{\mathsf{sw}})$ is also strictly decreasing, completing the proof.

□

Next, we would define a function $\psi(V_{\mathsf{sw}})$ and show that ψ is strictly decreasing, and there exists $\overline{\mathsf{V}}$ and $\underline{\mathsf{V}}$ such that $\psi(\overline{\mathsf{V}}) \le \beta$ and $\psi(\underline{\mathsf{V}}) \ge \beta$ hence there must exist a unique V_{sw} such that $\psi(V_{\mathsf{sw}}) = \beta$.

Define $\overline{\mathsf{V}}'$ as

$$\overline{\mathsf{V}}' \triangleq \max_{\theta \in \Theta} \frac{\|F_\theta\|}{1-\gamma}.$$

With a similar argument in Appendix A.1.4, we can show for each $\theta \in \Theta$, and for any $\kappa \ge 0$ and $V_{\mathsf{sw}} > 0$, $\|\mathcal{V}_{\mathsf{st}}^{\theta,\kappa,V_{\mathsf{sw}}}\|_\infty \le \overline{\mathsf{V}}'$. Therefore $\mathcal{V}_{\mathsf{arr}}^{\theta}(\kappa, V_{\mathsf{sw}}) \le \overline{\mathsf{V}}'$, and we have

$$\phi_\theta(\kappa, \overline{\mathsf{V}}') = \mathcal{V}_{\mathsf{arr}}^{\theta}(\kappa, \overline{\mathsf{V}}') - \overline{\mathsf{V}}' \le 0. \tag{B.3}$$

Specifically, consider $\kappa_\theta = \beta\lambda/p_\theta$. From (B.3) we have $\phi_\theta(\beta\lambda/p_\theta, \overline{\mathsf{V}}') \le 0$. On the other hand, it is trivially to observe $\phi_\theta(\kappa, 0) > 0$ for any $\kappa > 0$ and hence $\phi_\theta(\beta\lambda/p_\theta, 0) > 0$. From Lemma B.1.1, for any $\theta \in \Theta$, $\phi_\theta(\beta\lambda/p_\theta, \cdot)$ is a strictly decreasing function, hence there exists a unique $\underline{\mathsf{V}}_\theta \in (0, \overline{\mathsf{V}}']$ such that

$$\phi_\theta\left(\frac{\beta\lambda}{p_\theta}, \underline{\mathsf{V}}_\theta\right) = 0.$$

Next, consider $\kappa = \beta\lambda(1-\gamma)$. For each $\theta \in \Theta$, since $\phi_\theta(\cdot, \underline{\mathsf{V}}_\theta)$ is decreasing, we have $\phi_\theta(\beta\lambda(1-\gamma), \underline{\mathsf{V}}_\theta) \ge \phi_\theta(\beta\lambda/p_\theta, \underline{\mathsf{V}}_\theta) = 0$. Also, from (B.3) we have $\phi_\theta(\beta\lambda(1-\gamma), \overline{\mathsf{V}}') \le 0$. Since $\phi_\theta(\beta\lambda(1-\gamma), \cdot)$ is strictly decreasing, there exists unique $\overline{\mathsf{V}}_\theta \in [\underline{\mathsf{V}}_\theta, \overline{\mathsf{V}}']$ such that $\phi_\theta(\beta\lambda(1-\gamma), \overline{\mathsf{V}}_\theta) = 0$.

Let $\underline{\mathsf{V}} \triangleq \max_{\theta\in\Theta} \underline{\mathsf{V}}_\theta$ and $\overline{\mathsf{V}} \triangleq \max_{\theta\in\Theta} \overline{\mathsf{V}}_\theta$. Note that there exists $\theta' \in \Theta$ such that $\underline{\mathsf{V}}_{\theta'} = \underline{\mathsf{V}}$, and hence $\underline{\mathsf{V}} = \underline{\mathsf{V}}_{\theta'} \le \overline{\mathsf{V}}_{\theta'} \le \overline{\mathsf{V}}$.

For any $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$, and for each $\theta \in \Theta$, we define $\kappa_\theta(V_{\mathsf{sw}})$ as follows. First, notice that since $\phi_\theta(\beta\lambda/p_\theta, \cdot)$ is a decreasing function, we have

$$\phi_\theta\left(\frac{\beta\lambda}{p_\theta}, V_{\mathsf{sw}}\right) \leq \phi_\theta\left(\frac{\beta\lambda}{p_\theta}, \underline{\mathsf{V}}\right) \leq \phi_\theta\left(\frac{\beta\lambda}{p_\theta}, \underline{\mathsf{V}}_\theta\right) = 0.$$

On the other hand, if $V_{\mathsf{sw}} \geq \overline{\mathsf{V}}_\theta$, we have

$$\phi_\theta(\beta\lambda(1-\gamma), V_{\mathsf{sw}}) \geq \phi_\theta(\beta\lambda(1-\gamma), \overline{\mathsf{V}}_\theta) = 0.$$

By Lemma B.1.1, $\phi_\theta(\cdot, V_{\mathsf{sw}})$ is strictly decreasing, so there exists a unique $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda/p_\theta]$ such that $\phi_\theta(\kappa, V_{\mathsf{sw}}) = 0$, and we let $\kappa_\theta(V_{\mathsf{sw}}) = \kappa$. Otherwise $\overline{\mathsf{V}}_\theta < V_{\mathsf{sw}} \leq \overline{\mathsf{V}}$, and for all $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda/p_\theta]$, it must be that $\phi_\theta(\kappa, V_{\mathsf{sw}}) \leq \phi_\theta(\beta\lambda(1-\gamma), V_{\mathsf{sw}}) < \phi_\theta(\beta\lambda(1-\gamma), \overline{\mathsf{V}}_\theta) = 0$. In this case we let $\kappa_\theta(V_{\mathsf{sw}}) = \beta\lambda(1-\gamma)$.

Having defined $\kappa_\theta(V_{\mathsf{sw}})$, we let

$$\beta_\theta(V_{\mathsf{sw}}) \triangleq \sum_{z\in\mathbb{Z}, n\in\mathbb{N}_0} n\pi^{\kappa_\theta(V_{\mathsf{sw}})}(z, n).$$

We then let

$$\psi(V_{\mathsf{sw}}) \triangleq \sum_{\theta\in\Theta} \beta_\theta(V_{\mathsf{sw}}) p_\theta.$$

Consider $\theta' \in \arg\max_{\theta\in\Theta} \underline{\mathsf{V}}_\theta$, we have

$$\psi(\underline{\mathsf{V}}) = \sum_{\theta\in\Theta} \beta_\theta(\underline{\mathsf{V}}) p_\theta \geq \beta_{\theta'}(\underline{\mathsf{V}}) p_{\theta'}. \tag{B.4}$$

Since $\theta' \in \arg\max_{\theta\in\Theta} \underline{\mathsf{V}}_\theta$, we have $V_{\mathsf{sw}}(\theta') = \underline{\mathsf{V}}$ and $\phi_{\theta'}(\beta\lambda/p_{\theta'}, \underline{\mathsf{V}}) = \phi_{\theta'}(\beta\lambda/p_{\theta'}, V_{\mathsf{sw}}(\theta')) = 0$. Therefore $\kappa_{\theta'}(\underline{\mathsf{V}}) = \beta\lambda/p_{\theta'}$.

Since $\beta_{\theta'}(\underline{\mathsf{V}})$ is the expected number of agents of $\mathsf{MC}(\xi_\theta, \kappa_{\theta'}(\underline{\mathsf{V}}))$ in steady state, from Lemma A.1.19, we can lower bound $\beta_{\theta'}(\underline{\mathsf{V}})$ with the expected number of agents in an $M/M/\infty$ queue $\mathsf{M/M/}\infty(\kappa_{\theta'}(\underline{\mathsf{V}}), \lambda)$, which gives

$$\beta_{\theta'}(\underline{\mathsf{V}}) \geq \frac{\kappa_{\theta'}(\underline{\mathsf{V}})}{\lambda} = \frac{\beta}{p_{\theta'}}. \tag{B.5}$$

Combining (B.4) and (B.5), we have $\psi(\underline{\mathsf{V}}) \geq \beta$.

On the other hand, for each $\theta \in \Theta$, either $\overline{\mathsf{V}} = \overline{\mathsf{V}}_\theta$ or $\overline{\mathsf{V}} > \overline{\mathsf{V}}_\theta$. If $\overline{\mathsf{V}} = \overline{\mathsf{V}}_\theta$, we have $\phi(\beta\lambda(1-\gamma), \overline{\mathsf{V}}) = \phi(\beta\lambda(1-\gamma), \overline{\mathsf{V}}_\theta) = 0$ and $\kappa_\theta(\overline{\mathsf{V}}) = \beta\lambda(1-\gamma)$; otherwise $\overline{\mathsf{V}} > \overline{\mathsf{V}}_\theta$, and for all $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda/p_\theta]$, $\phi_\theta(\kappa, \overline{\mathsf{V}}) < \phi_\theta(\kappa, \overline{\mathsf{V}}_\theta) \leq \phi_\theta(\beta\lambda(1-\gamma), \overline{\mathsf{V}}_\theta) = 0$ hence $\kappa_\theta(\overline{\mathsf{V}}) = \beta\lambda(1-\gamma)$. Therefore for each $\theta \in \Theta$, $\kappa(\overline{\mathsf{V}}) = \beta\lambda(1-\gamma)$.

For each $\theta \in \Theta$, since $\beta_\theta(\overline{\mathsf{V}})$ is the expected number of agents of $\mathsf{MC}(\xi_\theta, \kappa_\theta(\overline{\mathsf{V}})$, by Lemma A.1.19 we can upper bound $\beta_\theta(\overline{\mathsf{V}})$ with the expected number of agents of an $M/M/\infty$ queue $\mathsf{M/M/}\infty(\kappa_\theta(\overline{\mathsf{V}}), \lambda(1-\gamma))$, which gives

$$\psi(\overline{\mathsf{V}}) \leq \sum_{\theta \in \Theta} \frac{\kappa_\theta(\overline{\mathsf{V}})}{\lambda(1-\gamma)} p_\theta = \sum_{\theta \in \Theta} \beta p_\theta = \beta.$$

In the next lemma, we show $\psi(V_{\mathsf{sw}})$ is a continuous and strictly decreasing function on $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$.

Lemma B.1.2. *ψ is continuous and strictly decreasing on* $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$.

Proof. We first show ψ is strictly decreasing. Consider $V_{\mathsf{sw}}, V'_{\mathsf{sw}}$ that $\underline{\mathsf{V}} \leq V'_{\mathsf{sw}} < V_{\mathsf{sw}} \leq \overline{\mathsf{V}}$, and we would like to show for any $\theta \in \Theta$, $\kappa_\theta(V_{\mathsf{sw}}) \leq \kappa_\theta(V'_{\mathsf{sw}})$, and there exists $\theta' \in \Theta$ such that $\kappa_\theta(V_{\mathsf{sw}}) < \kappa_\theta(V'_{\mathsf{sw}})$.

For each θ, if for any $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda/p_\theta]$, $\phi_\theta(\kappa, V'_{\mathsf{sw}}) < 0$, then $\kappa_\theta(V'_{\mathsf{sw}}) = \beta\lambda(1-\gamma)$, and for any $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda/p_\theta]$, since $\phi(\kappa, \cdot)$ is strictly decreasing, we have $\phi_\theta(\kappa, V_{\mathsf{sw}}) < \phi_\theta(\kappa, V'_{\mathsf{sw}}) < 0$, which gives $\kappa_\theta(V_{\mathsf{sw}}) = \beta\lambda(1-\gamma) = \kappa_\theta(V'_{\mathsf{sw}})$. On the other hand, assume there exists $\kappa_\theta(V'_{\mathsf{sw}}) \in [\beta\lambda(1-\gamma), \beta\lambda/p_\theta]$ such that $\phi_\theta(\kappa_\theta(V'_{\mathsf{sw}}), V'_{\mathsf{sw}}) = 0$. In this case, if for all $\kappa \in [\beta\lambda(1-\gamma), \beta\lambda/p_\theta]$, $\phi_\theta(\kappa, V_{\mathsf{sw}}) < 0$, then $\kappa_\theta(V_{\mathsf{sw}}) = \beta\lambda(1-\gamma)$ and clearly $\kappa_\theta(V_{\mathsf{sw}}) \leq \kappa_\theta(V'_{\mathsf{sw}})$; otherwise there exists $\kappa_\theta(V_{\mathsf{sw}}) \in [\beta\lambda(1-\gamma), \beta\lambda/p_\theta]$ such that $\phi_\theta(\kappa_\theta(V_{\mathsf{sw}}), V_{\mathsf{sw}}) = 0$, and we have

$$\phi_\theta(\kappa_\theta(V_{\mathsf{sw}}), V_{\mathsf{sw}}) = 0 = \phi_\theta(\kappa_\theta(V'_{\mathsf{sw}}), V'_{\mathsf{sw}}) > \phi_\theta(\kappa_\theta(V'_{\mathsf{sw}}), V_{\mathsf{sw}}). \tag{B.6}$$

Since $\phi_\theta(\cdot, V_{\mathsf{sw}})$ is decreasing, we have $\kappa_\theta(V_{\mathsf{sw}}) < \kappa_\theta(V'_{\mathsf{sw}})$. Therefore, for each $\theta \in \Theta$, we have $\kappa_\theta(V_{\mathsf{sw}}) \leq \kappa_\theta(V'_{\mathsf{sw}})$.

For $\theta' \in \arg\max_{\theta\in\Theta} \overline{\mathsf{V}}_\theta$, since $V'_{\mathsf{sw}} < \overline{\mathsf{V}}$, we have $\phi_{\theta'}(\beta\lambda(1-\gamma), V'_{\mathsf{sw}}) > \phi_{\theta'}(\beta\lambda(1-\gamma), \overline{\mathsf{V}}) = \phi_{\theta'}(\beta\lambda(1-\gamma), \overline{\mathsf{V}}_\theta) = 0$, hence it must hold that $\phi_\theta(\kappa_\theta(V'_{\mathsf{sw}}), V'_{\mathsf{sw}}) = 0$. If $V_{\mathsf{sw}} < \overline{\mathsf{V}}$, we can similarly show $\phi_{\theta'}(\kappa_{\theta'}(V_{\mathsf{sw}}), V_{\mathsf{sw}}) = 0$; otherwise $V_{\mathsf{sw}} = \overline{\mathsf{V}} = \overline{\mathsf{V}}_{\theta'}$, we also have $\kappa_{\theta'}(V_{\mathsf{sw}}) = \beta\lambda(1-\gamma)$ and $\phi_{\theta'}(\kappa_{\theta'}(V_{\mathsf{sw}}), V_{\mathsf{sw}}) = 0$. Thus,

$$\phi_{\theta'}(\kappa_{\theta'}(V_{\mathsf{sw}}), V_{\mathsf{sw}}) = 0 = \phi_{\theta'}(\kappa_{\theta'}(V'_{\mathsf{sw}}), V'_{\mathsf{sw}}) > \phi_{\theta'}(\kappa_{\theta'}(V'_{\mathsf{sw}}), V_{\mathsf{sw}}), \tag{B.7}$$

and therefore $\kappa_{\theta'}(V_{\mathsf{sw}}) < \kappa_{\theta'}(V'_{\mathsf{sw}})$.

By Lemma A.1.5, $\beta_\theta(V'_{\mathsf{sw}}) \leq \beta_\theta(V_{\mathsf{sw}})$ for all $\theta \in \Theta$ and $\beta_{\theta'}(V'_{\mathsf{sw}}) < \beta_{\theta'}(V_{\mathsf{sw}})$ for $\theta' \in \arg\max_{\theta\in\Theta} \overline{\mathsf{V}}_\theta$. Therefore, we have $\psi(V'_{\mathsf{sw}}) < \psi(V_{\mathsf{sw}})$ so ψ is strictly decreasing.

Continuity is straightforward. For each θ, for $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}_\theta]$, since $\kappa_\theta(V_{\mathsf{sw}})$ uniquely maximizes the function $f(\cdot) \triangleq -|\phi_\theta(\cdot, V_{\mathsf{sw}})|$, by Berge's maximum theorem, $\kappa_\theta(V_{\mathsf{sw}})$ is continuous with respect to V_{sw}, and since $\kappa_\theta(V_{\mathsf{sw}}) = \beta\lambda(1-\gamma)$ for all

$V_{\mathsf{sw}} \in [\overline{\mathsf{V}}_\theta, \overline{\mathsf{V}}]$, $\kappa_\theta(V_{\mathsf{sw}})$ is continous on $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$. By Lemma A.1.7, $\beta_\theta(V_{\mathsf{sw}})$ is also continuous in $V_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$, and since Θ is finite $\psi(V_{\mathsf{sw}})$ is also continuous.

□

Since $\psi(V_{\mathsf{sw}})$ is continuous and strictly decreasing on $[\underline{\mathsf{V}}, \overline{\mathsf{V}}]$, and $\psi(\overline{\mathsf{V}}) \leq \beta$ and $\psi(\underline{\mathsf{V}}) \geq \beta$, there exists unique $V^*_{\mathsf{sw}} \in [\underline{\mathsf{V}}, \overline{\mathsf{V}}]$ such that $\psi(V^*_{\mathsf{sw}}) = \beta$. Let $V_{\mathsf{sw}}(\xi) = V^*_{\mathsf{sw}}$. For each $\theta \in \Theta$, let $\kappa_\theta(\xi) = \kappa(V^*_{\mathsf{sw}})$ as defined above, let $\beta_\theta(\xi) = \sum_{z\in\mathbb{Z}, n\in\mathbb{N}_0} n\pi^{\kappa_\theta(V_{\mathsf{sw}})}(z, n)$, and let $V^\theta_{\mathsf{arr}}(\xi) = \mathcal{V}^\theta_{\mathsf{arr}}(\kappa_\theta(V^*_{\mathsf{sw}}), V^*_{\mathsf{sw}})$.

Clearly $\{\kappa_\theta(\xi)\}_{\theta\in\Theta}$, $\{\beta_\theta(\xi)\}_{\theta\in\Theta}$, $\{V^\theta_{\mathsf{arr}}(\xi)\}_{\theta\in\Theta}$ and $V_{\mathsf{sw}}(\xi)$ so defined satisfies (3.1), (3.2) and (3.4) $\sim$ (3.7). Also, notice that for each $\theta \in \Theta$, if $\kappa_\theta(V^*_{\mathsf{sw}}) = \beta\lambda(1-\gamma)$, (3.8) holds directly. Otherwise, $\kappa_\theta(V^*_{\mathsf{sw}}) > \beta\lambda(1-\gamma)$, and it must hold that $\phi_\theta(\kappa_\theta(V^*_{\mathsf{sw}}), V^*_{\mathsf{sw}}) = 0$. In this case, since $\phi_\theta(\kappa_\theta(V^*_{\mathsf{sw}}), V^*_{\mathsf{sw}}) = \mathcal{V}^\theta_{\mathsf{arr}}(\kappa_\theta(V^*_{\mathsf{sw}}), V^*_{\mathsf{sw}}) - V^*_{\mathsf{sw}} = V^\theta_{\mathsf{arr}}(\xi) - V_{\mathsf{sw}}(\xi)$, we have $V^\theta_{\mathsf{arr}}(\xi) = V_{\mathsf{sw}}(\xi)$ and (3.8) holds. To show (3.9) holds, for each $\theta \in \Theta$, first notice by definition of $\kappa_\theta(V^*_{\mathsf{sw}})$, $0 \geq \phi_\theta(\kappa_\theta(V^*_{\mathsf{sw}}), V^*_{\mathsf{sw}}) = \mathcal{V}^\theta_{\mathsf{arr}}(\kappa_\theta(V^*_{\mathsf{sw}}), V^*_{\mathsf{sw}}) - V^*_{\mathsf{sw}} = V^\theta_{\mathsf{arr}}(\xi) - V_{\mathsf{sw}}(\xi)$ hence $V^\theta_{\mathsf{arr}}(\xi) \leq V_{\mathsf{sw}}(\xi)$. Finally we show there must exist θ' such that $V^{\theta'}_{\mathsf{sw}}(\xi) = V_{\mathsf{sw}}(\xi)$. Consider $\theta' \in \arg\max_{\theta\in\Theta} \overline{\mathsf{V}}_\theta$, if $\kappa_{\theta'}(V^*_{\mathsf{sw}}) > \beta\lambda(1-\gamma)$ then by definition of $\kappa_{\theta'}(V^*_{\mathsf{sw}})$ we have $\phi_{\theta'}(\kappa_{\theta'}(V^*_{\mathsf{sw}}), V^*_{\mathsf{sw}}) = 0$ and hence $V^{\theta'}_{\mathsf{sw}}(\xi) = V_{\mathsf{sw}}(\xi)$. Otherwise $\kappa_{\theta'}(V^*_{\mathsf{sw}}) = \beta\lambda(1-\gamma)$, and it must be the case that $V^*_{\mathsf{sw}} = \overline{\mathsf{V}} = \overline{\mathsf{V}}_{\theta'}$, since if $V^*_{\mathsf{sw}} < \overline{\mathsf{V}}_{\theta'}$, we would have

$$\phi_{\theta'}(\beta\lambda(1-\gamma), V^*_{\mathsf{sw}}) > \phi_{\theta'}(\beta\lambda(1-\gamma), \overline{\mathsf{V}}_{\theta'}) = 0,$$

which implies $\kappa_{\theta'}(V^*_{\mathsf{sw}}) > \beta\lambda(1-\gamma)$, leading to a contradiction. Now since $V^*_{\mathsf{sw}} = \overline{\mathsf{V}} = \overline{\mathsf{V}}_{\theta'}$, we have $\kappa_{\theta'}(V^*_{\mathsf{sw}}) = \beta\lambda(1-\gamma)$ and $\phi_{\theta'}(\kappa_{\theta'}(V^*_{\mathsf{sw}}), V^*_{\mathsf{sw}}) = 0$ and $V^{\theta'}_{\mathsf{sw}}(\xi) = V_{\mathsf{sw}}(\xi)$. Thus we always have $V^{\theta'}_{\mathsf{sw}}(\xi) = V_{\mathsf{sw}}(\xi)$ and (3.9) holds.

We finished STEP 1. Next, we proceed to STEP 2. First, notice for any Markovian strategy ξ and $V_{\mathsf{sw}} > 0$, for each θ, we have $\kappa_\theta(V_{\mathsf{sw}})$ as defined above satisfies $\kappa_\theta(V_{\mathsf{sw}}) \le \beta\lambda/p_\theta$.

We let

$$\overline{\mathsf{V}}^* \triangleq \frac{1}{1-\gamma}\|F_\theta\|_\infty,$$
$$\underline{\mathsf{V}}_\theta^* \triangleq \exp\left(-\frac{\beta}{(1-\gamma)p_\theta}\right)\sum_{(z,n)\in\mathbb{S}}\frac{\beta^n(1-\gamma)^n}{(1+\beta/p_\theta+\Psi_\theta)^{n+1}(n+1)!}\pi_{\mathsf{res}}^\theta(z)F(z,n+1) > 0, \forall\theta\in\Theta,$$
$$\underline{\mathsf{V}}^* \triangleq \min_{\theta\in\Theta}\left\{\underline{\mathsf{V}}_\theta^*\right\} > 0,$$

where $\Psi_\theta = \frac{1}{\lambda}\max_{z\in\mathbb{Z}}\sum_{y\neq z}\mu_{zy}^\theta \in (0,\infty)$, and $\pi_{\mathsf{res}}^\theta$ is the steady state distribution of the resource process in a type-θ location.

For each $\theta\in\Theta$, using a similar argument of the proof of Lemma A.1.10, we can show for any Markovian strategy profile ξ, any $V_{\mathsf{sw}} > 0$ and any $\kappa_\theta \le \beta\lambda/p_\theta$, we have $V_{\mathsf{arr}}^\theta(\xi) \le \overline{\mathsf{V}}^*$. Hence we have $V_{\mathsf{sw}}(\xi) \le \overline{\mathsf{V}}^*$ for any ξ.

On the other hand, for each $\theta\in\Theta$, using a similar argument of the proof of Lemma A.1.11, we can show for any Markovian strategy profile ξ, any $V_{\mathsf{sw}} > 0$ and any $\kappa_\theta \le \beta\lambda/p_\theta$, we have $V_{\mathsf{arr}}^\theta(\xi) \ge \underline{\mathsf{V}}_\theta^*$. Hence we have $V_{\mathsf{sw}}(\xi) \ge \underline{\mathsf{V}}^*$ for any ξ.

Similar to Appendix A.1.5, we define, for each $\theta\in\Theta$,

$$K_0(\theta) \triangleq \inf\left\{m : F_\theta(z,n) < \frac{(1-\gamma)^2}{2}\underline{\mathsf{V}}^* \text{ for all } z\in\mathbb{Z} \text{ and } n\ge m\right\},$$

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

and $K_{\max}(\theta) \triangleq \max\{4K_0(\theta)^2 + 1, K_1(\theta)\}$. We define the set $\widehat{\Pi}_\theta$ as:

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

where Π still denotes the set of all Markovian Strategies.

Using similar arguments to the proof of Lemma A.1.12, we can show $\mathcal{X}_\theta(\xi) \subseteq \widehat{\Pi}_\theta$ for each $\theta \in \Theta$, for any markovian strategy profile ξ. Let $\widehat{\Pi} \triangleq \times_{\theta\in\Theta}\widehat{\Pi}_\theta$, and we have for any Markovian strategy profile ξ, $\mathcal{X}(\xi) \in \widehat{\Pi}$.

Next, note that $\xi \mapsto V_{\mathsf{sw}}(\xi)$ is continuous since $V_{\mathsf{sw}}(\xi)$ maximizes $-|\psi(\cdot) - \beta|$ so Berge's maximum theorem implies continuity. Along with arguments similar to that of Appendix A.1.6, we can show $\xi \mapsto \mathcal{X}_\theta(\xi)$ is upper hemicontinuous for each $\theta \in \Theta$. Therefore $\xi \mapsto \mathcal{X}(\xi)$ is also upper hemicontinuous. By Fan-Glicksberg fixed point theorem, the correspondence $\xi \mapsto \mathcal{X}(\xi)$ has a fixed point, which completes the proof of the theorem.

Remark. This proof relies on the assumption that F_θ is a decreasing resource sharing function for each $\theta \in \Theta$, which ensures strict monotoncity of ψ, hence there exists a unique V^*_{sw} such that $\psi(V^*_{\mathsf{sw}}) = \beta$. When this assumption does not hold, there might exist more than one V^*_{sw} such that $\psi(V^*_{\mathsf{sw}}) = \beta$. In this case, the correspondence $\xi \mapsto \times_\theta\mathcal{X}_\theta(\xi)$ is not well defined. There may exist other ways to construct a correspondence whose fixed points correspond to mean field equilibria when not all F_θ is decreasing, and by proving existence of a fixed point for that correspondence one can show existence of mean field equilibria for more general resouce sharing functions. Currently, this remains an open problem.

APPENDIX C

APPENDIX OF CHAPTER 4

C.1 Proofs

C.1.1 Proof of Lemma 4.3.1

Proof. Let ϕ be an arbitrary persuasive straightforward signaling scheme, where for any $S \subseteq [N]$, $\phi(S|\theta)$ is the probability of recommending agents in S to move, condition on resource state is θ. We construct another straightforward signaling scheme ϕ' where $\phi'(S|1) = \phi(S|1)$ for all $S \subseteq [N]$; $\phi'(\emptyset|0) = 1$ and $\phi'(S|0) = 0$ for all $S \subseteq [N]$ and $S \neq \{\emptyset\}$. We would show ϕ' is also persuasive and achieves at least the same social welfare as ϕ. The expected social welfare under ϕ is

$$
\begin{aligned}
& \sum_{\theta=0,1,S\subseteq[N]} \mu(\theta)\phi(S|\theta)W(\theta, S) \\
= & \sum_{\theta=0,1,S\subseteq[N]} \phi(\theta, S)W(\theta, S) \\
= & \sum_{S\subseteq[N]} \phi(1, S)W(1, S) + \sum_{S\subseteq[N]} \phi(0, S)W(0, S) \\
= & \sum_{S\subseteq[N]} \phi(1, S)W(1, S) - \sum_{S\subseteq[N]} \phi(0, S) \sum_{i\in S} r(i) \\
\leq & \sum_{S\subseteq[N]} \phi(1, S)W(1, S) \\
= & \sum_{S\subseteq[N]} \phi'(1, S)W(1, S) + \phi'(0, \emptyset)W(0, \emptyset)
\end{aligned}
$$

$$= \sum_{\theta=0,1,S\subseteq[N]} \phi'(\theta, S)W(\theta, S),$$

which is the social welfare under ϕ'.

Next, we show (Σ, ϕ') satisfies the persuasive constraints (4.3) and (4.4). For each agent i, since $\phi'(S|0) = 0$ for $S \neq \emptyset$ and $\phi'(S|1) = \phi(S|1)$ for any subset S, we have

$$\begin{aligned}
&\sum_{\theta,S:i\in S} \phi'(\theta, S)(\theta F(|S|) - r(i)) \\
=&\sum_{S:i\in S} \phi'(1, S)(F(|S|) - r(i)) \\
=&\sum_{S:i\in S} \phi(1, S)(F(|S|) - r(i)) \\
\geq& - r(i)\sum_{S:i\in S} \phi(0, S) + \sum_{S:i\in S} \phi(1, S)(F(|S|) - r(i)) \\
=&\sum_{\theta,S:i\in S} \phi(\theta, S)(\theta F(|S|) - r(i)) \\
\geq& 0,
\end{aligned}$$

where the last inequality is because ϕ is persuasive. Therefore we have shown (4.3) holds for agent i.

On the other hand, for each agent i, we have

$$\begin{aligned}
&\sum_{\theta,S:i\notin S} \phi'(\theta, S)(\theta F(|S| + 1) - r(i)) \\
=& - r(i)\sum_{S:i\notin S} \phi'(0, S) + \sum_{S:i\notin S} \phi'(1, S)(F(|S| + 1) - r(i))
\end{aligned}$$

$$
\begin{aligned}
&= -r(i)\mu(0) + \sum_{S:i\notin S} \phi(1,S)(F(|S|+1) - r(i)) \\
&\leq -r(i)\sum_{S:i\notin S}\phi(0,S) + \sum_{S:i\notin S} \phi(1,S)(F(|S|+1) - r(i)) \\
&= \sum_{\theta,S:i\notin S} \phi(\theta,S)(\theta F(|S|+1) - r(i)) \\
&\leq 0,
\end{aligned}
$$

where the first inequality is because $\sum_{S:i\notin S}\phi(0,S) \leq \mu(0)$, and the last inequality holds because ϕ is persuasive. Thus we have shown (4.4) holds for agent i.

□

C.1.2 Proof of Lemma 4.3.2

Proof. Let $\mathcal{S}$ be the random subset of agents recommended to move under ϕ. For each agent i, let $X_i = \mathbb{1}\{i \in \mathcal{S}\}$ be the random variable denoting that agent i is recommended to move. Let $\mathbf{P}$ and $\mathbf{E}$ denote the probability measure and expectation induced by ϕ.

Since $p_{ik} = \mathbf{P}(|\mathcal{S}| = k, i \in \mathcal{S}|\theta = 1) = \mathbf{P}(|\mathcal{S}| = k|\theta = 1)\mathbf{P}(i \in \mathcal{S}||\mathcal{S}| = k, \theta = 1)$, and

$$
k = \mathbf{E}\left[\sum_{i=1}^{N} X_i \middle| |\mathcal{S}| = k, \theta = 1\right] = \sum_{i=1}^{N} \mathbf{E}[X_i||\mathcal{S}| = k, \theta = 1] = \sum_{i=1}^{N} \mathbf{P}(i \in \mathcal{S}||\mathcal{S}| = k, \theta = 1),
$$

we have

$$
\sum_{i=1}^{N} p_{ik} = \mathbf{P}(|\mathcal{S}| = k|\theta = 1)\sum_{i=1}^{N} \mathbf{P}(i \in \mathcal{S}||\mathcal{S}| = k, \theta = 1) = k\mathbf{P}(|\mathcal{S}| = k|\theta = 1),
$$

and therefore $\mathbf{P}(|\mathcal{S}| = k|\theta = 1) = \sum_{i=1}^{N} p_{ik}/k$.

The objective (4.2) of (LP.1) can be written as

$$\begin{aligned}
&\mu(0) \sum_{S \subseteq [N]} \phi(S|\theta) W(0, S) + \mu(1) \sum_{S \subseteq [N]} \phi(S|\theta) W(1, S) \\
=&\mu(0)\phi(\varnothing|0)W(0, \varnothing) + \mu(1) \sum_{S \subseteq [N]} \left(|S|F(|S|) - \sum_{i \in S} r(i) \right) \\
=&\mu(1) \sum_{k=1}^{N} \sum_{S:|S|=k} \phi(S|1) \left(kF(k) - \sum_{i \in S} r(i) \right) \\
=&\mu(1) \sum_{k=1}^{N} kF(k) \sum_{S:|S|=k} \phi(S|1) - \mu(1) \sum_{k=1}^{N} \sum_{S:|S|=k} \phi(S|1) \sum_{i \in S} r(i) \\
=&\mu(1) \sum_{k=1}^{N} kF(k) \mathbf{P}(|\mathcal{S}| = k|\theta = 1) - \mu(1) \sum_{k=1}^{N} \sum_{i=1}^{N} r(i) \sum_{S:|S|=k, i \in S} \phi(S|1) \\
=&\mu(1) \sum_{k=1}^{N} kF(k) \sum_{i=1}^{N} \frac{p_{ik}}{k} - \mu(1) \sum_{k=1}^{N} \sum_{i=1}^{N} r(i) p_{ik} \\
=&\mu(1) \sum_{k=1}^{N} \sum_{i=1}^{N} p_{ik}(F(k) - r(i)).
\end{aligned}$$

Next, we write the persuasive constraints in terms of the p_{ik}'s. For each agent i, the left hand side of (4.3) can be written as

$$\begin{aligned}
&\sum_{\theta=0,1} \mu(\theta) \sum_{S:i \in S} \phi(S|\theta)(\theta F(|S|) - r(i)) \\
=&\mu(1) \sum_{S:i \in S} \phi(S|1)(F(|S|) - r(i)) \\
=&\mu(1) \sum_{k=1}^{N} \sum_{S:|S|=k, i \in S} \phi(S|1)(F(k) - r(i)) \\
=&\mu(1) \sum_{k=1}^{N} (F(k) - r(i)) \sum_{S:|S|=k, i \in S} \phi(S|1)
\end{aligned}$$

$$=\mu(1)\sum_{k=1}^{N} p_{ik}(F(k)-r(i)).$$

Therefore for agent i, constraint (4.3) is equivalent as

$$\sum_{k=1}^{N} p_{ik}(F(k)-r(i)) \geq 0.$$

On the other hand, for each agent i, the left hand side of (4.4) can be written as

$$\begin{aligned}
&\sum_{\theta=0,1}\mu(\theta)\sum_{S:i\notin S}\phi(S|\theta)(\theta F(|S|+1)-r(i))\\
=&-\mu(0)r(i)+\mu(1)\sum_{S:i\notin S}\phi(S|1)(F(|S|+1)-r(i))\\
=&-\mu(0)r(i)+\mu(1)\sum_{k=0}^{N-1}\sum_{S:|S|=k,i\notin S}\phi(S|1)(F(k+1)-r(i))\\
=&-\mu(0)r(i)+\mu(1)\sum_{k=0}^{N-1}\left(\sum_{S:|S|=k}\phi(S|1)(F(k+1)-r(i))-\sum_{S:|S|=k,i\in S}\phi(S|1)(F(k+1)-r(i))\right)\\
=&-\mu(0)r(i)+\mu(1)\sum_{k=0}^{N-1}(F(k+1)-r(i))\sum_{S:|S|=k}\phi(S|1)-\mu(1)\sum_{k=1}^{N-1}(F(k+1)-r(i))\sum_{S:|S|=k,i\in S}\phi(S|1)\\
=&-\mu(0)r(i)+\mu(1)\sum_{k=0}^{N-1}(F(k+1)-r(i))\mathbf{P}(|\mathcal{S}|=k|\theta=1)-\mu(1)\sum_{k=1}^{N-1}(F(k+1)-r(i))p_{ik}\\
=&-\mu(0)r(i)-\mu(1)\sum_{k=1}^{N-1}(F(k+1)-r(i))p_{ik}+\mu(1)\sum_{k=1}^{N-1}(F(k+1)-r(i))\mathbf{P}(|\mathcal{S}|=k|\theta=1)\\
&+\mu(1)(F(1)-r(i))\mathbf{P}(|\mathcal{S}|=0|\theta=1)\\
=&-\mu(0)r(i)+\mu(1)\sum_{k=1}^{N-1}(F(k+1)-r(i))(\mathbf{P}(|\mathcal{S}|=k|\theta=1)-p_{ik})\\
&+\mu(1)(F(1)-r(i))\left(1-\sum_{k=1}^{N}\mathbf{P}(|\mathcal{S}|=k|\theta=1)\right)\\
=&-\mu(0)r(i)+\mu(1)\sum_{k=1}^{N-1}(F(k+1)-r(i))\left(\sum_{i=1}^{N}\frac{p_{ik}}{k}-p_{ik}\right)
\end{aligned}$$

$$+ \mu(1)(F(1) - r(i))\left(1 - \sum_{k=1}^{N}\sum_{i=1}^{N}\frac{p_{ik}}{k}\right)$$

$$= -\mu(0)r(i) + \mu(1)\left(\left(1 - \sum_{k=1}^{N}\frac{1}{k}\sum_{j=1}^{N}p_{jk}\right)(F(1) - r(i)) + \sum_{k=1}^{N-1}\left(\frac{1}{k}\sum_{j=1}^{N}p_{jk} - p_{ik}\right)(F(k+1) - r(i))$$

Therefore for agent i, constraint (4.4) is equivalent as

$$-\mu(0)r(i) + \mu(1)\left(\left(1 - \sum_{k=1}^{N}\frac{1}{k}\sum_{j=1}^{N}p_{jk}\right)(F(1) - r(i)) + \sum_{k=1}^{N-1}\left(\frac{1}{k}\sum_{j=1}^{N}p_{jk} - p_{ik}\right)(F(k+1) - r(i))\right) \leq 0$$

□

C.1.3 Proof of Lemma 4.3.3

Proof. For any $k > 0$, [68] presents a sequential elimination subroutine that eliminates one agent from agents $1, \ldots, N$ at each step, and output the remaining k agents, ensuring the probability each agent i is included in the sample is $kp_{ik}/\sum_{i=1}^{N} p_{ik}$, so long as (4.9) holds for p_{ik}'s.

Specifically, at each step $n = N-1, N-2, \ldots, k$, this subroutine computes $\pi(i|n)$ for each $i \in [N]$ in the following way. It first computes for each i the quantity

$$\frac{np_{ik}}{\sum_{i=1}^{N} p_{ik}}.$$

For each agent i such that this quantity is larger than 1, it assigns this quantity to $\pi(i|n)$. It then recomputes this quantity for remaining agents, with the multiplier in the numerator being the number of remaining agents instead of n. This process is repeated until for all $i \in [N]$, $\pi(i|n)$ is in $[0, 1]$. After this process, some $\pi(i|n)$ is equal to 1 and others are strictly proportional to p_{ik}.

This subroutine then computes for each $i \in [N]$

$$r_{ni} = \begin{cases} 1 - \pi(i|n), & \text{if } n = N - 1; \\ 1 - \frac{\pi(i|n)}{\pi(i|n+1)}, & \text{if } n < N - 1. \end{cases}$$

For each n, let S_n be the set of remaining agents at the beginning of step n, it can be verified $\{r_{ni}\}_{i \in S_n}$ is a valid probability distribution, and this subroutine samples one agent according to this distribution and eliminate this agent from the pool. If (4.9) holds, this subroutine ensures after the final step k, the probability each agent i is included in the sample is $kp_{ik} / \sum_{i=1}^{N} p_{ik} = q_{ik}$.

The rest of the proof of this lemma is given in the main paper.

□

C.1.4 Proof of Proposition 3

Proof. First we show $F(i^*) \geq r(i^*)$. This claim holds trivially when $i^* = 0$. Assume $i^* \geq 1$ and assume for contradiction $F(i^*) < r(i^*)$. We have

$$\begin{aligned} &\tilde{W}(i^*) - \tilde{W}(i^* - 1) \\ =& i^* F(i^*) - \sum_{j=1}^{i^*} r(j) - \left((i^* - 1)^* F(i^* - 1) - \sum_{j=1}^{i^*-1} r(j) \right) \end{aligned}$$

$$
\begin{aligned}
&=(i^* - 1)(F(i^*) - F(i^* - 1)) + F(i^*) - r(i^*)\\
&\leq F(i^*) - r(i^*)\\
&<0.
\end{aligned}
$$

Therefore i^* is not maximizer of $\tilde{W}$, leading to a contradiction.

Let x be the probability distribution over the subsets of $[N]$ such that $x(\{1, \ldots, i^*\}) = 1$ and $x(S) = 0$ for any other $S \subseteq [N]$. To show recommending the agents to follow the social optimal strategy profile is persuasive, it suffices to show x satisfies the constraints of the linear program (LP.1).

For agent $i \leq i^*$, $x(S) = 0$ for all S such that $i \notin S$ so the second constraint in (LP.1) is satisfied. Also,

$$
\sum_{S:i\in S} x(S)(F(|S|) - r(i)) = x(\{1, \ldots, i^*\})(F(i^*) - r(i)) = F(i^*) - r(i) \geq F(i^*) - r(i^*) \geq 0,
$$

hence the first constraint is also satisfied.

On the other hand, for agent $i > i^*$, $x(S) = 0$ for all S such that $i \in S$ so the first constraint in (LP.1) is satisfied. Also,

$$
\begin{aligned}
&\sum_{S:i\notin S} x(S)(F(|S| + 1) - r(i))\\
=&F(i^* + 1) - r(i)\\
\leq&F(i^* + 1) - r(i^* + 1)\\
\leq&\frac{r(i^* + 1)}{\mu(1)} - r(i^* + 1)
\end{aligned}
$$

$$
\begin{aligned}
&= \frac{\mu(0)}{\mu(1)} r(i^* + 1) \\
&\leq \frac{\mu(0)}{\mu(1)} r(i),
\end{aligned}
$$

so the second constraint is satisfied. Therefore x is persuasive.

The social welfare corresponding to x is $\mu(1)\tilde{W}(i^*)$. Note that for any $S \subseteq [N]$,

$$
|S|F(|S|) - \sum_{i \in S} r(i) \leq |S|F(|S|) - \sum_{i=1}^{|S|} r(i) = \tilde{W}(|S|) \leq \tilde{W}(i^*).
$$

Therefore for any signaling scheme ϕ, the social welfare with respect to ϕ is

$$
\begin{aligned}
&\mu(0) \sum_{S \subseteq \{[N]\}} \phi(S|0) \left(- \sum_{i \in S} r(i) \right) + \mu(1) \sum_{S \subseteq [N]} \phi(S|1) \left(|S|F(|S|) - \sum_{i \in S} r(i) \right) \\
\leq &\mu(1) \sum_{S \subseteq [N]} \phi(S|1) \tilde{W}(i^*) \\
\leq &\mu(1)\tilde{W}(i^*),
\end{aligned}
$$

hence x gives the optimal signaling mechanism.

□

C.1.5 Proof of Proposition 4

Proof. Since agent i is randomizing between moving and staying in equilibrium, assuming all other agents follow the given strategy profile, she should be indifferent in moving and staying, hence the expected payoff for agent i choosing to move should be equal to $r(i)$. Let N_{-ij} be the number of agents who choose to move besides agent i and j, we have

$$q(p_j\mathbf{E}[F(N_{-ij}+2)] + (1-p_j)\mathbf{E}[F(N_{-ij}+1)]) = r(i)$$
$$\Rightarrow p_j = \frac{\mathbf{E}[F(N_{-ij}+1)] - r(i)/q}{\mathbf{E}[F(N_{-ij}+1) - F(N_{-ij}+2)]}.$$

Similarly, agent j is indifferent in moving and staying given other agents' strategies, which gives

$$p_i = \frac{\mathbf{E}[F(N_{-ij}+1)] - r(j)/q}{\mathbf{E}[F(N_{-ij}+1) - F(N_{-ij}+2)]}. \tag{C.1}$$

Since F is decreasing, $\mathbf{E}[F(N_{-ij}+1) - F(N_{-ij}+2)] > 0$ and $r(i) \le r(j)$ implies $p_i \le p_j$. □

C.1.6 Proof of Lemma 4.4.1

Proof. For any $t \in [\underline{i}(q), \bar{i}(q)]$, let $p = t + 1 - \lceil t \rceil$. We consider the utility for each agent i choosing to move assuming all other agents follow their strategy in the threshold equilibrium.

For $i < \lceil t \rceil$, agent i' s expected utility for moving is

$$q(pF(\lceil t \rceil) + (1-p)F(\lceil t \rceil - 1)) \ge qF(\lceil t \rceil) \ge F(\bar{i}(q)) \ge r(\bar{i}(q)) \ge r(i),$$

hence she has no incentive to alter her strategy in the threshold equilibrium.

For agent $\lceil t \rceil$, her expected utility for moving is $qF(\lceil t \rceil)$. If $t > \underline{i}(q)$, we have $\underline{i}(q) < \lceil t \rceil \le \bar{i}(q)$, and by definition of $\underline{i}(q)$ and $\bar{i}(q)$, we have $qF(\lceil t \rceil) = r(\lceil t \rceil)$ so agent $\lceil t \rceil$ would not alter her strategy. On the other hand, if $t = \underline{i}(q)$, by definition of

$\underline{i}(q)$ we know agent $\lceil t \rceil$ would also choose to move, same as her strategy in the threshold equilibrium.

Finally, for agent $i > \lceil t \rceil$, her expected utility for moving is $q(pF(\lceil t \rceil + 1) + (1 - p)F(\lceil t \rceil))$. In the case $t > \underline{i}(q)$, this utility is upper bounded by $qF(\lceil t \rceil) \leq qF(\underline{i}(q) + 1)$; and in the case $t = \underline{i}(q)$, we have $p = 1$ and the utility for moving is also $qF(\underline{i}(q) + 1)$. We have $qF(\underline{i}(q) + 1) \leq r(\underline{i}(q) + 1) \leq r(i)$ so agent i would not alter her equilibrium strategy, which completes the proof for (i).

The "only if" part is straightforward and we omit the details. $\square$

C.1.7 Proof of Theorem 4.4.1

Proof. We first prove the following lemma.

Lemma C.1.1. *$W(q, n)$ is concave in n and $W(q, \underline{i}(q)) \geq W(q, \underline{i}(q) + 1)$.*

Proof. Concavity is straightforward since $nF(n)$ is concave and $r(i)$'s are increasing in i so $-\sum_{i=1}^{n} r(i)$ is also concave. To show $W(q, \underline{i}(q)) \geq W(q, \underline{i}(q) + 1)$, we have

$$\begin{aligned}
&W(q, \underline{i}(q) + 1) - W(q, \underline{i}(q)) \\
=&q(\underline{i}(q) + 1)F(\underline{i}(q) + 1) - \sum_{i=1}^{\underline{i}(q)+1} r(i) - \left(q\underline{i}(q)F(\underline{i}(q)) - \sum_{i=1}^{\underline{i}(q)} r(i) \right) \\
=&q\underline{i}(q)(F(\underline{i}(q) + 1) - F(\underline{i}(q))) + qF(\underline{i}(q) + 1) - r(\underline{i}(q) + 1) \\
\leq&0,
\end{aligned}$$

where the inequality is from $qF(\underline{i}(q)+1) \leq r(\underline{i}(q)+1)$ and F is decreasing. $\square$

Let $p = (p_1, \ldots, p_N)$ be an arbitrary equilibrium strategy profile under belief q. Define $\mathcal{I}_1$, $\mathcal{I}_0$ and $\mathcal{I}_{\text{mix}}$ as: $\mathcal{I}_1 \triangleq \{1 \leq i \leq N : p_i = 1\}$, $\mathcal{I}_0 \triangleq \{1 \leq i \leq N : p_i = 0\}$ and $\mathcal{I}_{\text{mix}} \triangleq \{1 \leq i \leq N : 0 < p_i < 1\}$, denoting the agents who moves, stays and randomizes between moving and staying. Recall $W(q,p)$ denote the principal's expected utility with respect to belief q and this strategy profile. We aim to show $W(q,p) \leq W(q, \underline{i}(q))$.

We have

$$\begin{aligned}
W(q,p) &= \sum_{S \subseteq \mathcal{I}_{\text{mix}}} \prod_{j \in S} p_j \prod_{j \in \mathcal{I}_{\text{mix}} \setminus S} (1-p_j) \left(q(|S| + |\mathcal{I}_1|) F(|S| + |\mathcal{I}_1|) - \sum_{j \in \mathcal{I}_1 \cup S} r(j) \right) \\
&\leq \sum_{S \subseteq \mathcal{I}_{\text{mix}}} \prod_{j \in S} p_j \prod_{j \in \mathcal{I}_{\text{mix}} \setminus S} (1-p_j) \left(q(|S| + |\mathcal{I}_1|) F(|S| + |\mathcal{I}_1|) - \sum_{j=1}^{|\mathcal{I}_1|+|S|} r(j) \right) \qquad \text{(C.2)} \\
&= \sum_{n=0}^{|\mathcal{I}_{\text{mix}}|} \left(q(|\mathcal{I}_1| + n) F(|\mathcal{I}_1| + n) - \sum_{j=1}^{|\mathcal{I}_1|+n} r(j) \right) \sum_{S \subseteq \mathcal{I}_{\text{mix}} : |S| = n} \prod_{j \in S} p_j \prod_{j \in \mathcal{I}_{\text{mix}} \setminus S} (1-p_j) \\
&= \sum_{n=0}^{|\mathcal{I}_{\text{mix}}|} W(q, |\mathcal{I}_1| + n) \sum_{S \subseteq \mathcal{I}_{\text{mix}} : |S| = n} \prod_{j \in S} p_j \prod_{j \in \mathcal{I}_{\text{mix}} \setminus S} (1-p_j).
\end{aligned}$$

The inequality is from that $r(i)$'s are increasing in i. Let X be the number of agents in $\mathcal{I}_{\text{mix}}$ that chooses to move. The right-hand side of the last equality of (C.2) equals $\mathbf{E}[W(q, |\mathcal{I}_1| + X)]$. Since W is concave, by Jensen's inequality we have

$$\mathbf{E}[W(q, |\mathcal{I}_1| + X)] \leq W(q, |\mathcal{I}_1| + \mathbf{E}[X]). \qquad \text{(C.3)}$$

Therefore, $U(p_1, \ldots, p_N) \leq W(q, |\mathcal{I}_1| + \mathbf{E}[X])$. Next, we show $|\mathcal{I}_1| + \mathbf{E}[X] \geq \underline{i}(q)$ in Lemma C.1.2.

Lemma C.1.2. $|\mathcal{I}_1| + \mathbf{E}[X] \geq \underline{i}(q)$.

Proof. Define $G(n) \triangleq qF(n), n \in [N]$. Since G is convex, by Jensen's inequality, we have

$$G(|\mathcal{I}_1| + \mathbf{E}[X] + 1) \leq \mathbf{E}[G(|\mathcal{I}_1| + X + 1)]. \tag{C.4}$$

For any agent $i \in \mathcal{I}_0$, her expected utility if she chooses to move is $\mathbf{E}[G(|\mathcal{I}_1| + X + 1)]$. Since she prefers staying, we have $\mathbf{E}[G(|\mathcal{I}_1| + X + 1)] \leq r(i)$. On the other hand, for any agent $i \in \mathcal{I}_{\text{mix}}$, let $X^{(i)}$ denote the number of agents in $\mathcal{I}_{\text{mix}}$ besides agent i who chooses to move. Agent i's expected utility if she chooses to move is $\mathbf{E}[G(|\mathcal{I}_1| + X^{(i)} + 1)]$. Since agent i is indifferent in moving and staying, we have $\mathbf{E}[G(|\mathcal{I}_1| + X^{(i)} + 1)] = r(i)$. Clearly $X^{(i)} \preccurlyeq_{\mathsf{sd}} X$ where "sd" denotes first-order stochastic dominance, and since G is a decreasing function, we have $G(|\mathcal{I}_1| + X + 1) \preccurlyeq_{\mathsf{sd}} G(|\mathcal{I}_1| + X^{(i)} + 1)$. Therefore $\mathbf{E}[G(|\mathcal{I}_1| + X + 1)] \leq \mathbf{E}[G(|\mathcal{I}_1| + X^{(i)} + 1)] = r(i)$. Summarizing the above, we have

$$\mathbf{E}[G(|\mathcal{I}_1| + X + 1)] \leq \min_{i \in \mathcal{I}_0 \cup \mathcal{I}_{\text{mix}}} r(i). \tag{C.5}$$

Let $j \triangleq \min \mathcal{I}_0 \cup \mathcal{I}_{\text{mix}}$ and we show it must hold that $j \leq \bar{i}(q)$. Assume, for contradiction, $j > \bar{i}(q)$. First observe that since j is the agent in $\mathcal{I}_0 \cup \mathcal{I}_{\text{mix}}$ with the smallest index, it must hold $j \leq |\mathcal{I}_1| + 1$, which implies $|\mathcal{I}_1| \geq j - 1 \geq \bar{i}(q)$. Meanwhile, if $j \in \mathcal{I}_0$, we have $r(j) = \mathbf{E}[G(|\mathcal{I}_1| + X^{(j)} + 1)] \leq G(|\mathcal{I}_1| + 1)$; and if $j \in \mathcal{I}_{\text{mix}}$, we have $r(j) \leq \mathbf{E}[G(|\mathcal{I}_1| + X + 1)] \leq G(|\mathcal{I}_1| + 1)$. Therefore,

$$r(\bar{i}(q) + 1) \leq r(j) \leq G(|\mathcal{I}_1| + 1) \leq G(\bar{i}(q) + 1).$$

However by definition of $\bar{i}(q)$ it must be that $G(\bar{i}(q)+1) < r(\bar{i}(q)+1)$, which leads to a contradiction.

From $j \leq \bar{i}(q)$, along with (C.4) and (C.5), we have

$$G(|\mathcal{I}_1| + \mathbf{E}[X] + 1) \leq \mathbf{E}[G(|\mathcal{I}_1| + X + 1] \leq r(j) \leq r(\bar{i}(q)).$$

By definition of $\bar{i}(q)$ we have $G(\bar{i}(q)) \geq r(\bar{i}(q))$. In the case $G(\bar{i}(q)) > r(\bar{i}(q))$, we have $G(\bar{i}(q)) > G(|\mathcal{I}_1| + \mathbf{E}[X] + 1)$ hence $|\mathcal{I}_1| + \mathbf{E}[X] + 1 > \bar{i}(q)$ since G is decreasing. Note when $G(\bar{i}(q)) > r(\bar{i}(q))$, $\underline{i}(q) = \bar{i}(q)$, so we have $|\mathcal{I}_1| + \mathbf{E}[X] \geq \underline{i}(q)$. On the other hand, if $G(\bar{i}(q)) = r(\bar{i}(q))$, we have $\underline{i}(q) < \bar{i}(q)$ so $|\mathcal{I}_1| + \mathbf{E}[X] + 1 \geq \bar{i}(q) > \underline{i}(q)$ and we also get $|\mathcal{I}_1| + \mathbf{E}[X] \geq \underline{i}(q)$. □

By Lemma C.1.1, $W(i)$ is decreasing on $i \geq \underline{i}(q)$. Since $|\mathcal{I}_1| + \mathbf{E}[X] \geq \underline{i}(q)$ we have $W(q, |\mathcal{I}_1| + \mathbf{E}[X]) \leq W(q, \underline{i}(q))$. Along with (C.2), (C.3) and Lemma C.1.2 we have

$$W(q,p) \leq \mathbf{E}[W(q, |\mathcal{I}_1| + X)] \leq W(q, \mathcal{I}_1| + \mathbf{E}[X]) \leq W(q, \underline{i}(q)),$$

completing the proof.

□

C.2 Upper Bound of $\mu(1)$

In this section, we compute the upper bound $r(i^* + 1)/F(i^* + 1)$ for $\mu(1)$ given in Proposition 3, which ensures recommending the agents following the social optimal strategy profile is persuasive, for several typical resource sharing functions and cost structures.

Specifically, we consider cost function $F(i) = 1/i^{\alpha}$ for $\alpha = 0.2, 0.4, 0.6, 0.8$, with different α controlling the curvature of F and representing different resource sharing scenarios. We consider 3 different cost structures: constant costs , linear costs, and quadratic costs. We fix the total number of agents $N = 20$ since i^* does not depend on N. The result is given in Table C.1.

C.3 Additional Computational Results

In this section, we present computational results on how much social welfare can be generated by the optimal private and public signaling mechanism under different priors. We consider three different resource sharing function $F(i) = 1/i^{\alpha}$ for $\alpha = 0.2, 0.5, 0.9$; and three different cost structures: constant costs where $r(i) = 0.25$, linear costs where $r(i) = 0.05i$ and quadratic costs where $r(i) = 0.01i^2$, and r is the cost coefficient. The results are given in Figure C.1.

α \ r	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.2	1	1	1	1	1	1	1	1	0.811	0.808
0.4	1	1	1	0.621	0.628	0.653	0.666	0.696	0.698	0.776
0.6	1	0.422	0.44	0.459	0.483	0.58	0.531	0.606	0.682	0.758
0.8	0.237	0.241	0.261	0.348	0.435	0.522	0.609	0.696	0.783	0.871

(a) Resource sharing function $F(i) = 1/i^{\alpha}$, with constant costs $r(i) = 0.5r$.

α \ r	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.2	1	1	0.836	0.869	0.888	0.838	0.849	0.826	0.93	0.859
0.4	0.617	0.648	0.65	0.735	0.762	0.737	0.666	0.761	0.857	0.696
0.6	0.464	0.45	0.527	0.525	0.459	0.551	0.643	0.464	0.521	0.58
0.8	0.252	0.243	0.364	0.289	0.361	0.433	0.506	0.279	0.313	0.348

(b) Resource sharing function $F(i) = 1/i^{\alpha}$, with linear costs $r(i) = 0.1ri$.

α \ r	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.2	0.891	0.947	0.951	1	0.97	0.868	1	0.824	0.927	1
0.4	0.631	0.78	0.64	0.854	0.737	0.885	0.666	0.761	0.857	0.952
0.6	0.446	0.63	0.633	0.525	0.657	0.441	0.515	0.588	0.662	0.735
0.8	0.302	0.362	0.291	0.388	0.485	0.26	0.303	0.347	0.39	0.433

(c) Resource sharing function $F(i) = 1/i^{\alpha}$, with quadratic costs $r(i) = 0.02ri^2$.

Table C.1: The upper bound $r(i^* + 1)/F(i^* + 1)$ of $\mu(1)$ for the private signaling mechanism that recommends every agent to stay when $\theta = 0$ and recommends the first i^* agents to move when $\theta = 1$ to be optimal, given by Proposition 3, for different resource sharing functions and cost structures.

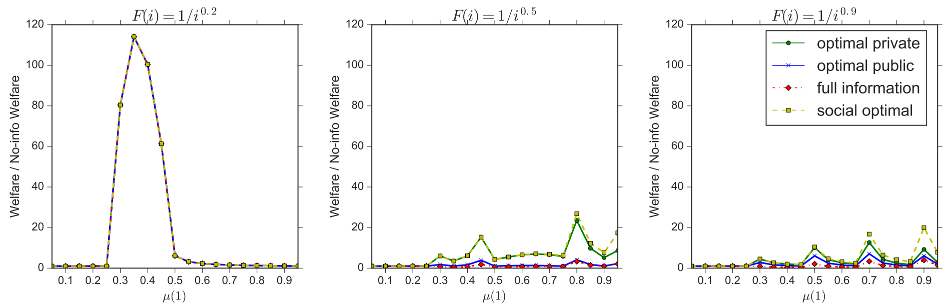

(a) Constant costs: $r(i) = 0.25$.

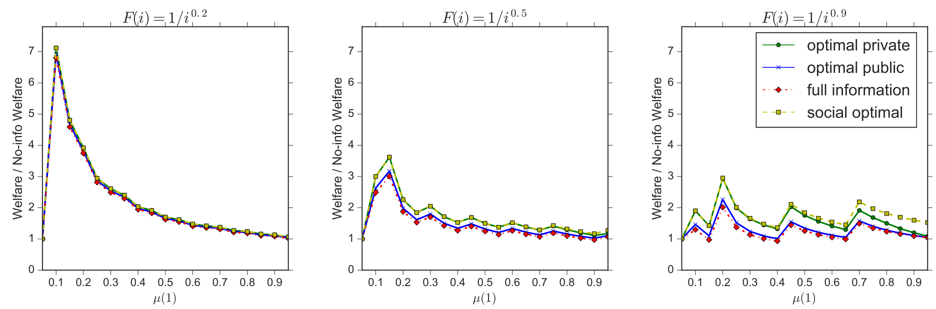

(b) Linear costs: $r(i) = 0.05i$.

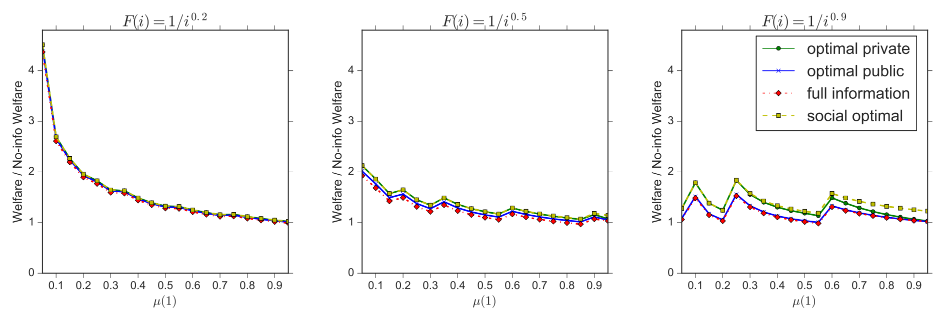

(c) Quadratic costs: $r(i) = 0.01i^2$.

Figure C.1: Social welfare of the optimal signaling mechanisms and the benchmarks, under different prior beliefs, for different resource sharing and cost functions. In all experiments, $N = 20$.

BIBLIOGRAPHY

[1] Sachin Adlakha, Ramesh Johari, and Gabriel Y. Weintraub. Equilibria of dynamic games with many players: Existence, approximation, and market structure. *Journal of Economic Theory*, 156:269–316, March 2015.

[2] Charalambos D. Aliprantis and Kim Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer Science & Business Media, 2006.

[3] Ricardo Alonso and Odilon Câmara. Persuading voters. *American Economic Review*, 106(11):3590–3605, 2016.

[4] Itai Arieli and Yakov Babichenko. Private bayesian persuasion. *Journal of Economic Theory*, 182:185–217, 2019.

[5] Nick Arnosti, Ramesh Johari, and Yash Kanoria. Managing congestion in decentralized matching markets. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 451–451, New York, NY, USA, 2014. ACM.

[6] W. Brian Arthur. Inductive reasoning and bounded rationality. *The American economic review*, pages 406–411, 1994.

[7] Yakov Babichenko and Siddharth Barman. Computational aspects of private bayesian persuasion. *arXiv preprint arXiv:1603.01444*, 2016.

[8] Yakov Babichenko and Siddharth Barman. Algorithmic aspects of private bayesian persuasion. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[9] Ashwinkumar Badanidiyuru, Kshipra Bhawalkar, and Haifeng Xu. Targeting

and signaling in ad auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2545–2563. SIAM, 2018.

[10] Santiago R. Balseiro, Omar Besbes, and Gabriel Y. Weintraub. Repeated auctions with budgets in ad exchanges: Approximations and design. *Management Science*, 61(4):864–884, 2015.

[11] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. Pricing and optimization in shared vehicle systems: An approximation framework. *CoRR*, abs/1608.06819, 2016.

[12] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. Pricing in ride-sharing platforms: A queueing-theoretic approach. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 639–639, New York, NY, USA, 2015. ACM.

[13] Arjada Bardhi and Yingni Guo. Modes of persuasion toward unanimous consent. *Theoretical Economics*, 13(3):1111–1149, 2018.

[14] Claude Berge. *Topological Spaces: Including a treatment of multi-valued functions, vector spaces, and convexity, translated by E. M. Patterson*. Dover, 1963.

[15] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–57, 2015.

[16] Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

[17] Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013.

[18] Kostas Bimpikis, Ozan Candogan, and Saban Daniela. Spatial pricing in ride-sharing networks. *Available at SSRN: https://ssrn.com/abstract=2868080*, 2016.

[19] Anton Braverman, J.G. Dai, Xin Liu, and Lei Ying. Empty-car routing in ridesharing systems. *arXiv preprint arXiv:1609.07219*, 2016.

[20] Bernard Caillaud and Jean Tirole. Consensus building: How to persuade a group. *American Economic Review*, 97(5):1877–1900, 2007.

[21] Ozan Candogan. Persuasion in networks: Public signals and k-cores. *Available at SSRN*, 2019.

[22] Ozan Candogan and Kimon Drakopoulos. Optimal signaling of content accuracy: Engagement vs. misinformation. *Misinformation (October 11, 2017)*, 2017.

[23] Juan Camilo Castillo, Dan Knoepfle, and Glen Weyl. Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC '17, pages 241–242, New York, NY, USA, 2017. ACM.

[24] Anindya-Sundar Chakrabarti, Bikas K. Chakrabarti, Arnab Chatterjee, and Manipushpak Mitra. The kolkata paise restaurant problem and resource utilization. *Physica A: Statistical Mechanics and its Applications*, 388(12):2420–2426, 2009.

[25] Bikas K. Chakrabarti. Kolkata restaurant problem as a generalised el farol bar problem. In *Econophysics of Markets and Business Networks*, pages 239–246. Springer, 2007.

[26] Archishman Chakraborty and Rick Harbaugh. Persuasive puffery. *Marketing Science*, 33(3):382–400, 2014.

[27] M. Keith Chen. Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, EC '16, pages 455–455, New York, NY, USA, 2016. ACM.

[28] Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982.

[29] New York City Taxi & Limousine Comission Trip Record Data. New york city taxi & limousine comission trip record data., 2017. `http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml`.

[30] Kimon Drakopoulos, Shobhit Jain, and Ramandeep S Randhawa. Persuading customers to buy early: The value of personalized information provisioning. 2018.

[31] Shaddin Dughmi. Algorithmic information structure design: a survey. 15(2):2–24, 2017.

[32] Shaddin Dughmi and Haifeng Xu. Algorithmic persuasion with no externalities. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 351–368. ACM, 2017.

[33] R. Durrett and S. Levin. The importance of being discrete (and spatial). *Theoretical Population Biology*, 46(3):363 – 394, 1994.

[34] Richard Durrett and Simon A. Levin. Stochastic spatial models: A user's guide to ecological applications. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 343(1305):329–350, 1994.

[35] Stewart N. Ethier and Thomas G. Kurtz. *Markov processes: characterization and convergence*. John Wiley & Sons, 1986.

[36] Joseph Farrell. Cheap talk, coordination, and entry. *The RAND Journal of Economics*, pages 34–39, 1987.

[37] Drew Fudenberg and Jean Tirole. *Game theory*. The MIT Press, Cambridge, Massachusetts, 1991.

[38] Drew Fudenberg and Jean Tirole. Perfect bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260, 1991.

[39] Matthew Gentzkow and Emir Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104:411–429, 2017.

[40] Asim Ghosh, Arnab Chatterjee, Manipushpak Mitra, and Bikas K. Chakrabarti. Statistics of the kolkata paise restaurant problem. *New Journal of Physics*, 12(7):075033, 2010.

[41] Sanford J Grossman. The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3):461–483, 1981.

[42] Weather History and Data Archive. Weather history and data archive. `https://www.wunderground.com/history/`, 2017.

[43] H. A. Hopenhayn. Entry, exit and firm dynamics in long run equilibrium. *Econometrica*, 60(5):1127 – 1150, 1992.

[44] Chamsi Hssaine and Siddhartha Banerjee. Information signal design for incentivizing team formation. *arXiv preprint arXiv:1809.00751*, 2018.

[45] M. Huang, P. E. Caines, and R. P. Malhamé. Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ϵ-Nash equilibria. *IEEE Transactions on Automatic Control*, 52(9):1560–1571, 2007.

[46] Krishnamurthy Iyer, Ramesh Johari, and Mukund Sundararajan. Mean field equilibria of dynamic auctions with learning. *Management Science*, 60(12):2949–2970, 2014.

[47] B. Jovanovic and R. W. Rosenthal. Anonymous sequential games. *Journal of Mathematical Economics*, 17:77–87, 1988.

[48] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[49] Juan E. Keymer, Pablo A. Marquet, Jorge X. VelascoHernndez, and Simon A. Levin. Extinction thresholds and metapopulation persistence in dynamic landscapes. *The American Naturalist*, 156(5):478–494, 2000.

[50] Ilan Kremer, Yishay Mansour, and Motty Perry. Implementing the wisdom of the crowd. *Journal of Political Economy*, 122(5):988–1012, 2014.

[51] Aimé Lachapelle and Marie-Therese Wolfram. On a mean field game approach modeling congestion and aversion in pedestrian crowds. *Transportation Research Part B: Methodological*, 45(10):1572 – 1589, 2011.

[52] J. M. Lasry and P. L. Lions. Mean field games. *Japanese Journal of Mathematics*, 2:229–260, 2007.

[53] Cuong Le Van and John Stachurski. Parametric continuity of stationary distributions. *Economic Theory*, 33(2):333–348, 2007.

[54] R Levin. Extinction. *Some mathematical problems in biology. American Mathematical Society, Providence, Rhode Island*, pages 77–107, 1970.

[55] Jian Li, Rajarshi Bhattacharyya, Suman Paul, Srinivas Shakkottai, and Vijay Subramanian. Incentivizing sharing in realtime d2d streaming networks: A mean field game perspective. *IEEE/ACM Trans. Netw.*, 25(1):3–17, February 2017.

[56] David Lingenbrink and Krishnamurthy Iyer. Optimal signaling mechanisms in unobservable queues with strategic customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 347–347. ACM, 2017.

[57] David Lingenbrink and Krishnamurthy Iyer. Signaling in online retail: Efficacy of public signals. In *Proceedings of the 13th Workshop on Economics of Networks, Systems and Computation*, NetEcon '18, pages 10:1–10:1, New York, NY, USA, 2018. ACM.

[58] M. Manjrekar, V. Ramaswamy, and S. Shakkottai. A mean field game approach to scheduling in cellular systems. In *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*, pages 1554–1562, April 2014.

[59] Paul R Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.

[60] Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge university press, 2005.

[61] Jane Molofsky. Population dynamics and pattern formation in theoretical populations. *Ecology*, 75(1):30–39, 1994.

[62] John A. Nelder and Roger Mead. A simplex method for function minimization. *The computer journal*, 7(4):308–313, 1965.

[63] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*. Cambridge University Press Cambridge, 2007.

[64] Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of political Economy*, 118(5):949–987, 2010.

[65] Robert W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

[66] James E. Smith and Kevin F. McCardle. Structural properties of stochastic dynamic programs. *Operations Research*, 50(5):796–809, 2002.

[67] Ina A Taneva. Information design. 2015.

[68] Yves Tillé. An elimination procedure for unequal probability sampling without replacement. *Biometrika*, 83(1):238–241, 1996.

[69] Uber. How surge pricing works, 2019.

[70] Yun Wang. Bayesian persuasion with multiple receivers. *Available at SSRN 2625399*, 2013.

[71] G. Y. Weintraub, C. L. Benkard, and B. VanRoy. Markov perfect industry dynamics with many firms. *Econometrica*, 76(6):1375–1411, 2008.

[72] Gabriel Y. Weintraub, C. Lanier Benkard, and Benjamin van Roy. Industry dynamics: Foundations for models with an infinite number of firms. *Journal of Economic Theory*, 146(5):1965 – 1994, 2011.

[73] Aihua Xia. Weak convergence of markov processes with extended generators. *The Annals of Probability*, 22(4):2183–2202, 10 1994.

[74] Haifeng Xu, Rupert Freeman, Vincent Conitzer, Shaddin Dughmi, and Milind Tambe. Signaling in bayesian stackelberg games. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 150–158. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

[75] Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi, and Milind Tambe. Exploring information asymmetry in two-stage security games. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

[76] Jiaming Xu and Bruce Hajek. The supermarket game. *Stochastic Systems*, 3(2):405–441, 2013.

[77] Pu Yang, Krishnamurthy Iyer, and Peter I Frazier. Mean field equilibria for competitive exploration in resource sharing settings. In *Proceedings of the 25th International Conference on World Wide Web*, pages 177–187. International World Wide Web Conferences Steering Committee, 2016.

[78] Pu Yang, Krishnamurthy Iyer, and Peter I Frazier. Mean field equilibria for resource competition in spatial settings. *Stochastic Systems*, 8(4):307–334, 2018.