Models of Competition for Intelligent Transportation Infrastructure: Parking, Ridesharing, and External Factors in Routing Decisions

Daniel Calderone

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2017-101

http://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-101.html

May 12, 2017

Models of Competition for Intelligent Transportation Infrastructure: Parking, Ridesharing, and External Factors in Routing Decisions

by

Daniel Joseph Calderone

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

S. Shankar Sastry, Chair
Alexandre M. Bayen
Roberto Horowitz

Spring 2017

Models of Competition for Intelligent Transportation Infrastructure: Parking, Ridesharing, and External Factors in Routing Decisions

Abstract

Models of Competition for Intelligent Transportation Infrastructure:
Parking, Ridesharing, and External Factors in Routing Decisions

by

Daniel Joseph Calderone

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley

S. Shankar Sastry, Chair

Competition underlies much of the complexity of modern transportation systems and accurately modeling the incentives transportation users face is critical in designing the smart cities of tomorrow. In this work, we extend traditional non-atomic routing games to model several different scenarios relevant to modern cities. In Chapter 2, we review results on continuous population potential games and routing games and discuss linear programming interpretations of individual drivers' decisions and their relationship to the Wardrop equilibria. We include discussion of the variable demand case and upper bounds on the price of anarchy. In Chapter 3, we combine an observable queueing game with a routing game to develop a queue-routing game useful for modeling traffic circling looking for parking in urban centers. We use this framework to model several parking situations in downtown Seattle. In Chapter 4, we discuss connections between linear programming formulations of shortest path problems and linear programming formulations of Markov decision processes (MDPs). We use these connections to motivate a stochastic population game we call a *Markov decision process routing game* where each infinitesimal agent solves an MDP as opposed to a shortest path problem. We develop this game in the infinite-horizon, average-cost case and in the finite-horizon, total-cost case. We comment on connections with traditional routing games as well as other stochastic population game formulations such as mean field games. Paralleling results for traditional routing games, we derive upper bounds on the price of anarchy and comment on the existence of Braess paradox. We apply this framework to model ridesharing drivers competing for fares and to develop a model of circling traffic competing for street parking. Finally, in Chapter 5, we consider a bi-criterion routing game where drivers consider travel time along with some external factor. Preference for this external factor is represented by general distribution over a type parameter that can be supported both above and below zero. We develop the appropriate equilibrium concept and show how this framework can be used to model the transportation data market, drivers' interest in their location privacy, and commuters comparing two different commuting options. In Chapter 6, we conclude and make final comments on modeling considerations and future research directions.

To Roy, Katie, Dorsa, Aaron, Bari, and Juli

Contents

List of Figures

List of Tables

Chapter 1

Introduction: Competition in Smart Cities

1.1 Motivation

Transportation systems are the backbone of cities, supporting many crucial interactions including economic transactions, resource distribution, and emergency response. Due to the *great urban sprawl* [1–3], transportation infrastructure in cities is being stretched to its limits. As a result, cities incur large economic costs from transport-related inefficiencies [3,4] particularly congestion. Across the U.S., traffic congestion is responsible for nearly 4 billion gallons of wasted fuel a year and nearly 7 billion extra hours of travel time [5]. Beyond economic costs, congestion has adverse effects on public health, the environment, and general quality of life in cities [6, 7].

At the same time that congestion is becoming a bigger problem than ever, the big data and autonomy revolutions are transforming transportation infrastructure in unprecedented ways. Apps like Google maps and Waze provide real time dynamic congestion information to drivers as well as information about public transportation. Parking apps, such as ParkWhiz, ParkPanda, and SpotHero are also appearing that help drivers find and, in some cases, reserve parking spots in cities [8]. Ride-sharing services have also become incredibly popular with Uber and Lyft growing to billion dollar companies in less than 7 years [9]. Autonomous car technology also has the potential to revolutionize urban transportation as many large companies, both traditional car manufacturers and others, including Ford, Tesla, Google, and Uber invest in self-driving technology [10–12].

These new technologies have the potential to help alleviate congestion but they also change the incentives that transportation users face. Access to more real time data allows users to make more sophisticated decisions to optimize their travel experience. Autonomous systems have the ability to combine this new data with much greater computation power to even further optimize their performance. In order to design transportation for these smart cities in the future, urban planners need models of transportation systems that account for these more complicated game theoretic scenarios.

One of the most widely used models of competitive behavior in transportation behavior is nonatomic routing games, a specific type of continuous population potential game that models the aggregate behavior of infinitesimal agents that each attempt to solve a shortest path problem where the path travel time is affected by congestion. This model was first introduced by John Wardrop in 1952 [13] and has been studied extensively since. In this work, we extend this classic framework to model the optimal behavior of intelligent agents in several new transportation scenarios including urban parking, ride-sharing, and agents making decisions when they make transportation decisions based on external factors other than travel time.

1.2 Overview

In Chapter 2, we review classic results on routing games. In Chapter 3, we extend a classic routing game to model on-street parking in urban areas and how circling traffic affects congestion. In Chapter 4, we present a new continuous population game that parallels traditional routing games except rather than seeking a shortest path each agents seeks to find an optimal policy for a Markov decision process (MDP). This framework has close connections with other stochastic population game formulations such as *anonymous sequential games* and *mean field games*. We present our model and use it to model the behavior of ride-sharing drivers and traffic circling city blocks looking for parking. In Chapter 5, we examine how varying preferences for different transportation options among a population of users can change their travel decisions. We present a bi-criterion routing framework that allows us to examine how members of a population evaluate these trade-offs. We show how this model could be applied to study how drivers' concern over their location privacy can change their travel choices as well as to study how users make tradeoffs between different modes of travel (such as driving or taking the metro). Throughout, we compare the game theoretic or *equilibrium* behavior of agents in these different scenarios with the *socially optimal* behavior agents would implement if they were collaborating to minimize congestion. Specifically in Chapter 4, we present upper bounds on the well studied *price of anarchy*. Finally, in Chapter 6, we make final comments on modeling considerations and connections between the various chapters and present directions for future work. We present relevant literature review and connections with our work in each section as applicable. All simulations were done in Matlab using the optimization package YALMIP by Löfberg [14].

Chapter 2

Fundamentals of Routing Games

The *traffic assignment problem* (TAP), understanding how traffic patterns form in transportation networks, has been a major focus of research in both the transportation, economics, and game theory communities for over 75 years. One of the most popular approaches to modeling this problem is to treat traffic as a continuous fluid that spreads out across a road network according to the equilibrium principle that any route some portion of the population chooses has minimal travel time. These equilibrium principles are generally attributed to John Wardrop and bear his name.

Definition 1 (Wardrop Equilibrium Principles [13])

1. *The journey times on all the routes actually used are equal, and less than those which would be experienced by a single vehicle on any unused route.*

2. *The average journey time is a minimum.*

In this chapter, we will provide a brief overview of these *nonatomic routing games* specifically focusing on results that will be important for understanding our work. For a more thorough discussion of routing games, we refer the reader to Patrkisson's excellent monograph [15]. Rather than proceeding chronologically through the literature, we will discuss nonatomic routing games as a special case of *continuous population potential games* introduced by Sandholm in [16].

Remark 1 *For notational simplicity, we will present all the results in terms of one population of agents. However, all the arguments go through for multiple populations under certain assumptions. We will comment on these at select points.*

2.1 Continuous Population Potential Games

We first review several results from continuous population potential games from Sandholm's presentation [16]. We start by assuming a population of identical agents with total mass m. Each member of the population has the option of choosing from a discrete set of strategies denoted by the set $\mathcal{R}$. Let $z \in \mathbb{R}^{|\mathcal{R}|}$ be the overall population mass vector where z_r is the portion of the population that chooses strategy $r \in \mathcal{R}$. We have that $\sum_r z_r = m$ and that $z_r \geq 0$ for all r. When an agent chooses strategy r they receive loss $\ell_r(z)$ which is a function of the overall population distribution. (Let $\ell(z) : \mathbb{R}^{|\mathcal{R}|} \to \mathbb{R}^{|\mathcal{R}|}$ represent the vector of all losses.) Each agent seeks to minimize their loss, and Sandholm defines the Nash equilibrium of the game as the population distribution where no infinitesimal member of the population has an incentive to deviate from their strategy. This definition of a Nash equilibrium is the same as Wardrop's equilibrium concept and we refer to it as a *Wardrop equilibrium.* We can state the equilibrium condition succinctly as follows.

Definition 2 (Wardrop Equilibrium [13]) *A population distribution $z \in \mathbb{R}^{|\mathcal{R}|}$ is called a* Wardrop equilibrium *if for any two strategies $r, r' \in \mathcal{R}$ such that $z_r > 0$, we have that*

$$\ell_r(z) \leq \ell_{r'}(z). \tag{2.1}$$

Intuitively, the loss associated with any strategy with non-zero mass must be less than or equal to any other possible strategy, otherwise that portion of mass would switch away from that strategy. It follows that any strategy chosen by any member of the population must have equal loss.

We now define continuous population potential games as detailed by Sandholm.

Definition 3 (Continuous Potential Game [16]) *We say that a continuous population game is a* potential game *if there exists a C^1 function $F : \mathbb{R}^{|\mathcal{R}|} \to \mathbb{R}$ such that*

$$\frac{\partial F}{\partial z_r} = \ell_r(z) \tag{2.2}$$

The derivative of the potential function with respect to a particular subpopulation captures the marginal cost of choosing that strategy. The usefulness of this fact is given in the next result.

Theorem 1 *A population distribution $z \in \mathbb{R}^{|\mathcal{R}|}$ is a Wardrop equilibrium if and only if it satisfies the KKT necessary conditions for minimizing the potential function F with respect to the constraints $\sum_r z_r = m$ and $z_r \geq 0$. [16]*

We reproduce the proof hear since it is illuminating.

Proof 1 *First, z is a Wardrop equilibrium if it satisfies the KKT conditions. The first order necessary conditions give that*

$$\ell_r(z) = \lambda + \mu_r, \qquad \mu_r \geq 0, \qquad z_r \mu_r = 0 \qquad \forall r \in \mathcal{R} \tag{2.3}$$

where $\lambda \in \mathbb{R}$ and $\mu \in \mathbb{R}_+^{|\mathcal{R}|}$ are the Lagrange multipliers for the mass conservation constraint and positivity constraints respectively. For any two strategies r and r' such that $z_r > 0$, it follows that

$$\ell_r(z) = \lambda \leq \lambda + \mu_{r'} = \ell_{r'}(z) \tag{2.4}$$

Conversely, given a Wardrop equilibrium distribution z, let $\lambda = \min_r \ell_r(z)$. By Condition (2.1) *for any r such that $z_r > 0$, $\ell_r(z) = \lambda$. Setting $\mu_r = \ell_r(z) - \lambda$ yields the result.*

As detailed in the proof, the Lagrange multipliers have specific interpretations that are illustrated in Figure 2.1b. λ is the loss incurred by a member of the population (equal to the loss of each strategy actually chosen by members of the population) and μ_r is the inefficiency of choosing strategy r.

Linear programming interpretation of individuals' decisions

We note here another interpretation that arises when we model each individual agent in the population as choosing from a mixed strategy rather than a pure strategy. Given the overall population distribution z each agent solves the following linear program to choose a mixed strategy ζ.

$$\min_{\zeta} \quad l(z)^T \zeta \tag{2.5a}$$

$$\text{s.t.} \quad \sum_r \zeta_r = 1, \quad \zeta \geq 0 \tag{2.5b}$$

The first order optimality conditions for (2.5) are

$$\ell_r(z) = \lambda + \mu_r, \qquad \mu_r \geq 0, \qquad \xi_r \mu_r = 0 \qquad \forall r \in \mathcal{R} \tag{2.6}$$

Notice that these are closely related to the first order conditions for minimizing the potential function with respect to z detailed in (2.3). Specifically the first condition is identical. This indicates a strong relationship between the equilibrium mass distribution z_{eq} and any ζ that minimizes (2.5). This is no surprise since z models the aggregate effect of each individual's choice ζ.

As illustrated in Figure 2.1b in general at equilibrium, the population mass will be distributed over several strategies, for example r_1 and r_2, with equal losses. Choosing either r_1 or r_2 yields the same loss and thus a mixed strategy that is a convex combination of them yields the same loss. Thus an optimal ζ could be any convex combination of pure strategies with minimal losses. The key is that at equilibrium, the support of an optimal ζ is included in the support of z_{eq}, that is $\zeta_r > 0$ only if $(z_{\text{eq}})_r > 0$. Otherwise, several individuals in the population solving (2.5) might choose strategies with mass outside the equilibrium distribution and the equilibrium would shift. This is roughly the definition of Nash equilibrium given by Sandholm and this line of reasoning is the motivation for defining the potential function.

2.2 Routing Games: Path and Edge Formulations

We now turn to non-atomic routing games as a special case of continuous population games. In the classic routing game, populations of drivers seek to travel from origin nodes to destination nodes in a road network in a way that minimizes their travel time. In the next two sections, we present two formulations of the routing game, the *path* or *route* formulation and the *edge* formulation. The route formulation is a straight forward application of continuous population potential games but is often computationally intractable since it involves enumerating each route through a road network. The edge formulation is slightly more complicated but more computationally tractable.

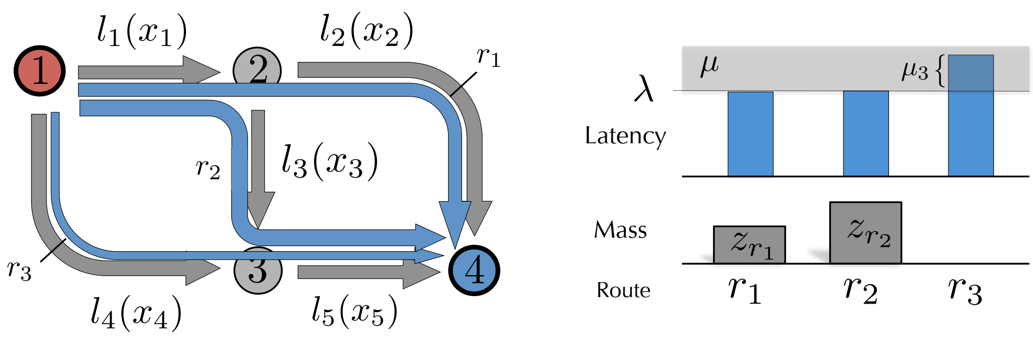

(a) (b)

Figure 2.1: (a) Sample graph with routes from origin to destination. (b) Illustration of the Wardrop equilibrium condition.

Path formulation

We start by defining a road network represented by a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with nodes and edges respectively. To simplify the presentation, we will assume one population of drivers with mass m that starts at a single origin node $o \in \mathcal{N}$ and travels to a single destination node $d \in \mathcal{N}$, but all the results hold for multiple populations with different origins and destinations. Drivers in the population choose from a set of routes through the network $\mathcal{R}$ and let $z \in \mathbb{R}^{|\mathcal{R}|}$ be the population distribution vector as before. We will refer to z as the *route flows* or *route masses*. When a portion of the population chooses a particular route, their mass gets added to each edge on that route. For each route $r \in \mathcal{R}$, we define the set of edges in that route $\mathcal{E}_r$ and an indicator vector $\mathbf{E}_r \in \{0,1\}^{|\mathcal{E}|}$ for the edges in route r.

$$[\mathbf{E}_r]_e = \begin{cases} 1 & ; \text{ if } e \in \mathcal{E}_r \\ 0 & ; \text{ otherwise} \end{cases} \tag{2.7}$$

We also define a *routing matrix* $\mathbf{E}_\mathcal{R} \in \{0,1\}^{|\mathcal{E}| \times |\mathcal{R}|}$ whose rth column is $\mathbf{E}_r$.

Let $x \in \mathbb{R}^{|\mathcal{E}|}$ be a vector of the total population of drivers using each edge or the *edge flows*. The route flows and edge flows are related by

$$x = \mathbf{E}_\mathcal{R} z \tag{2.8}$$

Each edge has an associated *latency function* $l_e(x_e)$ that indicates the time required to traverse the edge and is an increasing function of the flow on that edge. A typical latency function is illustrated in Figure 2.2.

Let $l(x) \in \mathbb{R}^{|\mathcal{E}|}$ represent the vector of edge latencies. The loss for a particular route or *route latency*, $\ell_r(z)$, is the total route travel time, i.e. the sum of the latencies of each edge

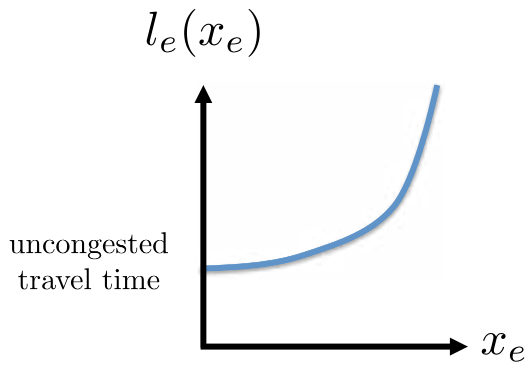

Figure 2.2: Typical latency function.

in that route. We can compute the vector of route latencies $\ell(z) \in \mathbb{R}^{|\mathcal{R}|}$ as

$$\ell(z) = \mathbf{E}_{\mathcal{R}}^T l(x) = \mathbf{E}_{\mathcal{R}}^T l(\mathbf{E}_{\mathcal{R}} z) \tag{2.9}$$

Our goal is to find the population distribution across the routes, z, that is a Wardrop equilibrium for the route latencies $\ell(z)$. We next show that this game is a continuous population potential game with potential function given by

$$F(z) = \sum_e \int_0^{x_e} l_e(u)\ du = \sum_e \int_0^{[\mathbf{E}_{\mathcal{R}} z]_e} l_e(u)\ du \tag{2.10}$$

which was first introduce in [17]. We will often abuse notation and write $F(x)$ as well as $F(z)$.

Proposition 1 *The non-atomic routing game is a continuous population potential game with potential function given by* (2.10)*.*

Proof 2 *The potential function is designed such that $\frac{\partial F}{\partial x} = l(x)^T$. Differentiating with respect to z by using Leibniz rule and then applying the chain rule gives*

$$\frac{\partial F}{\partial z} = l(x)^T \mathbf{E}_{\mathcal{R}} = \ell(z)^T \tag{2.11}$$

as desired.

The *path* or *route formulation* for finding the equilibrium of the routing game is given by

$$\min_z \quad F(z) \tag{2.12a}$$

$$\text{s.t.} \quad \sum_r z_r = m, \quad z \geq 0 \tag{2.12b}$$

and is a straightforward application of continuous population potential games.

Remark 2 *As mentioned previously, the main issue with this formulation is the need to compute the routing matrix* $\mathbf{E}_{\mathcal{R}}$ *which depending on the size of the graph is computationally intractable. For this reason, in practice we use the edge formulation detailed in the next section.*

Edge formulation

Rather than enumerating all possible paths from each origin to each destination, we can modify the optimization problem to solve for the edge flows directly. Instead of computing the routing matrix, we use an incidence matrix constraint. Let $G \in \{-1, 0, 1\}^{|\mathcal{N}| \times |\mathcal{E}|}$ be the node-edge incidence matrix of the directed graph.

$$[G]_{ne} = \begin{cases} 1 & ; \text{ if edge } e \text{ originates at node } n \\ -1 & ; \text{ if edge } e \text{ terminates at node } n \\ 0 & ; \text{ otherwise} \end{cases} \tag{2.13}$$

We will also find it useful to define "outgoing" and "incoming" incidence matrices, $I_o, I_i \in \{0, 1\}^{|\mathcal{N}| \times |\mathcal{E}|}$.

$$[I_o]_{ne} = \begin{cases} 1 & ; \text{ if edge } e \text{ originates at node } n \\ 0 & ; \text{ otherwise} \end{cases} \tag{2.14}$$

$$[I_i]_{ne} = \begin{cases} 1 & ; \text{ if edge } e \text{ terminates at node } n \\ 0 & ; \text{ otherwise} \end{cases} \tag{2.15}$$

Note that $G = I_o - I_i$.

We also define a source and sink vector, S, for the population of drivers traveling from node o to d.

$$S_n = \begin{cases} 1 & ; \ n = o \\ -1 & ; \ n = d \\ 0 & ; \text{ otherwise} \end{cases} \tag{2.16}$$

We can then minimize the potential function directly with respect to the edge flows

$$\min_x \quad F(x) \tag{2.17a}$$

$$\text{s.t.} \quad Gx = Sm, \quad x \geq 0 \tag{2.17b}$$

This is called *edge formulation* of the routing game.

Remark 3 *The set of feasible edge flows given the incidence matrix constraint includes all positive flows that route m mass from node o to node d as desired. However, it also includes flows in the null space of G, i.e. cyclic flows. Since $F(x)$ is strictly increasing in x_e, the minimizer of (2.17) will not include any cyclic component. For more details, we refer the reader to [15] (Theorem 2.2 and the following discussion) and the sources cited therein (Theorem 3.5 in [18] for a discussion of the feasible set and p. 154-155 of [19] for a discussion of the minimizer not including cyclic components).*

Remark 4 *For strictly increasing latency functions, the potential function in (2.10) is a strictly convex function (in the edge flows) and thus the KKT conditions ensure uniqueness of the equilibrium edge flows. Note this does not ensure uniqueness of the route flows since there may be multiple route flows that yield the same edges flows. This would be the case if for example there were more routes than edges (the routing matrix has a non-trivial nullspace.) Again, we refer the reader to [15] (Section 2.3) for further details.*

There is an alternative proof that if x minimizes (2.17) than it represents the edge flows associated with a Wardrop equilibrium of the routing game that is informative thus we present it here.

Proof 3 *The first order optimality conditions for (2.17) are*

$$l(x) = G^T\pi + \mu, \quad x \geq 0, \quad \mu \geq 0, \quad \mu^T x = 0 \tag{2.18}$$

with Lagrange multipliers $\pi \in \mathbb{R}^{|\mathcal{N}|}$ and $\mu \in \mathbb{R}_+^{|\mathcal{E}|}$ for the incidence matrix and positivity constraints respectively. This results in the following characterization of the latency of an edge that goes from node n_1 to node n_2.

$$l_e(x_e) = \pi_{n_2} - \pi_{n_1} + \mu_e, \qquad \textit{where } \mu_e \geq 0 \tag{2.19}$$

Summing the latencies along any route r from node o to node d, we get that

$$\ell_r = \sum_{e\in\mathcal{E}_r} l_e(x_e) = \pi_d - \pi_o + \sum_{e\in\mathcal{E}_r} \mu_e \tag{2.20}$$

Since $\mu_e \geq 0$ with equality achieved whenever $x_e > 0$ and since $x_e > 0$ for any edge in a route r such that $z_r > 0$, we have Condition (2.1).

As in the case of the route formulation, there is a specific interpretation of the Lagrange multipliers in the edge formulation. $\pi \in \mathbb{R}^{|\mathcal{N}|}$ can be thought of as a value function on the nodes that encodes the cost-to-go to the destination, i.e. for any given node n, $\pi_n - \pi_d$ is the minimum time to travel from node n to the destination node d. μ_e is the inefficiency of using edge e and thus $x_e = 0$ whenever $\mu_e > 0$.

Figure 2.3b illustrates the Wardrop equilibrium principle in the edge formulation framework for the sample graph shown in Figure 2.3a.

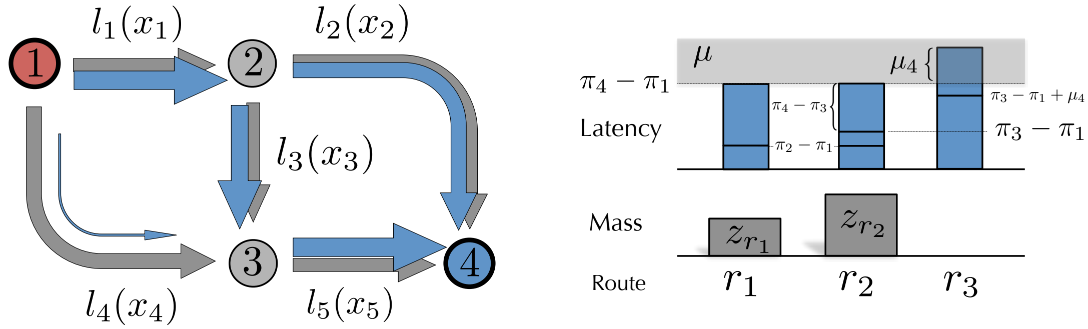

(a) (b)

Figure 2.3: (a) Sample graph. (b) Edge formulation of Wardrop equilibrium condition.

Remark 5 *As previously mentioned, multiple populations of drivers with different origin-destination pairs can be incorporated into the routing game. In the routing formulation case, we add a new population vector and a new routing matrix for each origin-destination pair. In the edge formulation case, we add a new population vector and incidence matrix constraint for each pair. The overall edge masses are then obtained by summing up the individual population masses. The gradient arguments for each population go through unchanged by applying the chain rule.*

Shortest path LP interpretation of individuals' decisions

As in the case of the simple potential game, there is a linear programming interpretation of each individual agent's decision related to the edge formulation of the routing game. Given the edge flows x and resulting edge travel times $l(x)$, each individual can solve for the shortest path from their origin to destination using the following linear program.

$$\begin{aligned} \min_{\xi} \quad & l(x)^T \xi \\ \text{s.t.} \quad & G\xi = S, \quad \xi \geq 0 \end{aligned} \tag{2.21}$$

The argument that this program solves for the shortest path given x is equivalent to the argument in Proof 3. Here $\xi \in [0,1]^{|\mathcal{E}|}$ and ξ_e is the probability of traveling on edge e. If there is only one shortest path, ξ will be an indicator vector for that path. If there are multiple shortest paths, ξ is a convex combination of those indicator vectors. Again, the important feature of the equilibrium edge flows x_{eq} is that the support of any minimizer ξ of (2.21) is contained in the support of x_{eq}. That is $\xi_e > 0$ only if $(x_{\text{eq}})_e > 0$. This principle is illustrated graphically in Figure 2.4b for the graph in Figure 2.4a. The picture is a cartoon of the positive orthant of the five dimensional space of edges flows but it illustrates

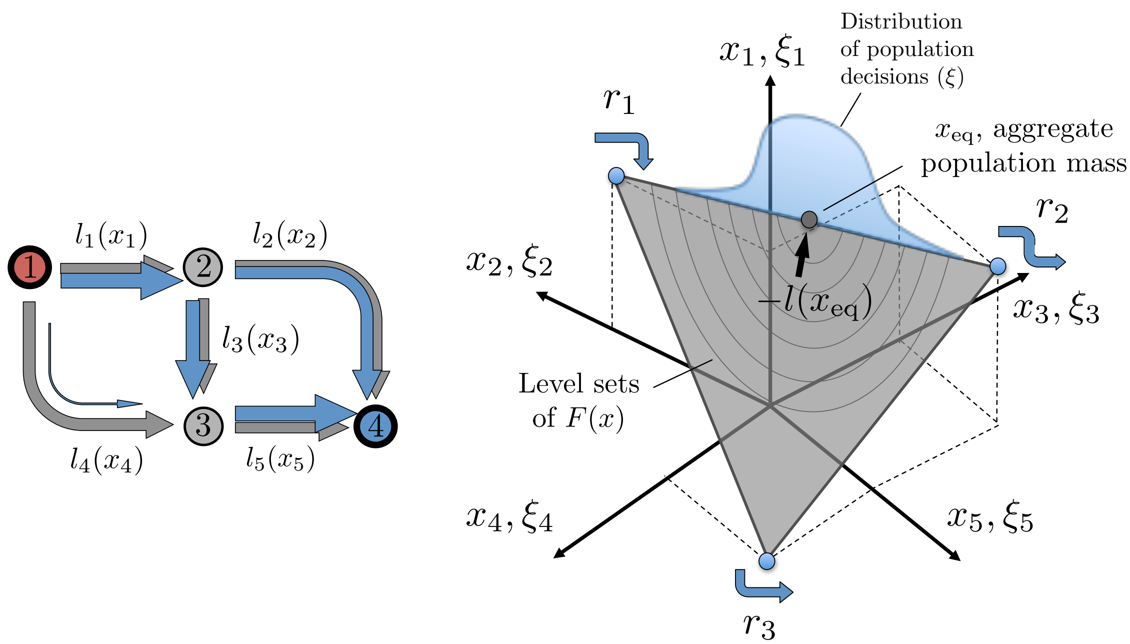

(a) (b)

Figure 2.4: Illustrations of potential game at equilibrium with an example population choice distribution. Note any other distribution for ξ centered around x_{eq} is possible.

that the negative gradient of the potential function at the equilibrium points directly into the edge between routes 1 and 2. An individual driver can choose between routes 1 and 2 with any probability and experience the same travel time (obtain the same value for (2.21)). The equilibrium condition is that the population mass must be balanced around a specific unique mixed strategy so that the aggregate population distribution does not shift. Individual agents can choose any mixed strategy along the edge but if non-negligible amount shift toward one path, the other path will become less congested and more appealing. This balancing interpretation is illustrated in Figure 2.5.

Variable demand routing game

An extension of the standard routing game model that is often considered is the *variable demand case* where travel through the network is modeled as some commodity with a demand curve that is a decreasing function of price. This type of model has been considered since early on in the literature [17, 20, 21]. A sample demand curve is illustrated in Figure 2.6. The price of traveling through the network is the travel time. In this case, we want to determine the population mass that plays the routing game for a given travel time. In this new formulation of the game, we include the total population mass as a new variable. Going back to the route formulation of the game, we know that the minimum travel time through

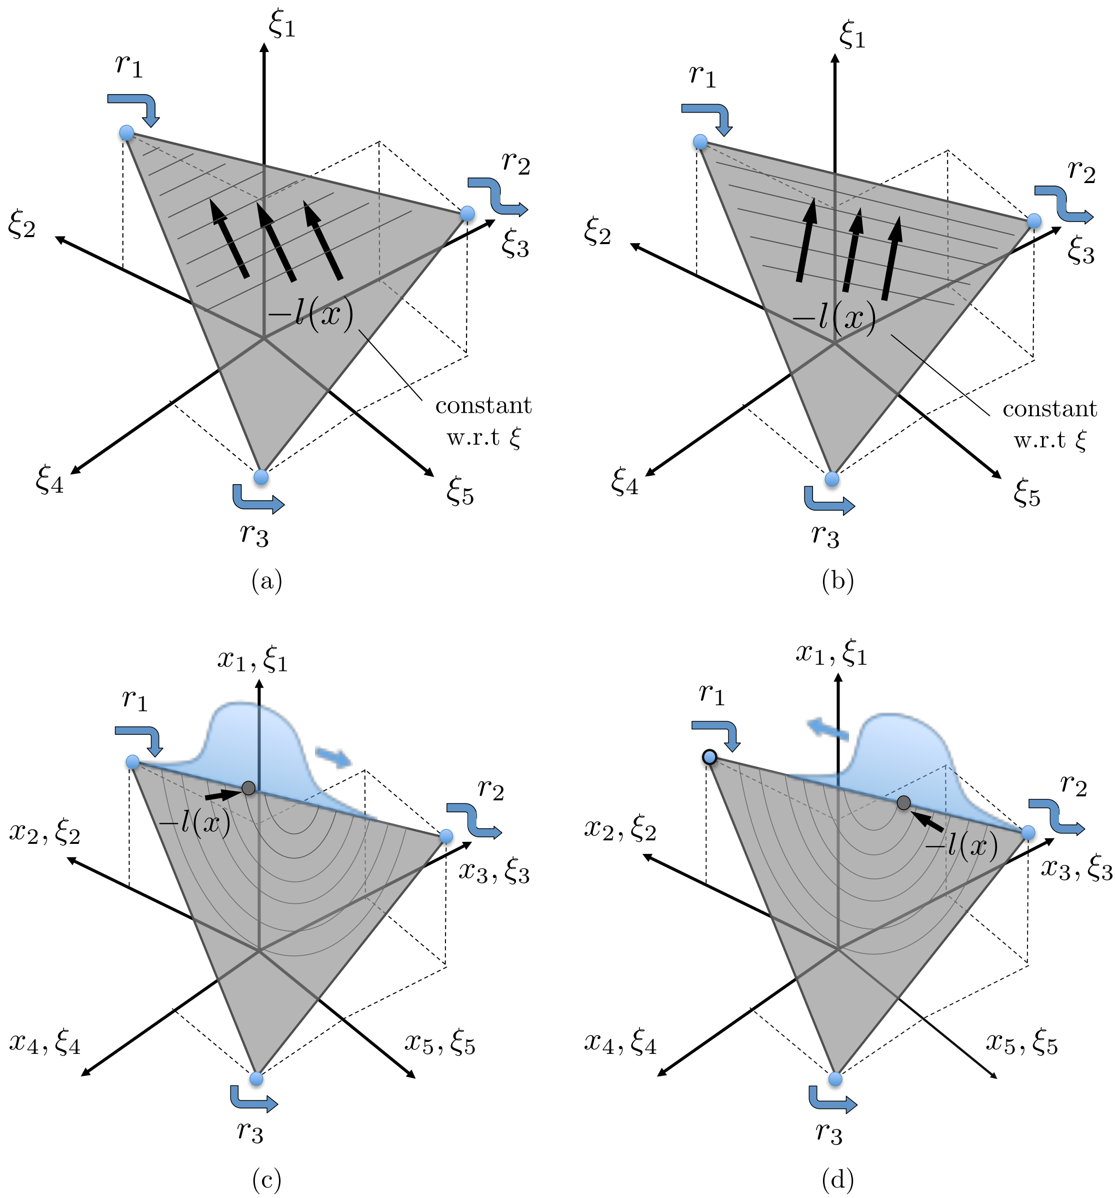

Figure 2.5: Illustrations of relationship between shortest path LP and edge formulation of routing game for the graph shown in Figure 2.4a. (a) Shortest path LP with minimizer r_1. (b) Shortest path LP with minimizers r_1 and r_2 ($\ell_{r_1} = \ell_{r_2}$). (c-d) Unbalanced mass distributions x and how they push the choices of population members.

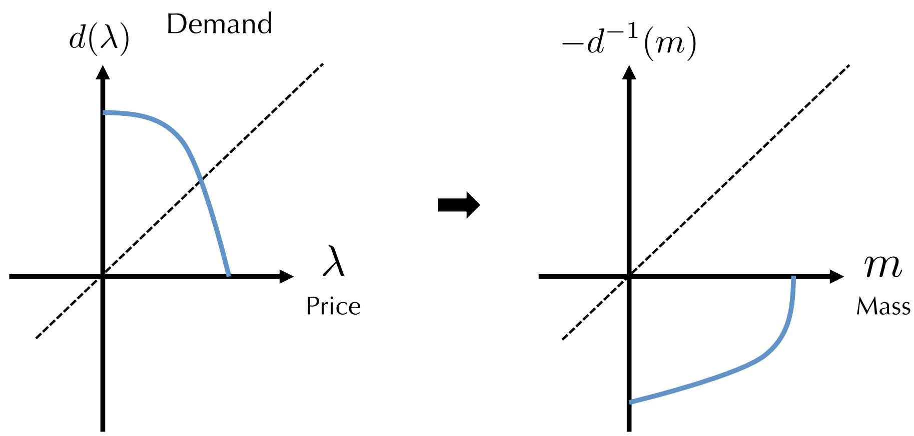

Figure 2.6: Sample demand curve and inverse curve for augmenting the potential function.

the network at equilibrium is encoded by the Lagrange multiplier λ. Thus our goal is to allow the total mass to vary and to design a new potential function so that $\lambda = d^{-1}(m)$. To this end we frame the following optimization problem

$$\min_{z,m} \quad F(z) - \int_0^m d^{-1}(u)\, du \tag{2.22a}$$

$$\text{s.t.} \quad \sum_r z_r = m, \quad z \geq 0, \quad m \geq 0 \tag{2.22b}$$

Considering the first order optimality conditions with respect to z_r and m, we get

$$\ell_r(z) = \lambda + \mu_r \tag{2.23a}$$

$$\lambda = d^{-1}(m) + \nu \tag{2.23b}$$

where ν is the Lagrange multiplier for the constraint $m \geq 0$. Equation (2.23b) says that the price is consistent with the demand whenever the mass is positive. If the mass equals zero, the price is greater than the demand dictates as expected.

2.3 Social Cost and Price of Anarchy

We minimize the potential function to compute the equilibrium distribution, but we will also be interested in the value of the social cost function which is the sum of the costs of the different strategies weighted by the portion of the population that experiences them. It can also be thought of as the average cost modulo dividing by the total mass. For the standard potential game, it is given by

$$J(z) = \sum_r z_r \ell_r(z) \tag{2.24}$$

In the routing game, we can compute it easily using the route flows or the edge flows. We will often abuse notation and write $J(x)$ as well as $J(z)$.

$$J(z) = z^T \ell(z) = z^T \mathbf{E}_{\mathcal{R}}^T l\big(\mathbf{E}_{\mathcal{R}} z\big) = x^T l(x) = J(x) \tag{2.25}$$

Along with the equilibrium flows, we can talk about the socially optimal flows which are calculated by minimizing the social cost function with respect to the constraints rather than the potential function. These flows are the ones that drivers would take if they were all collaborating to minimize the average travel time. We will use J_{eq} and J_{opt} to refer to the social cost at the equilibrium and socially optimal flows respectively.

In general, the Wardrop equilibrium does not optimize the social cost function. The well studied *price of anarchy* [22–24] defined for the routing game as

$$\text{PoA} = \frac{J_{\text{eq}}}{J_{\text{opt}}} \tag{2.26}$$

gives the inefficiency of the equilibrium. Clearly, PoA ≥ 1.

We also note here since the potential function is chosen so that it's gradient is given by $\nabla F = l(x)$, there is a close relationship between the social cost function and the potential function. Indeed, there is a variational inequality formulation of the Wardrop equilibrium that involves the social cost at equilibrium.

Theorem 2 (Variational Inequality Characterization of Wardrop Equilibrium [25]) *If z (or x in the edge flow formulation) is a Wardrop equilibrium then*

$$z^T \ell(z) \leq z'^T \ell(z) \tag{2.27}$$

$$x^T l(x) \leq x'^T l(x) \tag{2.28}$$

for any other feasible route flow z' (or edge flow x').

The proof is a direct application of the Wardrop equilibrium principle. We show it here for the route formulation but the argument is the same in the edge formulation.

Proof 4 *Since z is a Wardrop equilibrium, there exists $\lambda \in \mathbb{R}$ and $\mu \in \mathbb{R}_+^{|\mathcal{R}|}$ such that*

$$\ell(z) = \mathbf{1}\lambda + \mu, \qquad \mu^T z = 0 \tag{2.29}$$

It follows then that

$$z'^T \ell(z) = z'^T \mathbf{1}\lambda + z'^T \mu = m\lambda + z'^T \mu = z^T \mathbf{1}\lambda + z^T \mu + z'^T \mu \geq z^T \ell(z) \tag{2.30}$$

This characterization of the social cost at equilibrium will be useful in price of anarchy analysis.

In order to have some hope of bounding the price of anarchy, we will need to limit it ourselves to specific classes of latency functions such as linear or polynomial with positive coefficients. A useful bound for different latency function classes is the *Pigou bound* defined as follows.

Definition 4 (Pigou Bound [23, 26]) *Let $\mathscr{L}$ be a nonempty set of cost functions. The* Pigou bound *for the class of functions, $\alpha(\mathscr{L})$, is given by*

$$\alpha(\mathscr{L}) = \sup_{l \in \mathscr{L}} \sup_{x, x' \geq 0} \quad \frac{x \cdot l(x)}{x' \cdot l(x') + (x - x')l(x)} \tag{2.31}$$

$$\tag{2.32}$$

This bound seems complicated but it is actually finite and fairly easy to compute for several interesting classes of functions

Proposition 2 ([23, 26]) *Pigou bounds for various function classes*

- *[23] $\mathscr{L} := \{ax + b : a, b \geq 0\} \quad \Rightarrow \quad \alpha(\mathscr{L}) = \frac{4}{3}$*
- *[26] $\mathscr{L} :=$ concave functions $\quad \Rightarrow \quad \alpha(\mathscr{L}) = \frac{4}{3}$*
- *[23] $\mathscr{L} :=$ polynomials with nonnegative coefficients and degree at most p*

  $$\Rightarrow \quad \alpha(\mathscr{L}) = \left[1 - p(p+1)^{-(p+1)/p}\right]^{-1}$$

  which grows as $p \ln(p)$.

Using the Pigou bound and the variational characterization of the equilibrium, we can provide a straightforward upper bound on the price of anarchy.

Theorem 3 (Price of Anarchy Upper Bound [26]) *For a nonempty set of cost functions $\mathscr{L}$,*

$$PoA = \frac{J_{eq}}{J_{opt}} \leq \alpha(\mathscr{L}) \tag{2.33}$$

Proof 5 *Rearranging* (2.31) *gives*

$$x'^T l(x') \geq \frac{x^T l(x)}{\alpha(\mathscr{L})} + (x' - x)^T l(x) \tag{2.34}$$

Let x' be the socially optimal flow and let x be the equilibrium flows. Applying (2.34) *and Proposition 2 gives*

$$J_{opt} = x'^T l(x') \geq \frac{x^T l(x)}{\alpha(\mathscr{L})} + (x' - x)^T l(x) \geq \frac{J_{eq}}{\alpha(\mathscr{L})} \tag{2.35}$$

Earlier more complicated proofs of this bound appeared in [23, 27].

Chapter 3

Queue-Routing Game

Our goal in this chapter is to understand the impact that urban drivers looking for street parking have on overall congestion. To this end, we extend the routing game framework to include populations of city-goers that choose both their route through the network and which block face they want to park on. Previous attempts to analyze how parking can be used to impact congestion in urban centers has been limited to simple road-network topologies and or simple stochastic models of parking behaviour [28].

Studies have shown that a significant amount of peak congestion in many city centers is due to traffic circling looking for on street parking [29]. In our new model, an individual population of drivers travels to an abstract destination we call an *attraction* rather than a specific destination node in a network. Each attraction has several block faces (a collection of nodes and edges) called *parking areas* that drivers can choose from. Borrowing results from observable queueing games, we model each parking area as a continuous queue that parks join as they wait for a spot to open up. Drivers choose their parking area and the route they use to get there based on the likelihood of finding a space and the congestion along the route. At the moment, most drivers do not make travel decisions based on how crowded a parking area is but as parking apps become more ubiquitous, waiting time for a space to open up is very likely to become part of a driver's decision making process. When a certain amount of population mass travels to a parking area it is a distributed around the edges of that parking area modeling the effect of circling traffic on the congestion of those links.

Along with giving us a better picture of city congestion, this model gives us access to a ready made way to affect congestion in cities, namely parking prices in different areas. Currently, the most common way for urban municipalities to address congestion problems is through congestion charges. These have been implemented in many of the states and cities within the U.S. as well as in international cities such as London, Singapore, and Stockholm among others with varying levels of success [30–32]. Congestion charges have long been touted by economists to be a successful, if not *the* successful, mechanism for decreasing congestion, and yet the theory has confronted stiff opposition from the public, with criticism that they disproportionately target the poor, push traffic towards more residential neighborhoods, and have a negative impact on local economies by incentivising people to stay away from urban areas [33, 34]. In an area where much of the traffic is due to circling, increasing parking prices could significantly improve congestion in that area as drivers look to park elsewhere. Adjusting parking tolls is more palatable to the public than standard congestion prices since they are expecting to pay a price anyway. Parking pricing are also immediately implementable using largely existing infrastructure. Many pilot programs to test pricing schemes have been implemented (see, e.g., [35–37]). In one such pilot conducted in San Fransisco—*SFpark*—it was shown that changing the pricing of parking does lead to changes in parking-related behavior [38]. One of the main applications of our model would be solving the parking tolling problem.

The rest of this chapter is organized as follows. In Section 3.1, we present the queueing model for on-street parking. The queue model is used to inform the new queue-routing game that is presented in Section 3.2. We show that this new routing game is still a potential

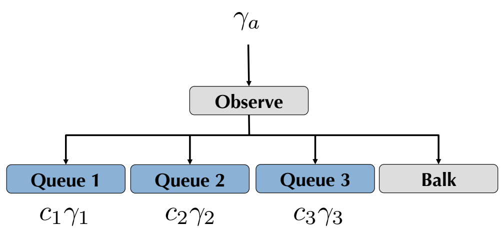

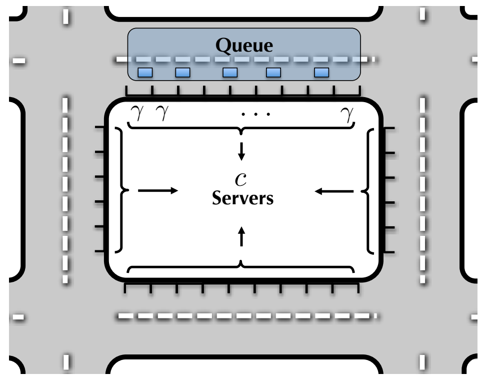

(a) (b)

Figure 3.1: Queue model. (a) Observable queueing game structure. (b) Block face as queue served by c parking spaces.

game and present an optimization problem for finding the equilibrium. In Section 3.3 we present simulation results using real-world networks taken from the downtown Seattle area. The results of this chapter were first published in [39].

3.1 Inspiration: Observable Queueing Game

Discrete observable queueing game

To model on-street parking, we draw inspiration from a simple observable queue game in which arriving customers observe the queue length and choose to join or not based on their expected utility. The expected utility of joining the parking queue is a function of the reward for having parked, the cost of additional wait time due to circling, and the cost of parking itself. The queueing game is illustrated in Figure 3.1a.

We use this queue game to inform the additional cost we will add to the routing game to account for populations of potential parkers. Abstractly, the queue length represents the amount of parking related congestion on a collection of roadways within a parking area. Each individual queue models a block face (illustrated in Figure 3.1b) with c available parking spaces that each become available at a rate of γ (1/average parking time). Together, γ and c define the overall service rate of the queue. In general customers also have the option of balking, i.e. choosing not to join the queue, and receiving some fixed cost (the cost of inconvenience for not joining). We can think of this option as simply a separate queue with a reward that does not depend on queue length.

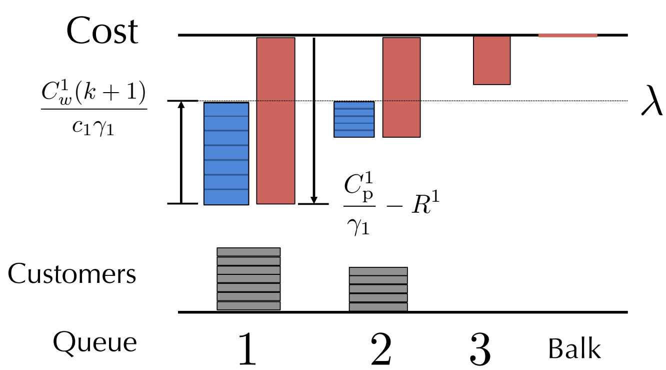

Figure 3.2: Observable queue game equilibrium condition.

We define the expected cost for parking (joining a queue) as

$$q(k) = -R + \tfrac{C_w(k+1)}{\gamma c} + \tfrac{C_{\rm p}}{\gamma}. \tag{3.1}$$

where k is the length of the queue, C_w is the cost of waiting per unit time, and $C_{\rm p}$ is the cost of parking per unit time.

Our goal in the classical queueing game is to determine the equilibrium length of each queue and the *balking level* which is the length of queues such that no more customers will join any of the queues since the cost of waiting is too high relative to the reward of joining. The equilibrium condition parallels the path formulation routing game equilibrium condition, each non-empty queue has equal cost and any empty queue has higher cost. We illustrate the equilibrium condition in Figure 3.2. As in the routing game, $\lambda = \min_i\{q_i(k)\}$ illustrated in Figure 3.2 is the cost of joining a queue at equilibrium. We can think of balking as joining a separate queue with a utility of 0. The *balking level* is the minimum total number of customers such that $\lambda \geq 0$. (Further details on queueing games can be found in [40].)

Continuous analog

In our setup, we want to incorporate the queueing cost in (4.76) into a continuos routing game framework. To this end, we relax the discrete nature of the queue length and define the cost experienced by a driver entering parking area p as

$$C^p(u) = \tfrac{C_{\rm p}^p}{\gamma^p} + \tfrac{C_w^p}{\gamma^p c^p} u \tag{3.2}$$

where u is the continuous mass of drivers in the parking area. Note that since $C_{\rm p}^p/(\gamma^p c^p)$ is positive, $C^p(\cdot)$ is strictly increasing.

Remark 6 *We note that this queueing model is rather stylized and slightly simplistic for modeling circling traffic. We detail a more accurate model that could be incorporated in the routing game framework in Section 4.4.*

3.2 Queue-Routing Game

Flow network model

We now use the queueing framework to extend a classical routing game to model the case where drivers choose their parking area as well as their route through the network. Along with the road graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, we define a set of *attractions*, $\mathcal{A}$, that drivers travel to as well as a set of *parking areas*, $\mathcal{P}$. A parking area $p \in \mathcal{P}$ consists of a set of nodes denoted $\mathcal{N}_p \subset \mathcal{N}$ and the edges that connect them denoted $\mathcal{E}_p \subset \mathcal{E}$. We will define indicator vectors for each of these sets respectively $\mathbf{N}_p \in \{0,1\}^{|\mathcal{N}|}$ and $\mathbf{E}_p \in \{0,1\}^{|\mathcal{E}|}$:

$$(\mathbf{N}_p)_i = \begin{cases} 1, & \text{if node } i \text{ is in parking area } p \\ 0, & \text{otherwise} \end{cases} \tag{3.3}$$

$$(\mathbf{E}_p)_e = \begin{cases} 1, & \text{if edge } e \text{ is in parking area } p \\ 0, & \text{otherwise} \end{cases} \tag{3.4}$$

An attraction $a \in \mathcal{A}$ is a drivers ultimate travel goal. Each attraction has an associated set of parking areas $\mathcal{P}^a$ that drivers can choose from. An individual population starts at a single origin node $o \in \mathcal{N}$ and travels to a specific attraction $a \in \mathcal{A}$. We will denote the population mass associated with this *origin-attraction* pair as m_o^a. The size of these populations are given *a priori*. We note that several attractions may share parking areas.

We use $d \in \mathcal{N}_p$ to denote a (destination) node drivers going to parking area p travel to in order to enter that parking area. For a given origin o and destination d, let $\mathcal{R}_{od}$ be the set of all routes from o to d. For a specific route $r \in \mathcal{R}_{od}$, we denote the set of edges in that route by $\mathcal{E}_r \subset \mathcal{E}$ and an indictor vector for that set of edges $\mathbf{E}_r \in \{0,1\}^{|\mathcal{E}|}$.

$$(\mathbf{E}_r)_e = \begin{cases} 1, & \text{if edge } e \text{ is in route } r \\ 0, & \text{otherwise} \end{cases} \tag{3.5}$$

We illustrate the sets of attractions $\mathcal{A}$, parking areas $\mathcal{P}$, and entry nodes $d \in \mathcal{N}_p$ in Figure 3.3. We note that this framework can include through traffic (as in a traditional routing game) by adding attractions that have a single parking area that consists of a single node.

For each population associated with an origin-attraction pair (o, a), we define a set of strategies

$$U_o^a = \{(p, d, r) |\ p \in \mathcal{P}^a,\ d \in \mathcal{N}_p,\ r \in \mathcal{R}_{od}\}. \tag{3.6}$$

Each driver traveling from an origin o to an attraction a chooses a parking area p, an entry node d, and a route to that entry node r. We use $(z_{od}^{ap})^r$ to denote the subpopulation of population m_o^a that chooses strategy $(p, d, r) \in U_o^a$.

Figure 3.3: Illustration of $\mathcal{A}$ and $\mathcal{P}$ and parking area nodes.

We will also need several other groupings of driver populations: z_{od}^{ap} the total population of drivers traveling to attraction a from o parking in area p entering through node d (but taking any route); z^{ap}, the total population traveling to attraction a and parking in area p; and z^p the total population parking in area p (since some attractions might share parking areas). These masses are calculated as follows.

$$z_{od}^{ap} = \sum_{r\in\mathcal{R}_{od}} (z_{od}^{ap})^r \tag{3.7}$$

$$z^{ap} = \sum_o \sum_{d\in\mathcal{N}_p} \sum_{r\in\mathcal{R}_{od}} (z_{od}^{ap})^r \tag{3.8}$$

$$z^p = \sum_a \sum_o \sum_{d\in\mathcal{N}_p} \sum_{r\in\mathcal{R}_{od}} (z_{od}^{ap})^r \tag{3.9}$$

We use z as a short hand for all populations of drivers.

We model the parking traffic as flowing across the network to their individual parking area and then spreading out uniformly within that parking area. With each population $(z_{od}^{ap})^r$, we can associate a vector $(x_{od}^{ap})^r \in \mathbb{R}^{|\mathcal{E}|}$ that gives the contribution of population $(z_{od}^{ap})^r$ to the flow on each edge of the graph. We define

$$(x_{od}^{ap})^r = \left[\mathbf{E}_r + \frac{1}{|\mathcal{E}_p|}\mathbf{E}_p\right](z_{od}^{ap})^r \tag{3.10}$$

Here, we add the population to each edge on the route through the network and then evenly distribute the mass over the edges in the parking area to model circling behavior. Note that a different model of circling, i.e. drivers spending more time on some edges than others, could easily be incorporated (see Section 4.4 for a more complex model of circling behavior.)

Summing over the various routes and populations, we can compute the flow contributions from z_{od}^{ap} and the total flow

$$x_{od}^{ap} = \sum_{r \in \mathcal{R}_{od}} \left[\mathbf{E}_r + \tfrac{1}{|\mathcal{E}_p|} \mathbf{E}_p \right] (z_{od}^{ap})^r \tag{3.11}$$

$$\bar{x} = \sum_{a} \sum_{o} \sum_{p \in \mathcal{P}_a} \sum_{d \in \mathcal{N}_p} \sum_{r \in \mathcal{R}_{od}} \left[\mathbf{E}_r + \tfrac{1}{|\mathcal{E}_p|} \mathbf{E}_p \right] (z_{od}^{ap})^r \tag{3.12}$$

We will write $(x_{od}^{ap})_e^r$, $(x_{od}^{ap})_e$, and $\bar{x}_e$ to denote the eth element of each of these vectors, i.e. the portion of each flow on edge e.

We define the cost associated with a particular strategy $(p, d, r) \in U_o^a$ as

$$(\ell_{od}^{ap})^r(z) = \underbrace{\textstyle\sum_{e \in \mathcal{E}^r} \tau l_e(\bar{x}_e)}_{\text{travel latency}} + \underbrace{\tfrac{1}{|\mathcal{E}_p|} \textstyle\sum_{e \in \mathcal{E}_p} \tau l_e(\bar{x}_e)}_{\text{circling latency}} + \underbrace{(C^p(z^p) - R^{ap})}_{\text{parking cost}} \tag{3.13}$$

The *travel latency* is the cost associated with traveling to the parking area. The *circling latency* is the cost associated with the inconvenience of circling in a congested area; it can be interpreted as the average latency on a link within the parking area. The *parking cost* is the cost associated with parking in a particular area which is a combination of the waiting time for a space to open up, the monetary cost of parking, and the reward for parking (proximity to attraction, etc) in a particular area. τ (in units of money/time) is a parameter that represents the population's time money tradeoff, R^{ap} is the reward for parking in area p for drivers traveling to attraction a, and $C^p(u)$ is the cost of parking derived from the queueing model defined by Equation (3.2). Note that having separate rewards for parking-attraction pairs allows us to model parking areas shared between several attractions to be convenient for some and inconvenient for others. We also note that the circling latency is separate from the parking cost. It is meant to represent the inconvenience of circling in a congested area as opposed to the time waiting for a space to open up.

Queue-routing equilibrium

Definition 5 *We say a population distribution z is a* Wardrop Equilibrium *of the queue-routing game if and only if it is feasible and for any (o, a) pair, any two strategies (p, d, r), $(p', d', r') \in U_o^a$ such that $(z_{od}^{ap})^r > 0$ and*

$$(\ell_{od}^{ap})^r(z) \leq (\ell_{od'}^{ap'})^{r'}(z) \tag{3.14}$$

We note that this is the standard definition of a Wardrop equilibrium applied to the costs in Equation (3.13).

We now present a potential function for the queue-routing game. We define the function.

$$F(\bar{x}, z) = \textstyle\sum_e \int_0^{\bar{x}_e} \tau l_e(u)\, du + \sum_p \int_0^{z^p} C^p(u)\, du + \sum_{a,p} \int_0^{z^{ap}} -R^{ap}\, du \tag{3.15}$$

We will abuse notation sometimes and write $F(z)$ since $\bar{x}$ is a function of z.

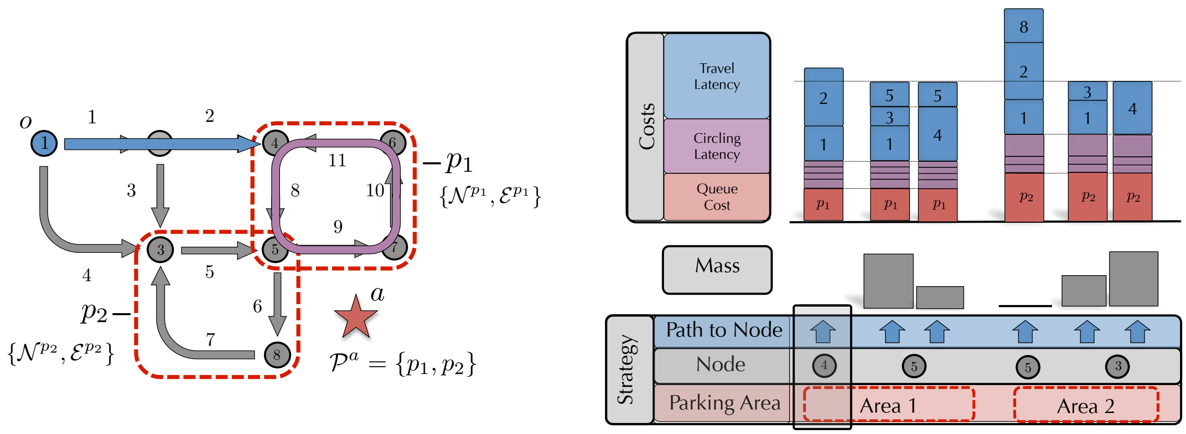

(a) (b)

Figure 3.4: Illustration of the queue-routing Wardrop equilibrium condition. (a) Sample graph with parking areas labeled. (b) Equilibrium condition.

Theorem 4 *Equation* (3.15) *defines a potential function for the queue-routine game.*

Proof 6 *Differentiating Equation* (3.15) *with respect to* $(z_{od}^{ap})^r$*, applying Leibniz integral rule using Equations* (3.8)*,* (3.9)*, and* (3.12)*, we find that derivative is equal to the cost defined in Equation* (3.13) *thus* (3.15) *is a potential function.*

Remark 7 *We note that the circling latency component of the cost is critical to making this game a potential game. It arises from the fact that parking traffic contributes to the overall congestion vector* $\bar{x}$ *and thus differentiating the first set of terms in the potential function with respect to the mass in each parking area produces the circling latency terms. Intuitively, it is important that each member of the population experience any cost to which their mass contributes in order for the game to be a potential game. Whether this is a strength or weakness of the model depends on whether or not one thinks circling drivers actually care about congestion within their parking area.*

We now formulate an optimization problem for finding the equilibrium flow distributions. Rather then enumerating all possible routes for each (o, d) pair and solving for $(z_{od}^{ap})^r$, we can use the edge formulation of the routing game and solve directly for x^{ap}_{od} and z^{ap}_{od}.

$$\begin{aligned}
\min_{x,z} \quad & F(x,z) && (3.16a)\\
\text{s.t.} \quad & Gx^{ap}_{od} = S_{od} z^{ap}_{od}, \quad \forall\ o,d,a,p && (3.16b)\\
& m^a_o = \textstyle\sum_{p\in\mathcal{P}^a}\sum_{d\in\mathcal{N}_p} z^{ap}_{od}, \quad \forall\ o,a && (3.16c)\\
& x^{ap}_{od} \geq 0, \quad \forall\ o,d,a,p && (3.16d)\\
& z^{ap}_{od} \geq 0, \quad \forall\ o,d,a,p && (3.16e)\\
& \bar{x} = \textstyle\sum_{a,o}\sum_{p\in\mathcal{P}^a}\sum_{d\in\mathcal{N}_p}\left[x^{ap}_{od} + \frac{1}{|\mathcal{E}_p|}\mathbf{E}_p z^{ap}_{od}\right] && (3.16f)\\
& z^p = \textstyle\sum_a\sum_o\sum_{d\in\mathcal{N}_p} z^{ap}_{od}, \quad \forall\ p && (3.16g)
\end{aligned}$$

where S_{od} is a source-sink vector defined as

$$(S_{od})_i = \begin{cases} 1 & ;\ \text{if } i = o \\ -1 & ;\ \text{if } i = d \\ 0 & ;\ \text{otherwise} \end{cases} \tag{3.17}$$

Again note that m^a_o is given a priori.

For the queue-routing game, we can calculate the social cost as

$$J(z) = \sum_{a,o,p,d,r} (z^{ap}_{od})^r (\ell^{ap}_{od})^r(z) \tag{3.18}$$

We will also be interested in just the routing portion of the social cost which captures the average congestion that drivers experience. We can calculate this value as

$$J_{\mathcal{R}}(z) = \sum_e \bar{x}_e l_e(\bar{x}_e). \tag{3.19}$$

As a social planner, we might consider adjusting parking prices in order to lower the average congestion.

3.3 Examples: Downtown Seattle

To demonstrate the usefulness of the routing game for queue–flow networks, we explore two examples using different regions in the Seattle downtown area and its arterials as the basis for the network topologies.

In each example, we measure flows on each edge ($\bar{x}_e$) in cars per unit time. We use linear latencies that were derived from the Bureau of Public Roads (BPR) *link performance function* which is given by

$$l^{\text{BPR}}_e(\bar{x}_e) = t_e\left(1 + 0.15\left(\frac{\bar{x}_e}{\kappa_e}\right)^4\right) \tag{3.20}$$

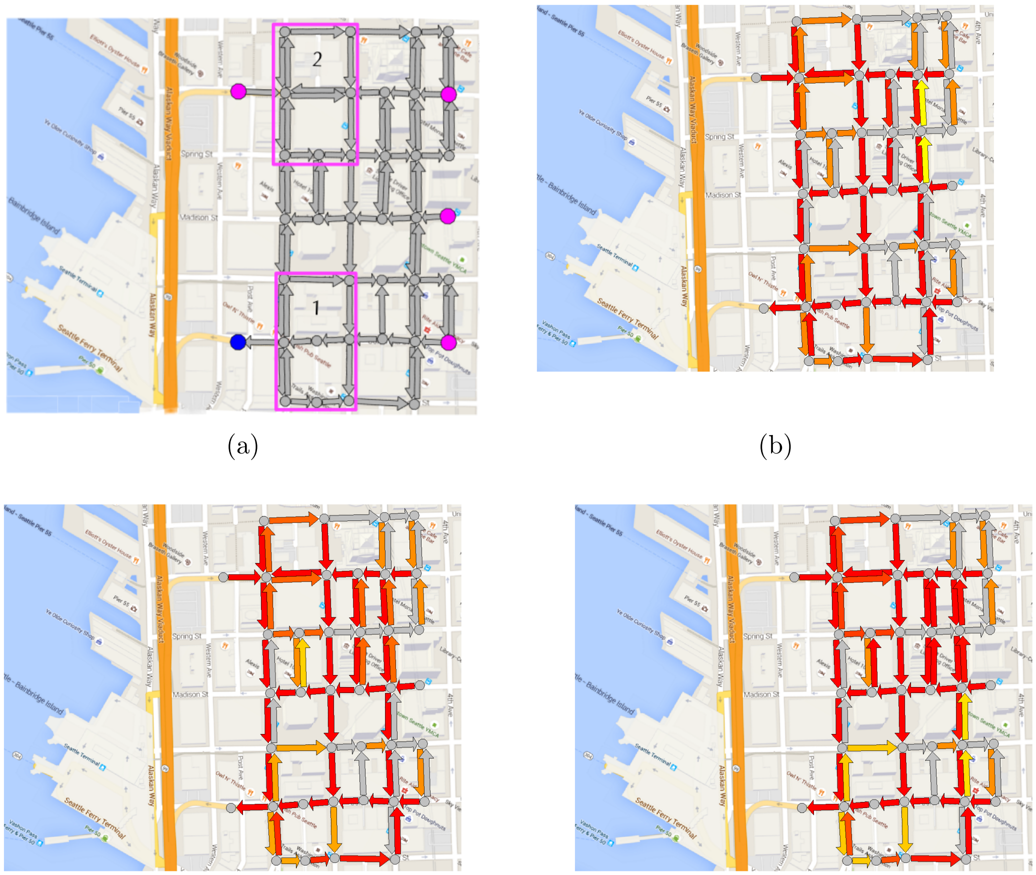

(a) (b)

(c) (d)

Figure 3.5: Setup and results for Example 1 (SR-99). (a) Queue–flow network: Magenta nodes are parking population sources and the magenta boxes are the parking areas. Through traffic begins at every node and flows to the destination node shown in blue. (b-d) Flow for Example 1 (SR-99) with $C_{\rm p}^2 = 0.01$ in area 2 and (b) $C_{\rm p}^1 = 0.01$, (c) $C_{\rm p}^1 = 0.124$, and (d) $C_{\rm p}^1 = 0.2$ in area 1.

where t_e is the free-flow travel time on link e (length/speed limit) and κ_e is the capacity of link per unit time [41]. We heuristically take κ_e to be

$$\kappa_e = \tfrac{50\ \text{cars}}{\text{mi}} \times \tbinom{\text{speed}}{\text{limit}} \tag{3.21}$$

assuming cars are travelling in *free-flow* with approximately 50 cars per mile (approximately 100 feet per car). We chose the linear latency that agrees with this function at $\bar{x}_e = 0$ and $\bar{x}_e = 3\kappa_e$.

$$l_e(\bar{x}_e) = t_e \left(1 + 4\tfrac{\bar{x}_e}{\kappa_e}\right). \tag{3.22}$$

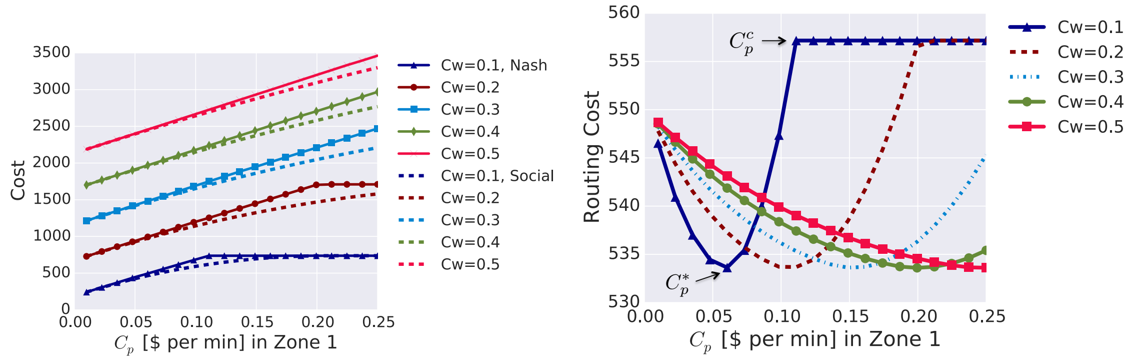

(a) (b)

Figure 3.6: (b) The social cost evaluated at the equilibrium and the socially optimal solution as a function of the cost of parking $C_{\rm p}^1$ in area 1 (with $C_{\rm p}^2 = 0.01$ fixed) for $C_w = C_w^1 = C_w^2 \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. (c) The portion of the social cost due to routing under the Nash equilibrium as a function of the cost of parking $C_{\rm p}^1$ in parking area 1 (with $C_{\rm p}^2 = 0.01$ fixed) for $C_w \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$. We point out for $C_w = 1$ the critical point $C_{\rm p}^c$ at which the routing demand is no longer flexible to increases in price and the price $C_{\rm p}^*$ that minimizes the routing costs .

Example 1: SR-99

In the first example, we construct a queue–flow network using a portion of the Seattle downtown area near state route 99 (SR-99), a heavily traversed road, detailed in Figure 3.5a and consider a possible evening rush hours scenario. We consider thru traffic that originates at each node in the grid and heads towards the entrance to freeway SR-99 (the blue node) and we consider parking traffic that originates at each of the magenta nodes and heads towards the two parking areas (the magenta boxes). This would consistent with constituents exiting the downtown area after work while others seek to find parking close to a restaurant near the two parking areas.

In Figures 3.6a and 3.6b, we show the social cost at the Wardrop equilibrium and socially optimal strategies and the routing portion of social cost at equilibrium. The Wardrop-induced cost is always higher than the socially optimal cost as expected. Both costs increase with the cost of waiting $C_w = C_w^1 = C_w^2$ and with the cost of parking $C_{\rm p}^1$. Intuitively, the Wardrop-induced cost plateaus when the price of parking in Zone 1 gets high enough that all parkers park in Zone 2.

In addition, in Figure 3.6b, we see that for each value of C_w, the routing cost obtains a minimum for some value of $C_{\rm p}^1$ (indicated for $C_w = 0.1$ by $C_{\rm p}^*$). These points represent the optimal price of parking that a municipal service provider should charge if its objective is to minimize the latency experienced by the total population—note this objective might not align with minimizing social cost. Moreover, as the cost of waiting C_w increases, the price of parking that minimizes the routing cost becomes larger. We remark that in Seattle,

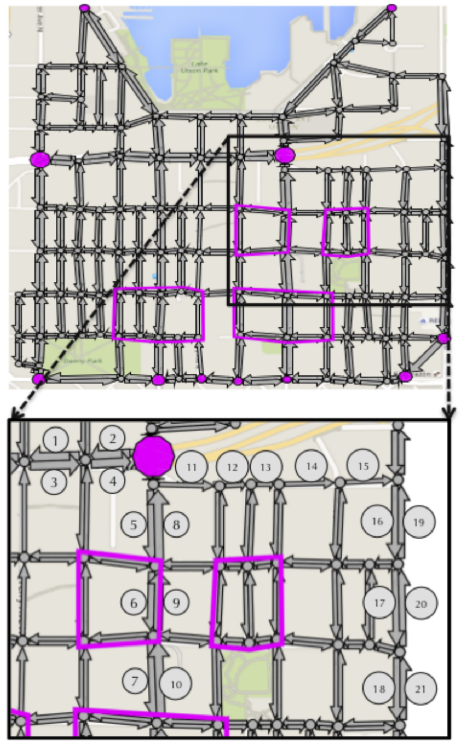

(a)

Figure 3.7: Structure of Amazon area simulation. Magenta dots are traffic sources (with radius indicating size relative to other sources.) Parking areas are polyhedra. The lower figure shows location of the edges from Figure 3.9

like many municipalities, there are regulatory constraints on the maximum value that can be charged per hour for on-street parking. This value tends to be around \$7; hence, if the cost of waiting is too large, then it may not be possible to optimally design the price of parking to minimize the latency experienced by the total population. This suggests that understanding preferences of money over waiting (time spent circling or in traffic) should be better understood and perhaps, incorporated into regulatory policies that cap parking prices.

In Figure 3.5, we show the total traffic flow (parking plus throughput populations) for three different values of C_{p}^1. It can be seen that as C_{p}^1 increases, the flow shifts from being evenly distributed between the two parking areas to largely being in Zone 2 (the top area).

Example 2: Amazon campus

The next example we explore is the effect of parking on congestion in the region around Amazon's headquarters. This area has seen increased congestion over the last half decade

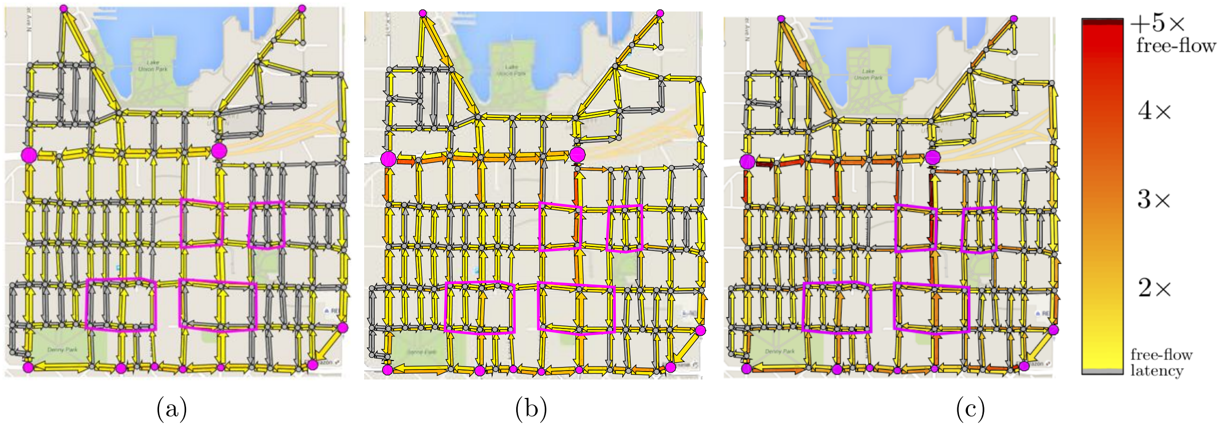

Figure 3.8: Evolution of the congestion in the Amazon Campus area network as the parking populations increase. Through traffic population is fixed at 75 cars/min (a) No parking population (b) Total parking population of 150 cars/min (c) Total parking population of 300 cars/min.

due to Amazon's presence. We hold the total throughput traffic fixed and increase the amount of parking traffic to simulate a potential rush hour scenario as Amazon employees are driving to work.

In Figure 3.7, we show the network graph along with a key section of the network which has interesting routing behaviors that emerge as the proportion of parking-related traffic increases. All traffic enters through the magenta nodes. The through traffic travels (uniformly) to other magenta nodes and the parking traffic travels to a parking area indicated by the boxed magenta regions. The size of the magenta nodes indicates the relative magnitude of both the throughput and parking traffic coming from those nodes. For each parking region, we use the queueing parameters shown in the following table:

R^p($)	C^p_w ($/min)	μ^p (spots/min)	c^p (# spots)	C^p_{p} ($/min)	τ ($/hr)
100	0.1	1/120	50	0.01	30

We fix the total throughput traffic to be 75 cars/minute distributed among the source nodes according to their relative size and we vary the amount of parking traffic in the interval $[0, 300]$ cars/minute.

In Figure 3.9, we show the total flow, the throughput traffic flow, and the parking flow for key edges in the graph enumerated and depicted in Figure 3.7. We remark that the throughput traffic decreases as parking traffic decreases in the parking regions. Also, note how even links that are significantly removed from the parking regions are affected by parking traffic as the network adjusts to the demand. In Figure 3.8, we show qualitatively how the network becomes more congested as the parking populations grow for a fixed through-traffic population size.

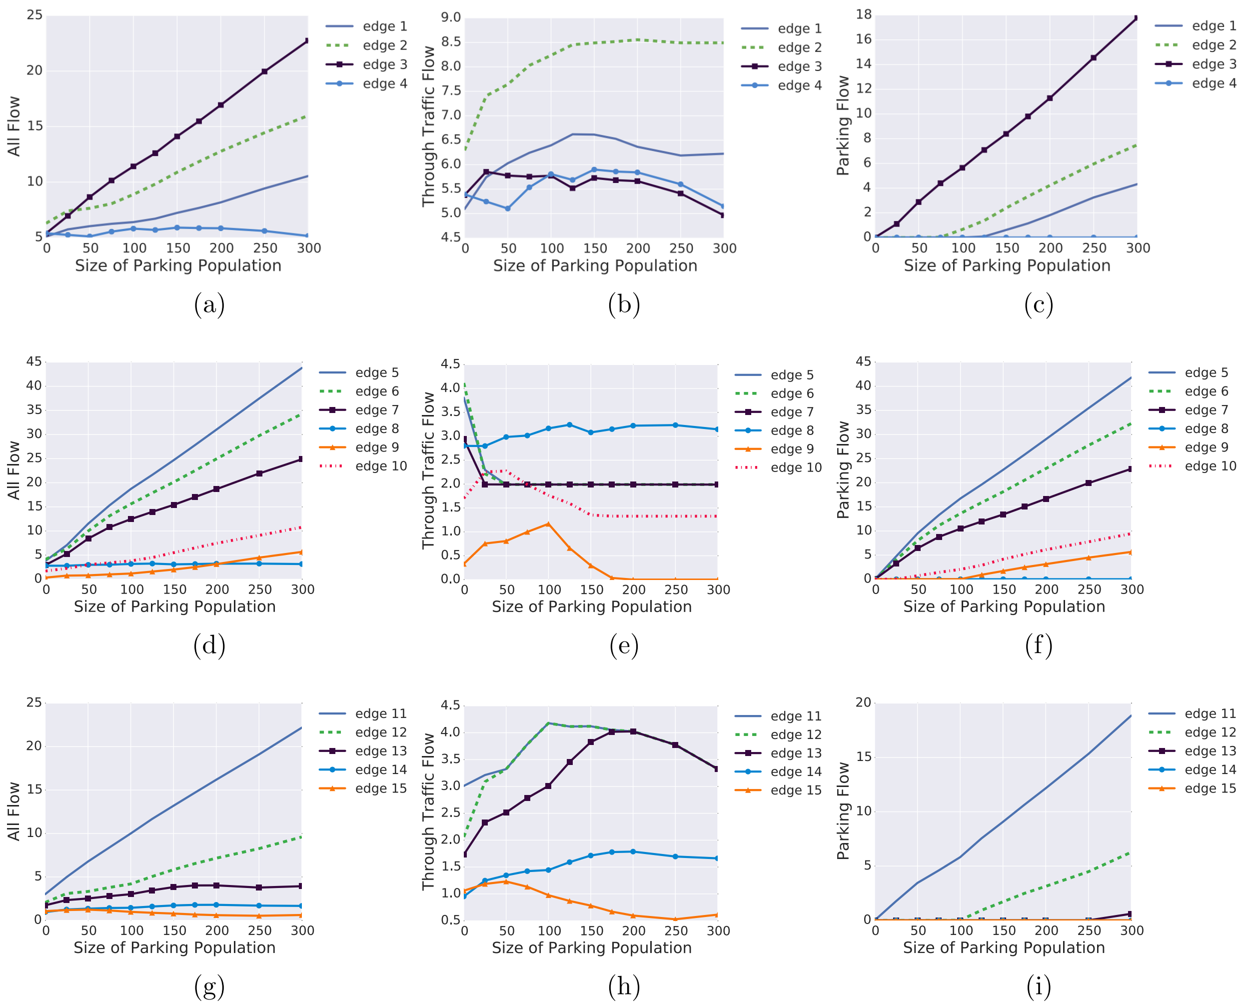

Figure 3.9: Total traffic, through traffic, and parking traffic for Example 2 (Amazon) along various edges of the network. (See Fig. 3.7 for the edge labels.)

The work in the paper is the first steps towards the development of a modeling paradigm for urban mobility that accounts for drivers having different objectives and intended uses of transportation infrastructure. There is still a significant amount of work to create a comprehensive theoretical framework for queue–flow networks. In future, we plan to extend the framework to include more complex queueing behaviors such as balking and jockeying between queues. We also plan to formulate and solve a bilevel optimization problem for pricing design as well as potentially validate a version of the model against real data.

Chapter 4

Markov Decision Process Routing Games

We now study a stochastic continuous population potential game where members of the population solve a Markov decision process (MDP) and the losses (traditionally rewards) are functions of the population distribution. This formulation is analogous to classical routing games where agents solve an MDP as opposed to a shortest path problem.

Stochastic games, games where individual players solve an MDP, have been around as long as routing games. They were first studied by Shapley in the two-player zero sum case [42] and by others in the finite player non-zero sum case [43–46]. In the finite player case, potential functions have been used to ensure the existence of pure strategy equilibria [47, 48].

Effort has been made on several fronts to simplify the analysis of stochastic games with large numbers of players. *Oblivious equilibria* were introduced by Weintraub [49–52]. In these stochastic games, the payoff of each player is determined by the average state of the other players. Stochastic continuous population games have also been considered, first as *anonymous sequential games* introduced by Rosenthal and Jovanovic [53]. Results have mostly focused on existence and uniqueness of equilibria [54–57] and specific applications [58, 59].

Recently, stochastic continuous population games have been studied as *mean field games* introduced by Lasry and Lions [60–62] and concurrently by Huang, Malhame, and Caines [63, 64]. The classic mean field game formulation assumes a continuous state space and a pair of coupled partial differential equations (PDEs): one backward time PDE that defines the value function or "cost-to-go" for the population of agents and one forward time PDE that defines the mass evolution of the population (the Fokker-Planck equation). As in our case, when agents' individual costs can be written as the gradient of some functional, the game is called a mean field potential game and both PDEs can be solved by solving a single optimal control problem. Mean field games have been studied significantly including numerical schemes for computing equilibria [65–70] and have been applied to model problems in wireless networks [71–77], oil production [78], smart grid technologies [79–81], and airline networks [82] to name a few.

Mean field games with a discrete state space (graph structure) have been specifically considered by Gomes, et.al. [83] in the stationary (infinite horizon) case and by Guéant who specifically focused on the finite horizon potential game case in [84] and congestion effects in [85]. In these discrete state space mean field potential games, the potential functional depends on the mass at each node. A major difference in our formulation is that the potential function will be allowed to depend on the mass at each node, the mass taking each action, and the mass on each edge of the graph. Formulations were the agents' costs depend on the other agents taking the same actions are called extended mean field games [86, 87].

Our formulation can be thought of as a routing game perspective on stochastic population games. Along with providing a more thorough understanding of the problem space, explicitly making this connection allows us to extend traditional routing game concepts on efficiency of equilibria namely the *price of anarchy* and *Braess' paradox* to the stochastic population game setting. Although there have been early results on efficiency of mean-field equilibria [64, 88–90], the literature is much more mature in the routing game community. We extend some of the most significant results from routing games, particularly upper bounds on the price of anarchy, to our framework at the end of the chapter.

The rest of this chapter is organized as follows. In Section 4.1, we present a linear programming formulation of standard Markov decision processes that parallels the linear programming version of the shortest path problem and serves as inspiration for the equilibrium of the game. We then present the game formulation in the infinite and finite horizon cases in Section 4.2 making explicit connections with classical routing games. In Section 4.3, we provide upper bounds for the price of anarchy and comment briefly on Braess' paradox in the MDP routing game setting. In Section 4.4, we explore three examples. The first two are related to ridesharing or taxi drivers competing for customers. We study this problem in both the finite horizon and infinite horizon settings with deterministic and stochastic transitions. In the third example, we consider a more complex model of circling traffic competing for parking spaces. The finite-horizon version of the MDP routing game formulation was first published in [91].

4.1 Inspiration: Linear Programming for Markov Decision Processes

We start out by presenting a linear programming formulation of Markov decision processes that has close parallels with the linear programming formulation of the shortest path problem (Problem (2.21)) and serves as the inspiration for the game we present in the next section. .

Let $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ be a graph with nodes (or states) and edges, and let $\mathcal{N}_j$ be the nodes that can be reached by an edge from node j. Let $\mathcal{A}_j$ be a set of *actions* associated with each node j and let $P_j^a = (P_{ij}^a)_{i \in \mathcal{N}_j} \in \Delta_{|\mathcal{N}_j|}$ be the transition probabilities associated with action a from node j to each node $i \in \mathcal{N}_j$. Here, $\Delta_{|\mathcal{N}_j|}$ is the simplex of dimension $|\mathcal{N}_j|$. We will sometimes refer to the set of all actions as $\mathcal{A} = \cup_j \mathcal{A}_j$ with the understanding that each action $a \in \mathcal{A}$ is only available from a specific node. We will use $\mathcal{P}_{\mathcal{A}} = \{P_j^a\}_{a \in \mathcal{A}}$ to refer to the set of all transition probabilities.

Define *node losses* (traditionally rewards) q_j associated with being in node j, *action losses* r_j^a, associated with taking each action $a \in \mathcal{A}_j$; and *edge losses*, l_{ij}, associated with transitioning from j to i. We will sometimes index the action losses simply by the action $a \in \mathcal{A}$, i.e. $r_a = r_j^a$ again with the understanding that a is only available from node j. We will also sometimes index the set of transition or edges losses by the edge only, i.e. $l_e = l_{ij}$ with the understanding that edge e runs from j to i.

The goal of solving a Markov decision process is to find the optimal mixed strategy over the actions $\mathcal{A}_j$ to choose whenever the agent is in node j. We will denote a mixed strategy at node j as $\eta_j = (\eta_j^a)_{a \in \mathcal{A}_j} \in \Delta_{|\mathcal{A}_j|}$ where $\Delta_{|\mathcal{A}_j|}$ is the simplex of dimension $|\mathcal{A}_j|$, and we will refer to a collection of mixed strategies, $\eta = (\eta_j)_{j \in \mathcal{N}}$, as a policy. We will sometimes index policies by simply the action set as well, i.e. $\eta_a = \eta_j^a$ for $a \in \mathcal{A}_j$. The optimal solution can be found over a finite time horizon or an infinite time horizon. In the finite time horizon case, the objective is generally to minimize the total expected loss. In the infinite horizon case, one can seek to minimize the discounted loss or the average loss. We consider here the average loss infinite horizon case.

A policy η gives rise to a transition matrix $P(\eta) \in [0,1]^{|\mathcal{N}|\times|\mathcal{N}|}$

$$\left[P(\eta)\right]_{ij} = \sum_{a\in\mathcal{A}_j} P^a_{ij}\eta^a_j \tag{4.1}$$

and the resulting stationary distribution $p(\eta) : \eta \mapsto [0,1]^{|\mathcal{N}|}$. We make the following standard assumption to guarantee that the stationary distribution exists, is unique, and describes the long term limit of the Markov chain's behavior.

Assumption 1 *Assume that $P(\eta)$ is irreducible and aperiodic for every pure strategy policy η.*

In the average loss infinite horizon case, we want to optimize the following program.

$$\begin{aligned}
\min_{\eta} \quad & \sum_j \left(\sum_i l_{ij} P_{ij}(\eta) + r^a_j \eta^a_j + q_j \right) p_j(\eta) && \text{(4.2a)}\\
\text{s.t.} \quad & \sum_{a\in\mathcal{A}_j} \eta^a_j = 1, \qquad \eta_j \geq 0 \qquad \forall j && \text{(4.2b)}\\
& P_{ij}(\eta) = \sum_{a\in\mathcal{A}_j} P^a_{ij}\eta^a_j \qquad \forall i,j && \text{(4.2c)}\\
& p_i(\eta) = \sum_j P_{ij}(\eta) p_j(\eta) && \text{(4.2d)}
\end{aligned}$$

The objective here is the expected loss at any given time which depends on the probability of being in a particular node (based on the stationary distribution) at that time and the probability of taking a particular action (based on the chosen policy). The constraints ensure that the stationary distribution is the one arising from the chosen policy. This problem as formulated is nonlinear and difficult to solve. However, if we solve for both the policy and the stationary distribution at the same time by applying a change of variables

$$\xi^a_j = p_j(\eta)\eta^a_j \tag{4.3}$$

we can transform the problem into a linear program. ξ^a_j is the probability of being in node j and choosing the action a. Problem (4.2) can be written as

$$\begin{aligned}
\min_{\xi} \quad & \sum_j \sum_{a\in\mathcal{A}_j} \left(\sum_i l_{ij} P^a_{ij} + r^a_j + q_j \right) \xi^a_j && \text{(4.4a)}\\
\text{s.t.} \quad & \sum_j \sum_{a\in\mathcal{A}_j} \xi^a_j = 1, \qquad \xi \geq 0 && \text{(4.4b)}\\
& \sum_{a\in\mathcal{A}_i} \xi^a_i = \sum_j \sum_{a\in\mathcal{A}_j} P^a_{ij}\xi^a_j \qquad \forall i,j && \text{(4.4c)}
\end{aligned}$$

Given ξ, it is straightforward to solve for η and $p(\eta)$ as

$$\eta_j^a = \frac{\xi_j^a}{\sum_{a\in\mathcal{A}_j}\xi_j^a}, \qquad\qquad p_j(\eta) = \sum_{a\in\mathcal{A}_a}\xi_j^a \tag{4.5}$$

More details of this formulation can be found in [92–94]

It is also straightforward to see that Equation (4.4c) guarantees that $p(\eta)$ is a stationary distribution for $P(\eta)$

$$\sum_{a\in\mathcal{A}_i}\xi_i^a = \sum_j\sum_{a\in\mathcal{A}_j}P_{ij}^a\xi_j^a \tag{4.6a}$$

$$\sum_{a\in\mathcal{A}_i}\xi_i^a = \sum_j\sum_{a\in\mathcal{A}_j}P_{ij}^a\frac{\xi_j^a}{\left(\sum_{a\in\mathcal{A}_j}\xi_j^a\right)}\left(\sum_{a\in\mathcal{A}_j}\xi_j^a\right) \tag{4.6b}$$

$$p_i = \sum_j\sum_{a\in\mathcal{A}_j}P_{ij}^a\eta_j^a p_j \tag{4.6c}$$

$$p_i = \sum_j P_{ij}p_j \tag{4.6d}$$

In order to write Problem (4.4) in matrix form and make connections with the shortest path problem and routing game, we can assign some ordering to the edges of the graph and as well as the action set $\mathcal{A}$. We can then define a transition matrix $\mathbf{P}_\mathcal{A} \in [0,1]^{|\mathcal{E}|\times|\mathcal{A}|}$ that maps ξ to the probability of taking edge e at any given transition.

$$\left[\mathbf{P}_\mathcal{A}\right]_{ea} = \begin{cases} P_{ij}^a & ;\ \text{if } a \text{ is available in node } j \text{ and edge } e \text{ connects } j \text{ to } i \\ 0 & ;\ \text{otherwise} \end{cases} \tag{4.7}$$

Note that $\mathbf{P}_\mathcal{A}$ is column stochastic. Using matrix $\mathbf{P}_\mathcal{A}$, we can compute the vector of probabilities of taking a specific edge in the graph at any given time, $\bar{\xi} \in [0,1]^{|\mathcal{E}|}$, as

$$\bar{\xi} = \mathbf{P}_\mathcal{A}\xi \tag{4.8}$$

We will often be interested in the *fully deterministic* case, where an agent can choose any edge originating from a given node with probability 1. In this case, we have that $\mathbf{P}_\mathcal{A} = I_{|\mathcal{E}|\times|\mathcal{E}|}$ (assuming the proper ordering on the edges and actions) and $\bar{\xi} = \xi$.

Indexing l by edges, q by nodes, and r and ξ by actions in $\mathcal{A}$, we can write Problem (4.4) in matrix form as

$$\min_{\xi} \quad (l^T\mathbf{P}_\mathcal{A} + r^T + q^T I_o\mathbf{P}_\mathcal{A})\xi \tag{4.9a}$$

$$\text{s.t.} \quad \mathbf{1}^T\xi = 1, \qquad \xi \geq 0 \tag{4.9b}$$

$$G\mathbf{P}_\mathcal{A}\xi = 0 \tag{4.9c}$$

Note that here we have used the fact that since $\mathbf{P}_{\mathcal{A}}$ is column stochastic and agrees with the graph structure, we have that

$$\Big[I_o \mathbf{P}_{\mathcal{A}}\Big]_{ja} = \begin{cases} 1 & ; \text{ if } a \in \mathcal{A}_j \\ 0 & ; \text{ otherwise} \end{cases} \tag{4.10}$$

In the fully deterministic case, Problem 4.9 becomes

$$\min_{\xi} \quad (l^T + r^T + q^T I_o)\xi \tag{4.11a}$$

$$\text{s.t.} \quad \mathbf{1}^T \xi = 1, \qquad \xi \geq 0 \tag{4.11b}$$

$$G\xi = 0 \tag{4.11c}$$

Remark 8 *It should be noted that this deterministic case breaks Assumption 1 (irreducibility and aperiodicity). We mention it here to draw connections with the routing game and we make further comments about it in the game context in Remark 13.*

In the form of Problem 4.11, we can see a clear connection with the shortest path linear program presented in Section 2.2. $\xi = \bar{\xi}$ must live in the nullspace of the incidence matrix G which means that ξ must be some positive linear combination of the cycles of the graph $\mathcal{G}$. Let $\mathcal{C}$ denote the set of cycles of the graph and let $\mathcal{E}_c$ be the edges in cycle $c \in \mathcal{C}$. If all the mass is concentrated on that cycle, the mass on each edge is $1/|\mathcal{E}_c|$. Thus the feasible set is the convex combinations of the cycles scaled by their lengths, i.e. convex combinations of the columns of the matrix $\mathbf{E}_{\mathcal{C}}$

$$\big[\mathbf{E}_{\mathcal{C}}\big]_{ec} = \begin{cases} \frac{1}{|\mathcal{E}_c|} & ; \text{ if } e \in \mathcal{E}_c \\ 0 & ; \text{ otherwise} \end{cases} \tag{4.12}$$

Given the scaled indicator matrix $\mathbf{E}_{\mathcal{C}}$, we could rewrite Problem (4.11) in terms of the probability mass assigned to each cycle, $\zeta \in \mathbb{R}_+^{|\mathcal{C}|}$ as

$$\min_{\zeta} \quad (l^T + r^T + q^T I_o)\mathbf{E}_{\mathcal{C}}\zeta \tag{4.13a}$$

$$\text{s.t.} \quad \mathbf{1}^T \zeta = 1, \qquad \zeta \geq 0 \tag{4.13b}$$

Here Problem (4.11) is analogous to the edge formulation of the shortest path problem and Problem (4.13) is analogous to the route formulation where cycles have taken the place of routes.

Remark 9 *One main conceptual difference with the shortest path problem is that in the MDP case, we are seeking to minimize the average loss over time and thus the length of each cycle matters. In the optimization problem, this results in the cycles being scaled by length and mass being divided up uniformly over the edges in a cycle. In the shortest path problem, the number of edges in each path is not important and thus mass conservation is modeled as the entire mass being added to each edge, i.e. the routes are not scaled by their length.*

In the nondeterministic case, we can think of the stationary probability masses in either the action mass space (ξ) or the edge mass space ($\bar{\xi}$). The stationary edge masses $\bar{\xi}$ can still be thought of as a linear combination of cycle flows since $\bar{\xi}$ must still live in the nullspace of G (Constraint (4.35c)). However, the probabilistic transitions put limits on the relative mass that is assigned to each cycle. To illustrate, this we consider the simple graph in Figure 4.3a with three cycles. We first consider fully deterministic transitions. Figure 4.1b illustrates the feasible space of the stationary probability masses ($\xi = \bar{\xi}$), i.e. the convex combination of the cycles of the graph, as well as the linear program for optimizing the MDP. We then consider the case where the transitions are probabilistic with action sets

$$\mathcal{A}_1 = \{a_1\}, \qquad \mathcal{A}_2 = \{a_2\}, \qquad \mathcal{A}_3 = \{a_3, a_4\}, \qquad \mathcal{A}_4 = \{a_5, a_6\} \tag{4.14}$$

with transition probabilities

$$\begin{aligned} \mathcal{P}_{\mathcal{A}} &= \left\{ P_j^a \right\}_{a \in \mathcal{A}} \\ &= \left\{ P_1^{a_1} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix}, P_2^{a_2} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, P_3^{a_3} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}, P_3^{a_4} = \begin{bmatrix} 0 \\ \frac{1}{2} \\ 0 \\ \frac{1}{2} \end{bmatrix}, P_4^{a_5} = \begin{bmatrix} \frac{9}{10} \\ 0 \\ \frac{1}{10} \\ 0 \end{bmatrix}, P_4^{a_6} = \begin{bmatrix} \frac{1}{2} \\ 0 \\ \frac{1}{2} \\ 0 \end{bmatrix} \right\} \end{aligned} \tag{4.15}$$

Given these transition probabilities, the optimizer cannot make the choice to simply circle on cycle 2 anymore. The only way to transition from node 3 to node 2 is to choose action 4 which causes a transition to node 4 fifty percent of the time. Similarly, there is no way to simply circle on cycle 3. Because of this, the space of edge probability masses ($\bar{\xi}$) is limited to the blue region illustrated in Figure 4.1c. The vertices of this region represent the pure strategy policies (where the agent choose a single action with probability 1 at each node). We can also draw the space of feasible action probability masses (ξ) illustrated in Figure 4.1d. The ξ and $\bar{\xi}$ spaces are related by the linear transformation $\mathbf{P}_{\mathcal{A}} \in \mathbb{R}^{|\mathcal{E}| \times |\mathcal{A}|} = \mathbb{R}^{6 \times 6}$ which we produce here for clarity.

$$\mathbf{P}_{\mathcal{A}} = \begin{array}{c} \\ e_1 \\ e_2 \\ e_3 \\ e_4 \\ e_5 \\ e_6 \end{array} \left\{ \overbrace{\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{2} & 0 & 0 \\ 0 & 0 & 1 & \frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{9}{10} & \frac{1}{2} \\ 0 & 0 & 0 & 0 & \frac{1}{10} & \frac{1}{2} \end{bmatrix}}^{a_1 \quad a_2 \quad a_3 \quad a_4 \quad a_5 \quad a_6} \right. \tag{4.16}$$

Note that in this example $\mathbf{P}_{\mathcal{A}}$ is invertible which simplifies the relationship between the two spaces. This need not be the case in general.

This linear programming view illuminates certain properties of MDPs such as the fact that there always exists a pure strategy optimal policy. In general the descent direction (the negative gradient) of the objective will point into a vertex of the feasible space. Even when the negative gradient points directly into a face of the feasible space, any vertex of that face still represents an optimal policy.

This linear program view of MDPs and the connections with the edge flow formulation of the routing game as well as the route flow formulation with the proper analogy between paths and cycles also suggests that we can define a population potential game where the loss functions are functions of the mass of the population and each infinitesimal member of the population is seeking to optimize an MDP. In the next section, we define this game, present the appropriate Wardrop-type equilibrium condition, the appropriate potential function and optimization problem for finding the equilibrium in the infinite horizon and then the finite horizon cases.

4.2 Markov Decision Process Routing Games: Infinite and Finite Horizon Cases

We now present the Markov decision process routing game in the infinite and finite horizon cases. Again for simplicity of notation, we only assume a single population of agents; however, as in classic routing games, the results hold for multiple populations in both the infinite and finite horizon cases. We comment on this more after presenting the model.

Infinite horizon, average cost case

We first consider the infinite horizon case where agents optimize their average expected loss. As in the routing game, we allow the losses to be functions of the population mass distribution. We assume the same form of the action sets and probabilistic transitions introduced in Section 4.1. Let x_j be the steady state population mass in node j and let x_j^a be the portion of that population that chooses action a. We have that

$$x_j = \sum_{a \in \mathcal{A}_j} x_j^a \tag{4.17}$$

We will denote the vector of all masses on nodes as $\mathbf{x} \in \mathbb{R}_+^{|\mathcal{N}|}$ where $\mathbf{x}_j = x_j$. We will denote the vector of masses choosing each action at each node as $x \in \mathbb{R}_+^{|\mathcal{A}|}$. Note that $x_a = x_j^a$ with the implication that action a is only available from node j. We note that we can take advantage of the structure of $\mathbf{P}_\mathcal{A}$ to write Equation (4.17) in matrix form as

$$\mathbf{x} = I_o \mathbf{P}_\mathcal{A} x \tag{4.18}$$

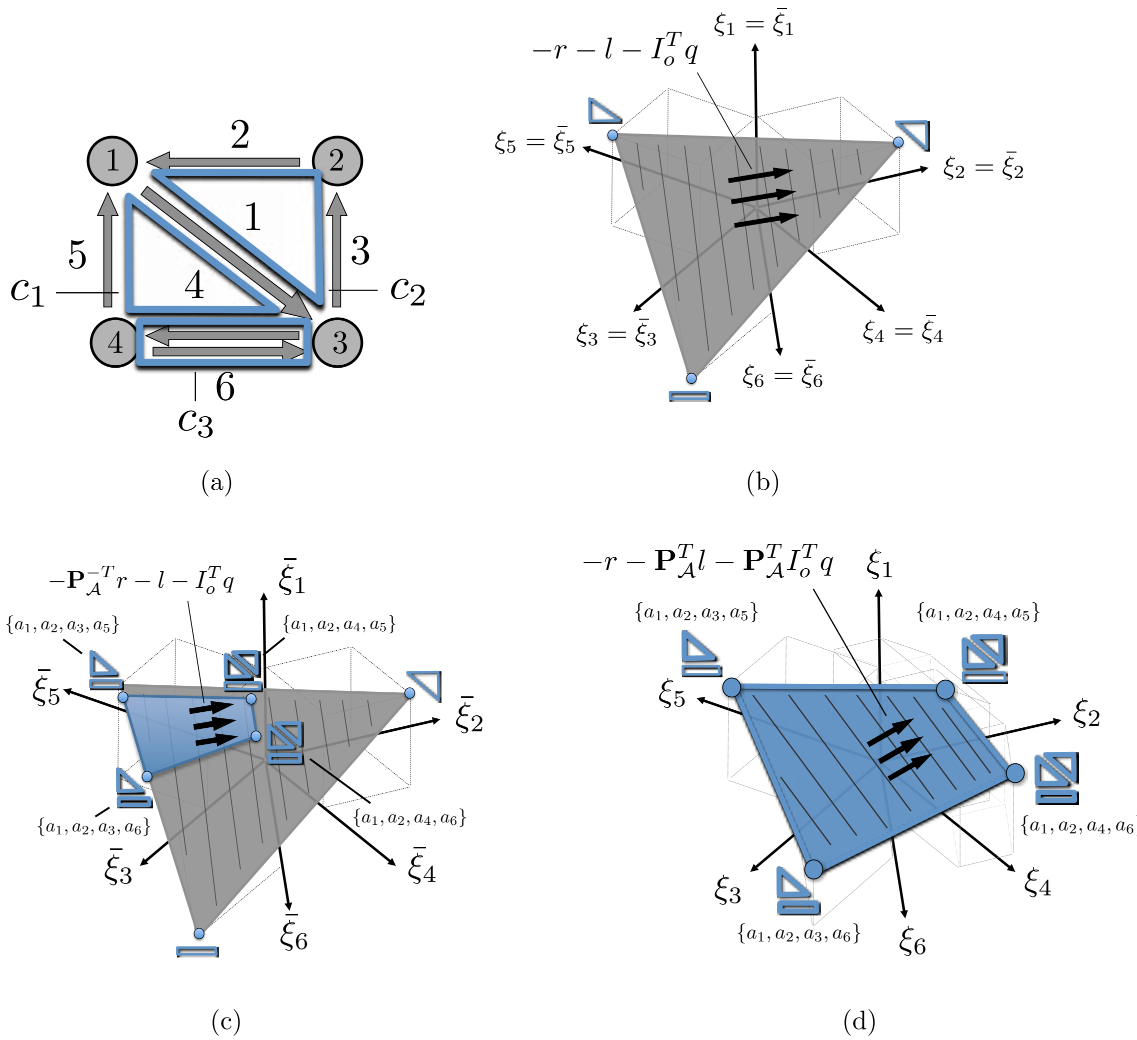

Figure 4.1: (a) Sample MDP graph. (b) Illustration of fully deterministic LP (edge space = action space). (c) Illustration of stochastic LP (edge space) (d) Illustration of stochastic LP (action space).

We will denote the mass transitioning from j to i as x_{ij} which we can compute from the action masses as

$$x_{ij} = \sum_{a \in \mathcal{A}_j} P_{ij}^a x_j^a \tag{4.19}$$

We will refer to vector of all transition or edge masses as $\bar{x} = \mathbb{R}_+^{|\mathcal{E}|}$. Note that $\bar{x}_e = x_{ij}$ for the edge e that runs from j to i. In matrix form, we can write Equation (4.19) as

$$\bar{x} = \mathbf{P}_{\mathcal{A}} x \tag{4.20}$$

At steady state, mass conservation at each node gives

$$\sum_{a \in \mathcal{A}_i} x_i^a = x_i = \sum_j x_{ij} = \sum_j \sum_{a \in \mathcal{A}_j} P_{ij}^a x_j^a \tag{4.21}$$

In addition, the total mass of the population, m, is conserved and must be positive everywhere.

$$m = \sum_j \sum_{a \in \mathcal{A}_j} x_j^a, \qquad x_j^a \geq 0, \qquad\qquad \forall a \in \mathcal{A}_j, \quad \forall j \in \mathcal{N}, \tag{4.22}$$

In matrix form, (4.21) and (4.22) can be written as

$$I_o \mathbf{P}_{\mathcal{A}} x - I_i \mathbf{P}_{\mathcal{A}} x = G \mathbf{P}_{\mathcal{A}} x = 0, \qquad \mathbf{1}^T x = m, \qquad x \geq 0 \tag{4.23}$$

Analogous to the routing game, we now allow the losses to depend on the population masses. In general, we assume the action losses are functions of the action masses, $r(x) : \mathbb{R}_+^{|\mathcal{A}|} \to \mathbb{R}^{|\mathcal{A}|}$; the transition losses are functions of the edges masses $l(\bar{x}) : \mathbb{R}_+^{|\mathcal{E}|} \to \mathbb{R}^{|\mathcal{E}|}$; and the node losses are functions of the node masses, $q(\mathbf{x}) : \mathbb{R}_+^{|\mathcal{N}|} \to \mathbb{R}^{|\mathcal{N}|}$. We will often abuse notation and write the edge and node losses as functions of the action masses, $l(x)$ and $q(x)$, with the understanding that $l(x) = l(\mathbf{P}_{\mathcal{A}} x)$ and $q(x) = q(I_o \mathbf{P}_{\mathcal{A}} x)$. Again, depending on context, we will sometimes index $r(x)$ by nodes and actions, i.e. $r_j^a(x)$, and sometimes simply by actions, i.e. $r_a(x)$. We will also sometimes index $l(\bar{x})$ by the starting and ending nodes, i.e. $l_{ij}(\bar{x})$, and sometimes by edges, i.e. $l_e(\bar{x})$.

As in Section 4.1, let η represent the policy of an individual agent within the population and $P(\eta)$ the resulting transition matrix. Intuitively, we are looking for the mass distribution x that is consistent with the individual policies that agents can choose similar to the mass distribution in the classic routing game being consistent with the shortest path problem each individual population member solves.

Given the action mass distribution of the other population members x, we take the cost of the feedback policy η to be the average expected loss over the infinite horizon. Given

Assumption 1 (irreduciblibility and aperiodicity), this can be calculated as

$$\begin{aligned}\ell(x,\eta) &= \sum_j \Big(\sum_i l_{ij}(x)P_{ij}(\eta) + \sum_{a\in\mathcal{A}_j} r_j^a(x)\eta_j^a + q_j(x)\Big)p_j(\eta)\\ &= \sum_j \sum_{a\in\mathcal{A}_j}\Big(\sum_i l_{ij}(x)P_{ij}^a + r_j^a(x) + q_j(x)\Big)\eta_j^a p_j(\eta)\end{aligned} \tag{4.24}$$

We will use the tuple $\mathbb{G} = (\mathcal{G}, r(x), l(x), q(x), \mathcal{A}, \mathbf{P}_{\mathcal{A}})$ to refer to the game induced by the graph structure, the set of loss functions, actions, and transition probabilities when individual agents within the population seek to optimize Equation (4.24).

We can now define the appropriate version of a Wardrop equilibrium for the stationary infinite horizon case as follows.

Definition 6 *A population mass distribution x is a* stationary infinite horizon Wardrop equilibrium *for the MDP routing game $\mathbb{G}$ if for any two policies η, η' such that $\eta_j^a > 0$ only if $x_j^a > 0$ for all j.*

$$\ell(x,\eta) \le \ell(x,\eta') \tag{4.25}$$

Intuitively, no population member can improve their expected average reward by changing their action profile at any state.

We now define the notion of a potential game in the infinite horizon case.

Definition 7 *We say $\mathbb{G}$ is a* infinite-horizon potential game *if there exists a C^1 function $F : x \mapsto \mathbf{R}$ such that*

$$\frac{\partial F}{\partial x_j^a}(x) = \sum_i l_{ij}(x)P_{ij}^a + r_j^a(x) + q_j(x) \tag{4.26}$$

Intuitively, the derivative of the potential function with respect to the mass taking a particular action captures the immediate payoff of that action.

Remark 10 *We note that if we can find a function $F(\cdot)$ that is also explicitly a function of the edge and node masses as well as the action masses, and $F(\cdot)$ satisfies*

$$\frac{\partial F}{\partial x_j^a} = r_j^a(x), \qquad \frac{\partial F}{\partial x_{ij}} = l_{ij}(x), \qquad \frac{\partial F}{\partial x_j} = q_j(x) \tag{4.27}$$

then Condition (4.26) *is satisfied by applying the chain rule and Equations* (4.17) *and* (4.19).

Remark 11 *This definition of a potential function is a substantial deviation from mean-field games on graphs where the potential function differentiation condition is defined with respect to the mass on the nodes as opposed to the mass taking a particular action. See [83] for comparison in the stationary case.*

Remark 12 *In the special case where each $r_j^a(\cdot)$ is simply a function of x_j^a and $l_{ij}(\cdot)$ is simply a function of x_{ij}, and $q_j(\cdot)$ is simply a function of x_j, we can use the potential*

$$F(x) = \sum_j \left[\sum_i \int_0^{x_{ij}} l_{ij}(u)\ du + \sum_{a\in\mathcal{A}_j} \int_0^{x_j^a} r_j^a(u)\ du + \int_0^{x_j} q_j(u)\ du \right] \tag{4.28}$$

which bears obvious similarities to the standard routing game potential function.

We now show that we can find a Wardrop equilibrium by minimizing the potential function.

Theorem 5 *Given a potential function F for the infinite horizon game $\mathbb{G}$, if x satisfies the KKT first order necessary conditions for minimizing F, then x is an infinite horizon Wardrop equilibrium.*

Proof 7 *The optimization problem and corresponding Lagrangian are given by*

$$\min_{x\geq 0} \quad F(x) \tag{4.29a}$$

$$s.t. \quad \sum_{a\in\mathcal{A}_i} x_i^a = \sum_j \sum_{a\in\mathcal{A}_j} P_{ij}^a x_j^a \tag{4.29b}$$

$$m = \sum_j \sum_{a\in\mathcal{A}_j} x_j^a \tag{4.29c}$$

$$\mathcal{L}(x,\pi,\lambda,\mu) = F(x) - \sum_i \pi_i \left(\sum_{a\in\mathcal{A}_i} x_i^a - \sum_j \sum_{a\in\mathcal{A}_j} P_{ij}^a x_j^a \right) - \lambda \left(\sum_j \sum_{a\in\mathcal{A}_j} x_j^a - m \right) - \sum_{j,a} \mu_j^a x_j^a \tag{4.30}$$

with Lagrange multipliers $\pi \in \mathbb{R}^{|\mathcal{N}|}$, $\lambda \in \mathbb{R}$, and $\mu_j \in \mathbb{R}_+^{|\mathcal{A}_j|}$. Since F is a potential function, the first order necessary conditions give

$$\frac{\partial \mathcal{L}}{\partial x_j^a} : \qquad \sum_i l_{ij}(x) P_{ij}^a + r_j^a(x) + q_j(x) - \pi_j + \sum_i \pi_i P_{ij}^a - \lambda - \mu_j^a = 0 \tag{4.31a}$$

$$\mu_j^a \geq 0 \tag{4.31b}$$

$$\mu_j^a x_j^a = 0 \tag{4.31c}$$

Using (4.31a), *we can compute the utility of any individual's strategy profile* η *as*

$$\ell(x,\eta) = \sum_j \sum_{a\in\mathcal{A}_j} \left(\sum_i l_{ij}(x) P_{ij}^a + r_j^a(x) + q_j(x) \right) \eta_j^a p_j(\eta) \tag{4.32a}$$

$$= \sum_j \sum_{a\in\mathcal{A}_j} \left(\pi_j - \sum_i \pi_i P_{ij}^a + \lambda + \mu_j^a \right) \eta_j^a p_j(\eta) \tag{4.32b}$$

$$= \sum_j \pi_j p_j(\eta) - \sum_i \pi_i \sum_j \sum_{a\in\mathcal{A}_j} \eta_j^a P_{ij}^a p_j(\eta) + \lambda + \sum_j \sum_{a\in\mathcal{A}_j} \eta_j^a \mu_j^a p_j(\eta) \tag{4.32c}$$

$$= \sum_j \pi_j p_j(\eta) - \sum_i \pi_i p_i(\eta) + \lambda + \sum_j \sum_{a\in\mathcal{A}_j} \eta_j^a \mu_j^a p_j(\eta) \tag{4.32d}$$

$$= \lambda + \sum_j \sum_{a\in\mathcal{A}_j} \eta_j^a \mu_j^a p_j(\eta) \tag{4.32e}$$

It follows that

$$\ell(x,\eta) - \sum_j \sum_{a\in\mathcal{A}_j} \eta_j^a \mu_j^a p_j(\eta) = \ell(x,\eta') - \sum_j \sum_{a\in\mathcal{A}_j} {\eta'}_j^a \mu_j^a p_j(\eta') \tag{4.33}$$

for any two action profiles η, η'. *If* $\eta_j^a > 0$ *only if* $x_j^a > 0$ *for all* j, *than by complementary slackness we have that* $\eta_j^a \mu_j^a = 0$. *It follows that*

$$\ell(x,\eta) \leq \ell(x,\eta') \tag{4.34}$$

since ${\eta'}_j^a \mu_j^a p_j(\eta') \geq 0$ *for all* j *and* $a \in \mathcal{A}_j$.

Remark 13 *Many interesting cases break the irreducible, aperiodic assumption (Assumption 1), the deterministic transition case being the most obvious. We note that the steady state equilibrium concept may be valid even in these cases. A reducible Markov chain will have several irreducible subsets of recurrent states. Problem* (4.4) *will assign a certain amount of mass to each of these irreducible subsets in order to minimize the expected loss. If the initial probability distribution of the agent is fixed, there may not be a policy that divides up the probability mass between the irreducible subsets in this optimal way. If, however, we think of the agent as choosing their initial probability distribution as well as a policy, they would be able to assign the appropriate amount of mass to each recurrent subset. In the case of a periodic Markov chain assuming strictly increasing losses, we note that oscillating solutions cause some members of the population to experience more congestion than others. If agents have the option of remaining in a node at any time, they will damp out these oscillations. We could recover the full aperiodic assumption by adding a small probability of remaining in the same node to each action; however, from this argument it seems to be enough to simply assume that agents have the option of remaining in a node. This argument is particularly applicable in the deterministic case where we are really thinking of agents as playing a routing game where they choose between cycles as opposed to routes.*

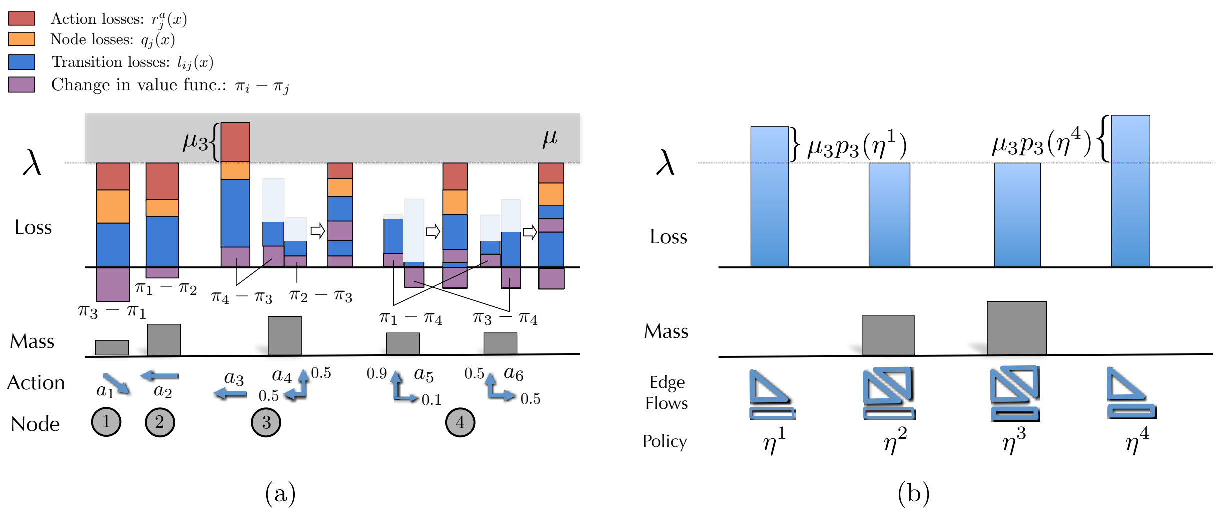

Figure 4.2: Infinite horizon equilibrium conditions. (a) Action balance condition with interpretation of the Lagrange multipliers (illustrating Equation (4.31a)). (b) Pure strategy policy balance condition (analogous to Figure 2.1b).

In Figure 4.2, we give two illustrations of the equilibrium condition in the infinite horizon case for the graph and equilibrium depicted in Figure 4.3 (see below). Specifically, it depicts how the Lagrange multipliers of the optimization problem capture the loss information at equilibrium. The multiplier λ encodes the average cost of taking each action, π can be thought of as value function on the nodes that encodes how any individual transition between nodes deviates from the average cost, and μ encodes the inefficiency of any particular action. The fact that $\mu_j^a = 0$ if and only if $x_j^a > 0$ means that no member of the population takes an inefficient action at any node. Figure 4.2a shows the first order optimality condition (Equation (4.31a)) for each action specifically detailing these Lagrange multipliers. Figure 4.2b shows the average cost of each pure strategy policy that members of the population can choose. In the standard routing game, the equilibrium condition is that each route that some portion of the population chooses has equal minimal latency. In the infinite horizon MDP routing game, the equilibrium condition is that each pure strategy policy that members of the population choose has equal minimal latency. This is illustrated in Figure 4.2b.

Parallels with classic routing games

The infinite horizon formulation just presented has clear parallels with the edge formulation of the classical routing game. Assuming the appropriate ordering of edges and actions, we can write Problem 4.29 in matrix form as

$$\min_{\mathrm{x}} \quad F(x) \tag{4.35a}$$
$$\text{s.t.} \quad \mathbf{1}^T x = m, \quad x \geq 0 \tag{4.35b}$$
$$G\mathbf{P}_{\mathcal{A}} x = 0 \tag{4.35c}$$

In the fully deterministic case where $\mathbf{P}_{\mathcal{A}} = I_{|\mathcal{E}|\times|\mathcal{E}|}$ and $\bar{x} = x$, this becomes

$$\min_{x} \quad F(x) \tag{4.36a}$$
$$\text{s.t.} \quad \mathbf{1}^T x = m, \quad x \geq 0 \tag{4.36b}$$
$$Gx = 0 \tag{4.36c}$$

In the deterministic case, there are parallels to the path formulation of the routing game as well. Analogous to the discussion in Section 4.1, given the set of cycles of the graph $\mathcal{C}$ and the scaled indicator matrix for the cycles $\mathbf{E}_{\mathcal{C}}$ defined in Equation (4.12), we can rewrite Problem (4.36) as

$$\min_{z} \quad F(x) \tag{4.37a}$$
$$\text{s.t.} \quad \mathbf{1}^T z = m, \quad z \geq 0 \tag{4.37b}$$
$$x = \mathbf{E}_{\mathcal{C}} z \tag{4.37c}$$

where $z \in \mathbb{R}_+^{|\mathcal{C}|}$ is a vector of the masses on each cycle.

The feasible set and equilibrium condition for the deterministic and stochastic cases are illustrated in Figure 4.3 in both the edge and action mass spaces.

Finite horizon, total cost case

In the finite horizon case, we consider a time horizon of T discrete steps. Let $\mathcal{T} = \{0, 1, \ldots, T\}$ refer to the set of times. Let $\mathcal{A}_j^t$ refer to the set of actions from node j at time step t. We will use $\mathcal{A} = \cup_{t\in\mathcal{T}} \cup_{j\in\mathcal{N}} \mathcal{A}_j^a$ to refer to the set of all actions over all time steps again with the assumption that $a \in \mathcal{A}$ is only available from a specific node at a specific time step. Let $P_j^{at} = (P_{ij}^{at})_{i\in\mathcal{N}_j} \in \Delta_{|\mathcal{N}_j|}$ be the transition probabilities associated with action a. P_{ij}^{at} is the probability that an agent who chooses action a in node j at time t will transition to node i at time $t+1$. Let $\mathcal{P}_{\mathcal{A}} = \left\{P_j^{at}\right\}_{a\in\mathcal{A}}$ be the set of all transition probabilities. Let x_j^t refer to the mass in node j at time t and x_j^{at} the portion of that mass taking action a. We have that

$$x_j^t = \sum_{a\in\mathcal{A}_j^t} x_j^{at} \tag{4.38}$$

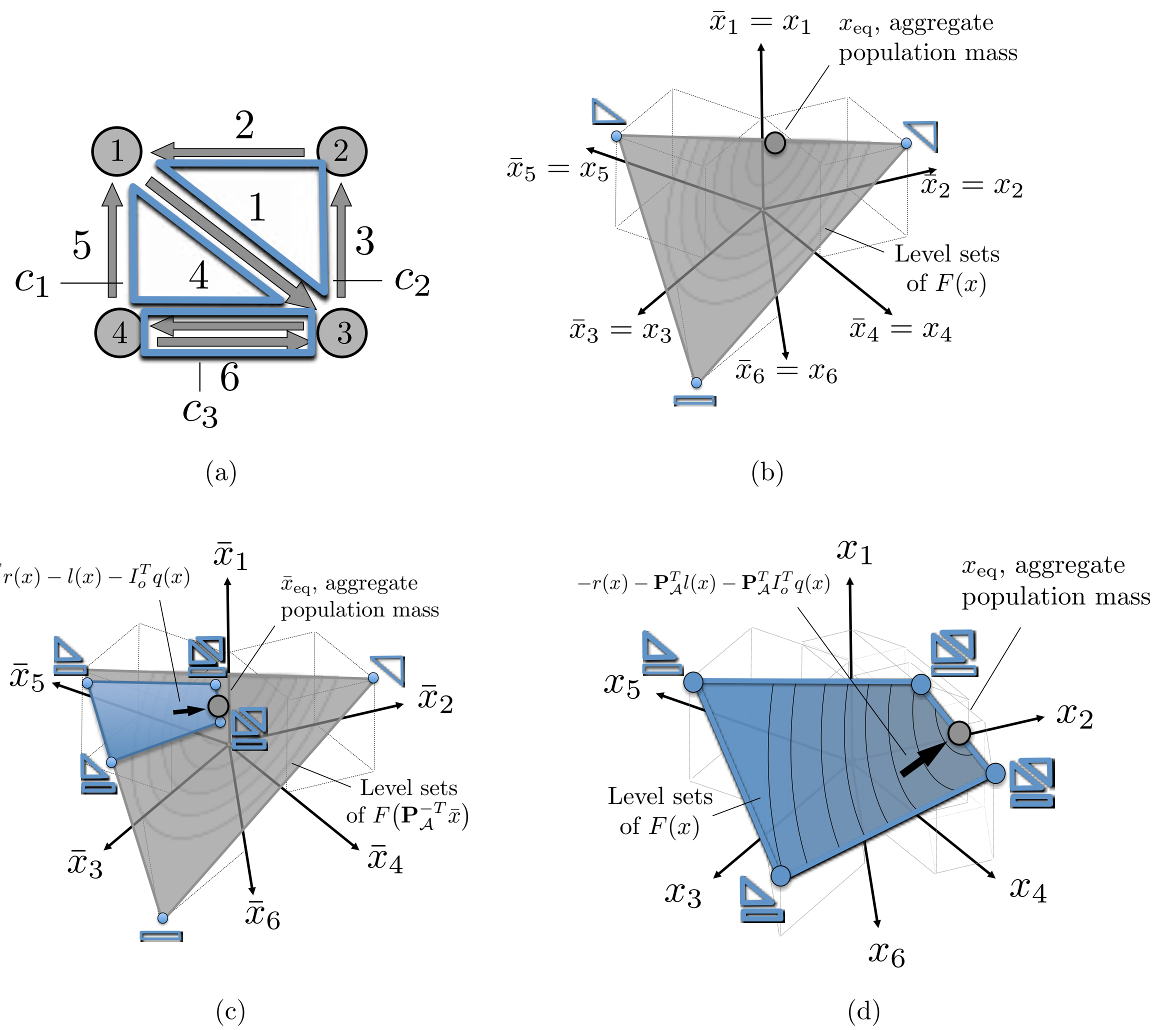

Figure 4.3: (a) Graph and cycles. Illustrations of minimizing the potential function for the MDP routing game in the (b) fully deterministic case (edge flows = action flows), (c) the stochastic transition case (edge flows), and (d) the stochastic transition case (action flows).

We will denote the mass transitioning from j to i at time t as x_{ij}^t. The edge masses are related to the action masses by the equation

$$x_{ij}^t = \sum_{a\in\mathcal{A}_j^t} P_{ij}^{at} x_j^{at} \tag{4.39}$$

We will use $x \in \mathbb{R}_+^{|\mathcal{A}|}$ to refer to the vector of all action masses at all time steps, $\bar{x} \in \mathbb{R}_+^{|\mathcal{E}|\cdot|\mathcal{T}|}$ to refer to the vector of all edge masses at all time steps, and $\mathbf{x} \in \mathbb{R}_+^{|\mathcal{N}|\cdot|\mathcal{T}|}$ to refer to the vector of all node masses at all time steps.

Summing over all nodes j gives us the new population in node i

$$\sum_{a\in\mathcal{A}_i^{t+1}} x_i^{a(t+1)} = x_i^{t+1} = \sum_j x_{ij}^t = \sum_j \sum_{a\in\mathcal{A}_j^t} P_{ij}^{at} x_j^{at} \tag{4.40}$$

The initial distribution of populations at the first time step is given a priori.

$$\sum_{a\in\mathcal{A}_i^0} x_i^{a0} = m_i \tag{4.41}$$

We now allow the loss functions to change with time as well. Let $r_j^{at}(x)$ refer to the loss of taking action a at node j at time t, let $l_{ij}^t(\bar{x})$ refer to the loss of transitioning from j to i at time t, and let $q_j^t(\mathbf{x})$ refer to the loss of being in node j at time t. Again in general, $q_j^t(\cdot)$ and $l_{ij}^t(\cdot)$ will be functions of the node masses and edge masses respectively but as in the infinite horizon case we will abuse notation and write them as functions of the action masses implicitly assuming Equations (4.39), (4.40), and (4.41). We will use $r(x)$, $l(x)$, and $q(x)$ to refer to vectors of all the loss functions.

Each individual agent in the population seeks to minimize their expected total loss over the entire time horizon. At each transition, an agent must consider the immediate reward they expect from the action as well as the expected cost-to-go from the state they transition to.

For a given population distribution x, let $v_j^t(x)$ be the expected optimal cost-to-go from node j at time t which can be defined backwards recursively from $t = T$ as

$$v_j^T(x) = \min_{a\in\mathcal{A}_j^T} \left(\sum_i P_{ij}^{aT} l_{ij}^T(x) + r_j^{aT}(x) + q_j^T(x) \right) \tag{4.42a}$$

$$v_j^{t-1}(x) = \min_{a\in\mathcal{A}_j^{t-1}} \left(\sum_i P_{ij}^{a(t-1)} \left[l_{ij}^{t-1}(x) + v_i^t(x) \right] + r_j^{a(t-1)}(x) + q_j^{t-1}(x) \right) \tag{4.42b}$$

Note that $v_j^t(x)$ is a function of x. We define the cost of taking action a from node j at time t as

$$\ell_j^{aT}(x) = \sum_i P_{ij}^{aT} l_{ij}^T(x) + r_j^{aT}(x) + q_j^T(x) \tag{4.43a}$$

$$\ell_j^{at}(x) = \sum_i P_{ij}^{at} \left[l_{ij}^t(x) + v_i^{t+1}(x) \right] + r_j^{at}(x) + q_j^t(x) \qquad t < T \tag{4.43b}$$

We will refer to the game induced by agents seeking to optimize their expected cost over the time horizon by the tuple $\mathbb{G} = (\mathcal{G}, r(x), l(x), q(x), \mathcal{A}, \mathcal{P}_{\mathcal{A}})$.

We now give the appropriate Wardrop equilibrium condition and the finite horizon potential game.

Definition 8 *We say that x is a finite-horizon* Wardrop equilibrium *for the game $\mathbb{G}$ if at every node j at every time step t for any two actions $a, a' \in \mathcal{A}_j^t$ such that $x_j^{at} > 0$*

$$\ell_j^{at}(x) \leq \ell_j^{a't}(x) \tag{4.44}$$

This is the standard definition of Wardrop equilibria applied to the appropriate decision that agents make at each node and at each time step. Intuitively, no agent has an incentive to deviate from their chosen strategy at each transition. We now define the notion of a potential function for the finite horizon game.

Definition 9 *We call $\mathbb{G}$ a finite horizon* potential game *if the following condition holds*

There exists a C^1 function $F : x \mapsto \mathbb{R}$ such that

$$\frac{\partial F}{\partial x_j^{at}}(x) = \sum_i l_{ij}^t(x) P_{ij}^{at} + r_j^{at}(x) + q_j^t(x) \tag{4.45}$$

Remark 14 *Again, we note that if we can find a function $F(\cdot)$ that is also explicitly a function of the edge and node masses as well as the action masses, and $F(\cdot)$ satisfies*

$$\frac{\partial F}{\partial x_j^{at}} = r_j^{at}(x), \qquad \frac{\partial F}{\partial x_{ij}^t} = l_{ij}^t(x), \qquad \frac{\partial F}{\partial x_j^t} = q_j^t(x) \tag{4.46a}$$

then Condition (4.45) *is satisfied by applying the chain rule using Equations* (4.38) *and* (4.39).

Remark 15 *Again, this definition of a potential function deviates from mean-field games on graphs where the potential function differentiation condition is defined with respect to the mass on the nodes as opposed to the mass taking particular actions. See [84] for comparison in the finite horizon case.*

Remark 16 *In the special case where each $l_{ij}^t(\cdot)$ is simply a function of x_{ij}^t, each $r_j^{at}(\cdot)$ is simply a function of x_j^{at}, and each $q_j^t(\cdot)$ is simply a function of x_j^t, we can use the potential*

$$F(x) = \sum_t \left(\sum_j \left[\sum_i \int_0^{x_{ij}^t} l_{ij}^t(u)\ du + \sum_{a \in \mathcal{A}_j^t} \int_0^{x_j^{at}} r_j^{at}(u)\ du. + \int_0^{x_j^t} q_j^t(u)\ du \right] \right) \tag{4.47}$$

which could be thought of as a discrete time functional version of the standard routing game potential.

We now show that the first order necessary conditions for minimizing the potential with respect to the mass conservation constraints guarantees that the mass distribution is a finite horizon Wardrop equilibrium.

Theorem 6 *Given a potential function $F(x)$ for the finite horizon game $\mathbb{G}$. x satisfies the KKT first order necessary conditions for minimizing F if and only if x is a finite horizon-Wardrop equilibrium for $\mathbb{G}$.*

Proof 8 *($\Rightarrow$) The optimization problem and corresponding Lagrangian are given by*

$$\begin{aligned} \min_{x\geq 0} \quad & F(x) && (4.48a)\\ s.t. \quad & \sum_{a\in\mathcal{A}_i^0} x_i^{a0} = m_i \qquad \forall i && (4.48b)\\ & \sum_{a\in\mathcal{A}_i^{t+1}} x_i^{a(t+1)} = \sum_j \sum_{a\in\mathcal{A}_j^t} P_{ij}^{at} x_j^{at} \qquad \forall i && (4.48c) \end{aligned}$$

$$\begin{aligned} L(x,\mu,\pi) = F(x) - \sum_i \pi_i^0 \Big(\sum_{a\in\mathcal{A}_i^0} x_i^{a0} - m_i\Big) - \\ \sum_{t=0}^{T-1}\sum_i \pi_i^{t+1}\Big(\sum_{a\in\mathcal{A}_i^{t+1}} x_i^{a(t+1)} - \sum_j\sum_{a\in\mathcal{A}_j^t} P_{ij}^{at}x_j^{at}\Big) - \sum_t\sum_j\sum_{a\in\mathcal{A}_j^t}\mu_j^{at}x_j^{at} \end{aligned} \tag{4.49}$$

with Lagrange multipliers $\pi \in \mathbb{R}^{|\mathcal{N}|\cdot|\mathcal{T}|}$ and $\mu \in \mathbb{R}_+^{|\mathcal{A}|\cdot|\mathcal{T}|}$ Computing the KKT first-order necessary conditions gives

$$\begin{aligned} \sum_i P_{ij}^{at} l_{ij}^t(x) + r_j^{at}(x) + q_j^t(x) - \pi_j^t + \sum_i \pi_i^{t+1} P_{ij}^{at} - \mu_j^{at} &= 0 && (4.50a)\\ \sum_i P_{ij}^{aT} l_{ij}^T(x) + r_j^{aT}(x) + q_j^T(x) - \pi_j^T - \mu_j^{aT} &= 0 && (4.50b)\\ \mu_j^{at} &\geq 0 && (4.50c)\\ \mu_j^{at} x_j^{at} &= 0 && (4.50d) \end{aligned}$$

In the following, we suppress the dependence of $l_{ij}^t(x)$, $r_j^{at}(x)$, $q_j^t(x)$, $u_j^{at}(x)$ and $v_j^t(x)$ on x for notational simplicity.
Starting at $t = T$, we have that

$$\begin{aligned} \pi_j^T &= \sum_i P_{ij}^{aT} l_{ij}^T + r_j^{aT} + q_j^T - \mu_j^{aT}\\ &= \ell_j^{aT} - \mu_j^{aT} \end{aligned} \tag{4.51}$$

For any two actions $a, a' \in \mathcal{A}_j^T$ such that $x_j^{aT} > 0$, we have that $\mu_j^{aT} = 0$ and $\mu_j^{a'T} \geq 0$ and thus

$$\pi_j^T = \ell_j^{aT} \leq \ell_j^{a'T} \tag{4.52}$$

Thus, we have that Condition (4.44) is satisfied at $t = T$. We also have that

$$\pi_j^T \leq \min_{a \in \mathcal{A}_j^T} \ell_j^{aT} = v_j^T \tag{4.53}$$

We would have equality except for the possibility that $x_j^T = 0$ and thus $x_j^{aT} = 0$ for all $a \in \mathcal{A}_j^T$. In this case, $\mu_j^{aT} > 0$ for all a and π_j^T could be shifted down by an arbitrary amount. In the case where $x_j^T > 0$, there must exist $a_1 \in \mathcal{A}_j^T$ such that $x_j^{a_1 T} > 0$. It follows that $\mu_j^{a_1 T} = 0$ and

$$\pi_j^T = \ell_j^{a_1 T} = \min_{a \in \mathcal{A}_j^T} \ell_j^{aT} = v_j^T. \tag{4.54}$$

Thus, we have that π_j^T is a lower bound on the optimal cost-to-go from node j at time T with equality achieved whenever $x_j^T > 0$.
Moving on to $t = T - 1$, we have that a

$$\pi_j^{T-1} = \sum_i P_{ij}^{a(T-1)} \left[l_{ij}^{T-1} + \pi_i^T \right] + r_j^{a(T-1)} + q_j^{a(T-1)} - \mu_j^{at} \tag{4.55}$$

$$\leq \sum_i P_{ij}^{a(T-1)} \left[l_{ij}^{T-1} + v_i^T \right] + r_j^{a(T-1)} + q_j^{a(T-1)} - \mu_j^{at} \tag{4.56}$$

$$\leq \ell_j^{a(T-1)} - \mu_j^{a(T-1)} \tag{4.57}$$

Since $\mu_j^{at} \geq 0$, we have that

$$\pi_j^{T-1} \leq \ell_j^{a(T-1)} \tag{4.58}$$

for all $a \in \mathcal{A}_j^{T-1}$ If $x_j^{T-1} > 0$, for any two actions $a_1, a_2 \in \mathcal{A}_j^{T-1}$ such that $x_j^{a_1(T-1)} > 0$, we have that $\mu_j^{a_1(T-1)} = 0$ and $\mu_j^{a_2(T-1)} \geq 0$. For any i such that $P_{ij}^{a_1(T-1)} > 0$, $x_i^T > 0$ and thus $\pi_i^T = v_i^T$ by (4.54) and it follows that

$$\pi_j^{T-1} = \ell_j^{a_1(T-1)} \leq \ell_j^{a_2(T-1)} \tag{4.59}$$

which is Condition (4.44) at time $t = T - 1$. In addition, we have that

$$\pi_j^{T-1} \leq \min_{a \in \mathcal{A}_j^{T-1}} \ell_j^{a(T-1)} = v_j^{T-1} \tag{4.60}$$

with equality achieved whenever $x_j^{T-1} > 0$. The result follows by induction.
($\Leftarrow$) *Conversely suppose x satisfies Condition (4.44). The primal feasibility conditions are*

satisfied a priori. We need to construct dual variables π_j^t and μ_j^{at} that satisfy dual feasibility, complementary slackness, and the gradient condition. Starting from $t = T$ for each j, define

$$\begin{aligned}\pi_j^T &= \min_{a \in \mathcal{A}_j^T} \quad \sum_i P_{ij}^{aT} l_{ij}^T + r_j^{aT} + q_j^{aT} \\ &= \min_{a \in \mathcal{A}_j^T} \quad \ell_j^{aT} = v_j^T \end{aligned} \tag{4.61}$$

By Condition (4.44)*, we have that $\pi_j^T = \ell_j^{a_1 T}$ for all $a_1 \in \mathcal{A}_j^T$ such that $x_j^{a_1 T} > 0$ and $\pi_j^T \leq \ell_j^{a_2 T}$ for all $a_2 \in \mathcal{A}_j^T$ such that $x_j^{a_2 T} = 0$. Defining*

$$\mu_j^{aT} = \ell_j^{aT} - \pi_j^T \tag{4.62}$$

satisfies dual feasibility, complementary slackness, and the gradient constraint for $t = T$. Moving on to $t = T - 1$, since π_j^T is the optimal cost-to-go from j at time T, we have that

$$\ell_j^{a(T-1)} = \sum_i P_{ij}^{a(T-1)} \left[l_{ij}^{T-1} + \pi_i^T \right] + r_j^{a(T-1)} + q_j^{a(T-1)} \tag{4.63}$$

Let

$$\pi_j^{T-1} = \min_{a \in \mathcal{A}_j^{T-1}} \ell_j^{a(T-1)} = v_j^{T-1} \tag{4.64}$$

By Condition (4.44)*, $\pi_j^{T-1} = \ell_j^{a_1(T-1)}$ for all $a_1 \in \mathcal{A}_j^{T-1}$ such that $x_j^{a_1(T-1)} > 0$ and $\pi_j^{T-1} \leq \ell_j^{a_2(T-1)}$ for all $a_2 \in \mathcal{A}_j^{T-1}$ such that $x_j^{a_2(T-1)} = 0$. Again, setting*

$$\mu_j^{a(T-1)} = \ell_j^{a(T-1)} - \pi_j^{T-1} \tag{4.65}$$

satisfies dual feasibility and complementary slackness and by substituting Equation (4.63) *into Equation* (4.65) *gives the gradient condition. The result follows by induction.*

Remark 17 *In the finite horizon case if the transition matrices are fully deterministic, the problem could be framed as a classic routing game by making T copies of the state space, connecting the proper nodes (where transitions are allowed) between each time step, and then enumerating all possible paths through this new network. It should be noted that this is not possible however when agents strategies consist of choosing a sequence of non-deterministic transitions. It might be the case for a specific sequence of transitions that a realization of the first t transitions make the $t + 1$ transition impossible even if it would have been possible for another realization. Indeed, this might be true for* all *possible sequences of transitions available to an agent.*

Existence and uniqueness

We comment briefly on existence and uniqueness of the equilibria in both the infinite and finite horizon cases. Existence is guaranteed in both cases by Weierstrass's Theorem (a continuous function attains it's minimum on a nonempty, closed and bounded set). This is simpler even then the classic routing game case (see Theorem 2.4 in [15]) since the feasible set is always closed and bounded. Uniqueness is guaranteed in each case if the potential function is a strictly convex function of the action masses. Assuming a potential function of the form of Equation 4.28 or 4.47, this is guaranteed if all the action losses are strictly increasing functions. If not, strict convexity still may be achieved by either the edge losses or node losses being strictly increasing. In these case, however, it will also depend on the properties of the matrices $\mathbf{P}_{\mathcal{A}}$ and $I_o\mathbf{P}_{\mathcal{A}}$. For example, in a game where all the edge losses are strictly increasing, the equilibrium still may not be unique if $\mathbf{P}_{\mathcal{A}}$ has a nontrivial nullspace. This could happen, for instance, if there were more actions available at one node than there were edges coming out of that node. Similar arguments apply to the node losses and the matrix $I_o\mathbf{P}_{\mathcal{A}}$.

Multiple populations

As in the traditional routing game, these results extend directly to multiple populations. Each new population would get its own state vector and copy of the constraints and could either share the same action set or have their own action set to choose from. Two populations having different action sets is analogous to two populations in a traditional routing game having different origin-destinations pairs and different routing matrices. (In the finite horizon case, we already implicitly consider different populations that start at each of the nodes in the network.)

If each population cares about the same loss functions then the loss functions must depend on the sum of the population masses. The results follow by considering the potential function differentiation condition with respect to each population mass separately and applying the chain rule. It would also be possible to consider a game where populations care about some of the same losses but not all of them. For example, the results would still go through in a game where each population cares about the same edge and node losses but separate action losses. The edge and node losses, in this case, would need to depend on the sum of the population masses and the action losses for each population would only depend on the mass of that population. Indeed any combination of populations caring or not caring about particular types of losses would work as along as when populations care about the same loss, that loss is dependent on the sum of the population masses.

Variable demand

We comment briefly that we could also easily formulate a variable demand MDP routing game similarly to the classical variable demand routing game. In the infinite horizon case, we would define a demand function for the entire mass using the network. In the finite horizon case as formulated above, we would define separate demand functions for the mass at each node at the initial time step. We could also extend the model so that mass is able to join the game at any time step and define demand functions for at each node at each time step. It should be noted that in this case, we are making the assumption that all agents participating in the game have full information about these demand functions, i.e. they know how much population mass will join the game at each node at each time step depending on the cost-to-go from that node.

Further discussion of modeling considerations

As previously mentioned, this formulation can be thought of as a routing game perspective on stochastic population games. Depending on the situation, it could provide a simple, more tractable modeling framework than the traditional mean-field game approach. Given the comparable form, it could be directly combined with traditional routing game models. One application of this would be improving the queue-routing model for parking presented in Chapter 3. In Section 4.4, we presented a more complex (realistic) model of cars competing for parking spaces around a set of block faces. This model could be combined with the queue-routing game replacing the assumption that cars spread out uniformly over the block faces in a parking area.

One of the weaknesses of the model is that the discrete time view of time is inconsistent with congestion effects that increase latency at any given node or on any given edge of the graph. In the ridesharing examples, for instance, all the mass that makes a specific transition at a given time step completes that transition by the next time step even though the cost of making a specific transition is related to the waiting time to pick up riders. This could be problematic depending on the required accuracy of the model.

4.3 Price of Anarchy and Braess' Paradox

One advantage of looking at stochastic population games from a routing game perspective is that it allows us to extend traditional routing game concepts into the stochastic game space. We consider the price of anarchy first and then briefly consider Braess' paradox.

Price of anarchy

In matrix form, the social cost in the infinite horizon case is computed as

$$J(x) = x^T\Big(\mathbf{P}_{\mathcal{A}}^T l(x) + r(x) + \mathbf{P}_{\mathcal{A}}^T I_o^T q(x)\Big) \tag{4.66}$$

$$= \bar{x}^T l(x) + x^T r(x) + \mathbf{x}^T q(x) \tag{4.67}$$

Just as in the standard routing game we can write a variational inequality characterization of the equilibrium which allows us to upper bound the price of anarchy using the Pigou bound.

Remark 18 *We present the variational inequality characterization in Proposition 3 and the price of anarchy bound in Theorem 7 in the infinite horizon case, but the results carry through directly in the finite horizon case as well.*

Proposition 3 (Wardrop Variational Inequality) *Let x be an infinite horizon stochastic Wardrop equilibrium and x' be any other feasible flow and let $\bar{x}$, $\bar{x}'$, $\mathbf{x}$, and $\mathbf{x}'$ be the corresponding edge and node flows respectively. The following inequality holds.*

$$x^T r(x) + \bar{x}^T l(x) + \mathbf{x}^T q(x) \leq x'^T r(x) + \bar{x}'^T l(x) + \mathbf{x}'^T q(x) \tag{4.68}$$

Proof 9 *Since both x and x' are feasible, we have that*

$$G\mathbf{P}_{\mathcal{A}} x = 0, \quad \mathbf{1}^T x = m, \quad x \geq 0, \qquad G\mathbf{P}_{\mathcal{A}} x' = 0, \quad \mathbf{1}^T x' = m, \quad x' \geq 0 \tag{4.69}$$

For the equilibrium flow x, the first order optimality conditions give

$$\mathbf{P}_{\mathcal{A}}^T l(x) + r(x) + \mathbf{P}_{\mathcal{A}}^T I_o^T q(x) = \mathbf{P}_{\mathcal{A}}^T G^T \pi + \mathbf{1}\lambda + \mu, \quad \mu \geq 0, \quad x^T \mu = 0 \tag{4.70}$$

Note that we also have $x'^T \mu \geq 0$ by positivity of x' and μ. It follows that

$$\begin{aligned} x'^T r(x) + \bar{x}'^T l(x) + \mathbf{x}'^T q(x) &= x'^T \left(\mathbf{P}_{\mathcal{A}}^T l(x) + r(x) + \mathbf{P}_{\mathcal{A}}^T I_o^T q(x) \right) && \text{(4.71a)} \\ &= x'^T \left(\mathbf{P}_{\mathcal{A}}^T G^T \pi + \mathbf{1}\lambda + \mu \right) = 0^T \pi + m\lambda + x'^T \mu && \text{(4.71b)} \\ &\geq x^T \left(\mathbf{P}_{\mathcal{A}}^T G^T \pi + \mathbf{1}\lambda + \mu \right) = x^T r(x) + \bar{x}^T l(x) + \mathbf{x}^T q(x) && \text{(4.71c)} \end{aligned}$$

Theorem 7 (Price of Anarchy Upper Bound for the MDP Routing Game) *Recall the Pigou bound from Definition 4.*

$$\alpha(\mathscr{L}) = \sup_{l \in \mathscr{L}} \sup_{x, x' \geq 0} \quad \frac{x \cdot l(x)}{x' \cdot l(x') + (x - x') l(x)} \tag{4.72}$$

For a nonempty class of loss functions $\mathscr{L}$ such that $r(\cdot), l(\cdot), q(\cdot) \in \mathscr{L}$,

$$PoA \leq \alpha(\mathscr{L}). \tag{4.73}$$

for the MDP routing game.

Proof 10 *(Again, we prove the result for the infinite horizon case and note that the same arguments go through in the finite horizon case.) Let x' and x be the socially optimal and*

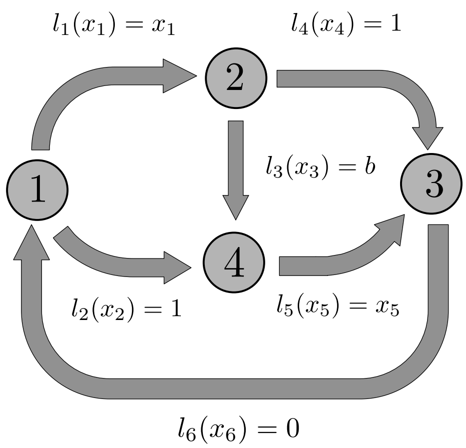

(a)

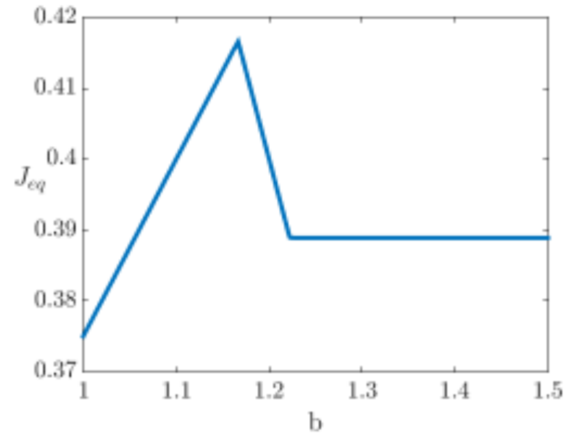

(b)

Figure 4.4: Illustration of Braess' paradox in the MDP routing game with deterministic transitions. (a) Braess graph. (b) Social cost at equilibrium for varying values of b.

equilibrium action masses and let $\bar{x}'$ and $\bar{x}$ and $\mathbf{x}'$ and $\mathbf{x}$ be the corresponding edge and node masses respectively. As in Proof 5, applying the Pigou bound to action masses gives

$$x'^T r(x') \geq \frac{x^T r(x)}{\alpha(\mathscr{L})} + (x' - x)^T r(x) \tag{4.74}$$

and similarly for the edge and node masses. Applying these inequalities and applying Proposition 3 gives

$$\begin{aligned}
J_{opt} &= x'^T r(x') + \bar{x}'^T l(\bar{x}') + \mathbf{x}'^T q(\mathbf{x}') && \text{(4.75a)}\\
&\geq \frac{x^T r(x) + \bar{x}^T l(\bar{x}) + \mathbf{x}^T q(\mathbf{x})}{\alpha(\mathscr{L})} + \left[x' - x\right]^T r(x) + \left[\bar{x}' - \bar{x}\right]^T l(\bar{x}) + \left[\mathbf{x}' - \mathbf{x}\right]^T q(\mathbf{x}) && \text{(4.75b)}\\
&\geq \frac{1}{\alpha(\mathscr{L})} J_{eq} && \text{(4.75c)}
\end{aligned}$$

Braess' paradox example

We now give a simple example of Braess' paradox in the MDP routing game case. We assume deterministic transitions and use the obvious graph derived from the graph used to illustrate Braess' paradox in the routing game case shown in Figure 4.4a. As can be seen in Figure 4.4b, increasing the latency on edge 3 causes the social cost to decrease.

4.4 Examples: Ridesharing and Urban Street Parking

Ridesharing: finite-horizon, deterministic transitions

To illustrate the model, we simulate a scenario where ridesharing drivers competing for customers on a weekend night in downtown San Francisco. We abstract the city as a set of neighborhoods (nodes) that drivers travel between. We assume the graph is fully connected and that all transitions are fully deterministic. x^t_{ij} is the population of drivers transitioning from neighborhood j to neighborhood i at time t. The loss functions that drivers consider are influenced by multiple factors including the fares they receive, their fuel costs, the time they spend traveling, and the time they spend waiting for customers. We use linear costs of the form

$$l^t_{ij}(x^t_{ij}) = -M^t_{ij} + (C^t_{ij})_{\text{travel}} + (C^t_{ij})_{\text{wait}} x^t_{ij} \tag{4.76}$$

The monetary reward of a trip M^t_{ij} has the form

$$M^t_{ij} = k \cdot \underbrace{\Big(\text{Rate}\Big)}_{\$/\text{mi}} \cdot \underbrace{\Big(\text{Dist}\Big)}_{\text{mi}} \tag{4.77}$$

where k is the surge pricing multiplication factor. The travel cost of the trip consists of travel time plus fuel costs.

$$(C^t_{ij})_{\text{travel}} = \tau \cdot \underbrace{\Big(\text{Dist}\Big)}_{\text{mi}} \cdot \underbrace{\Big(\text{Vel}\Big)^{-1}}_{\text{hr/mi}} + \underbrace{\Big(\substack{\text{Fuel}\\\text{Price}}\Big)}_{\$/\text{gal}} \cdot \underbrace{\Big(\substack{\text{Fuel}\\\text{Eff}}\Big)^{-1}}_{\text{gal/mi}} \cdot \underbrace{\Big(\text{Dist}\Big)}_{\text{mi}} \tag{4.78}$$

where τ is a time-money tradeoff parameter which we calculate by multiplying the ride rate (\$/mi) times the average distance between neighborhoods times the length of one time interval (20 min), assuming one trip per time interval. The last portion of the cost is the cost of waiting for jobs that depends on the other ridesharing drivers attempting to make the same transition. The coefficient $(C^t_{ij})_{\text{wait}}$ has units of \$ / driver and is defined as

$$(C^t_{ij})_{\text{wait}} = \tau \cdot \underbrace{\left(\frac{1}{\substack{\text{Customer}\\\text{Demand Rate}}}\right)}_{\text{hr/rides}} \tag{4.79}$$

The values that are not specifically edge dependent are listed in Table 4.1

We simulate the activity of ridesharing drivers in San Francisco over the course of a weekend evening from 7 pm to 1 am with every time step representing 20 min. A population of 20 drivers starts at each node. The various neighborhoods (modeled as nodes, and divided loosely into downtown and residential neighborhoods) are shown in Figure 4.5a and are listed in Table 4.2a. We assume that throughout the night there is at least a few customers (10

Rate	Velocity	Fuel Price	Fuel Eff
\$6 /mi	8 mph	\$2.5/gal	20 mi/gal

Table 4.1: Common values for latency calculations

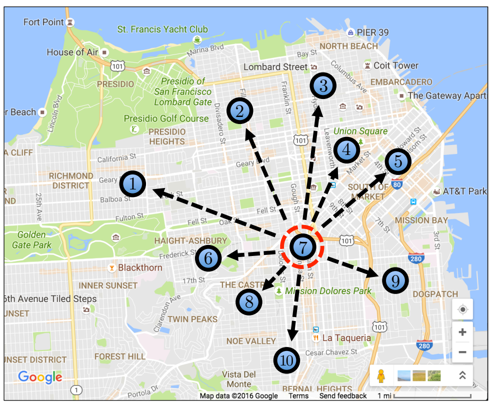

(a)

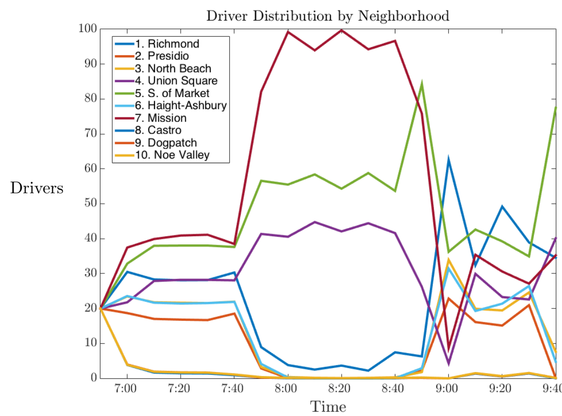

(b)

Figure 4.5: (a) Neighborhoods in San Francisco. (b) Population of drivers in each neighborhood at each time.

customers/hr) who want to travel between any two nodes. During the first few hours, most customers are traveling from residential neighborhoods to downtown neighborhoods. As the evening progresses, more customers are looking for rides among downtown neighborhoods, and then towards the end of the evening, most customers are looking to travel back to residential neighborhoods. The demand for rides between each of the different types of neighborhoods is detailed in Table 4.2b. We also add a surge pricing factor of 2 between the downtown nodes from 9-11pm and a surge pricing factor of 3 from downtown to residential nodes from 11pm-1am. We note that all the values in this simulation could be chosen much more accurately given google maps data and driver demand data. We solve the game by optimizing the potential function given in (4.47). We display the results in terms of the rewards drivers receive which are just the negative costs.

In Figure 4.5b, we show the population of drivers in each neighborhood over the course of the evening. Low number of drivers in certain neighborhoods could indicate a need to adjust the fares in those neighborhoods to maintain service for all customers throughout the evening. Given the population distribution at equilibrium, there are many optimal routes that drivers starting from each node can take over the course of the evening. In Figures 4.6a and 4.6b, we show the running reward and cumulative average reward for several optimal routes as well as several random routes starting from Node 1. Note that for the optimal

#	Neighborhood	Type
1	Richmond	Resident
2	Presidio	Resident
3	North Beach	Downtown
4	Union Square	Downtown
5	S. of Market	Downtown
6	Haight-Ashbury	Resident
7	Mission	Downtown
8	Castro	Resident
9	Dogpatch	Resident
10	Noe Valley	Resident

(a)

Rates $\left(\frac{\text{rides}}{\text{hr}}\right)$	Resident to Downtown	Downtown to Downtown	Downtown to Resident	Resident to Resident
7 pm – 9 pm	300	100	10	20
9 pm – 11 pm	100	200	100	20
11 pm – 1 am	10	50	300	20

(b)

Table 4.2: (a) Neighborhood Types. (b) Customer demand rates (rides/hr).

routes, the instantaneous running reward that drivers experience might go down or even be negative at one time step in order to setup for a large reward in the future. We note also that while the cumulative averages for each optimal route vary separately over time, they all become equal at the final time step. This has to be true at equilibrium for any two routes starting from the same node. Optimal routes starting at different nodes could have different total rewards. As expected, random routes achieve significantly less total reward over the time horizon. Two of the optimal traces are shown in Figures 4.6c and 4.6d.

Finally, we consider the decision that drivers make at an individual node at a specific time step. In Figures 4.7 and 4.7b, we show decision criteria that drivers face at node 7 (the Mission, Figure 4.5a) at time steps $t = 9$ and $t = 17$. This decision criteria includes both the immediate reward for a specific transition and the expected reward-to-go. Notice that population mass is only distributed among transitions choices that achieve the maximum expected reward.

Ridesharing: infinite-horizon, stochastic transitions

In this section, we simulate a modified version of the of the ridesharing game where drivers do not get to specifically choose which riders they want to take. This model could be more useful for studying the incentives of taxi drivers. We assume that at each node a certain percentage of riders want to travel to each of the other nodes. We assume that these percentages are given a priori.

At each node, the taxi driver can choose from several actions. The first action, a_r is to wait for a random rider and transition to whatever node that rider wants to go to. The transition probabilities of this action are determined by the percentages of riders at that node that want to make specific trips. The other actions the driver can choose from are to transition to some other node without a rider. We will refer to the driver transitioning to

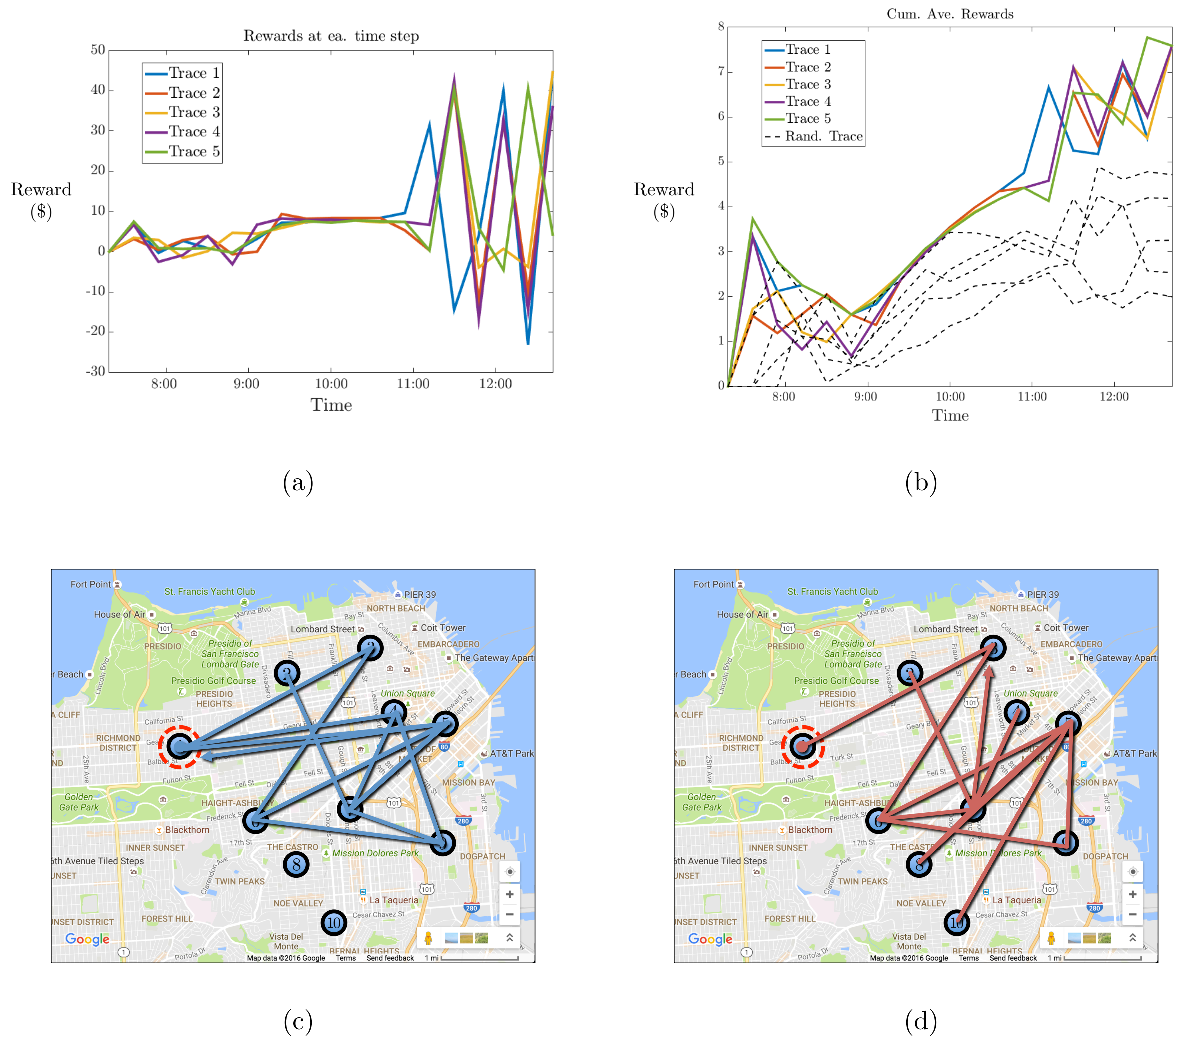

Figure 4.6: (a) Running reward for various optimal routes. (b) Cumulative average reward for optimal and suboptimal routes. (c) Trace 1. (d) Trace (2)

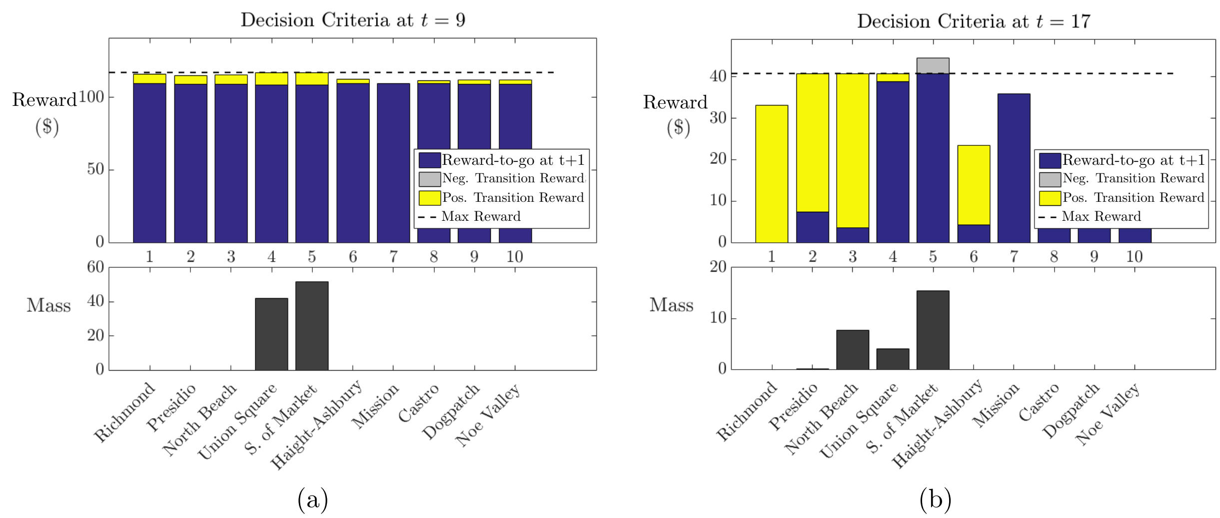

Figure 4.7: Decision criteria at (a) $t = 9$ and (b) $t = 17$ showing immediate reward and expected reward-to-go.

To ea. \ From ea.	Resident	Downtown
Resident	0.06	0.083
Downtown	0.175	0.167

Table 4.3: Probability of transitioning from one neighborhood to another when a driver waits for a rider (based on percentage of riders making each transition).

node i without a rider as a_i. In general, this would result in the driver paying the travel costs without receiving a fare; however, there is a small possibility that the driver will find a customer along the way.

In order to model this scenario, we define a graph with two sets of edges going between each node. The first set model transitioning with a rider. We will denote these edges $\mathcal{E}_{\text{rider}}$. The second set model transitioning without a rider which we will denote $\mathcal{E}_{\text{norider}}$. This allows us to differentiate between the rewards received for taking a rider and for driving without a rider.

We assume the same node structure as in the finite horizon example of downtown San Francisco presented above. The transition probabilities for each action a_r (based on rider percentages) are shown in Table 4.3.

Note that drivers who take a_r at each node travel on edges in $\mathcal{E}_{\text{rider}}$. If a driver takes the action of transitioning to another without a rider, they take the appropriate edge in $\mathcal{E}_{\text{norider}}$ with probability 0.82 and they take each of the edges in $\mathcal{E}_{\text{rider}}$ (coming from that node) with probability 0.02. This is meant to represent the small chance that they might pick up a rider along the way.

The costs on the actions are given as

$$r_j^{a_r} = (C_j)_{\text{wait}} x_j^{a_r} \tag{4.80a}$$
$$r_j^{a_i} = 0 \tag{4.80b}$$

i.e. drivers who wait for a rider have to compete with other drivers taking that action while drivers who transition without a rider pay no action cost. Here, as in the finite horizon example, we take the waiting cost coefficient to be

$$(C_j)_{\text{wait}} = \tau \cdot \underbrace{\left(\frac{1}{\substack{\text{Customer}\\ \text{Demand Rate}}}\right)}_{\text{hr/rides}} \tag{4.81}$$

Driver demand rates at each node are shown in the table below.

Neighborhood type	Resident	Downtown
Demand (rides/hr)	20	50

The transition costs differ depending on whether the drivers take a rider or not (whether they take an edge in $\mathcal{E}_{\text{rider}}$ or an edge in $\mathcal{E}_{\text{norider}}$). The transition costs are given by

$$l_e(x) = \begin{cases} -M_e + (C_e)_{\text{travel}} & ; \text{ if } e \in \mathcal{E}_{\text{rider}} \\ (C_e)_{\text{travel}} & ; \text{ if } e \in \mathcal{E}_{\text{norider}} \end{cases} \tag{4.82}$$

where M_e is the fare for a trip on edge e and $(C_e)_{\text{travel}}$ is the cost of travel on that edge. We assume the same form and values for M_e as in the finite-horizon example (with a surge pricing factor of $k = 1$, Equation (4.77)) and we assume the same form and values for $(C_e)_{\text{travel}}$ as in Equation (4.78).

We compute both the equilibrium strategies and the socially optimal strategies in the infinite horizon game. Figure 4.8 shows the steady state distribution of drivers at the nodes in both cases including the portion that take riders and the portion that do not. Figure 4.9 shows the portion of the population that is making transitions without riders. Interestingly, the socially optimal strategies involve more drivers making trips without riders.

Circling for parking: infinite-horizon, deterministic transitions

Another application of this model is determining the optimal strategies for urban drivers circling city blocks looking for places to park. We consider a sample set of city blocks shown in Figure 4.10a. We solve the problem on the dual graph as it allows us more freedom to restrict certain transitions (left turns, U-turns, etc.). We write down a loss function for each node (block face) that represents the expected waiting time to find an open space.

The expected waiting time for a given member of the population is related to the rate at which spots become available and the probability that when a spot becomes available that

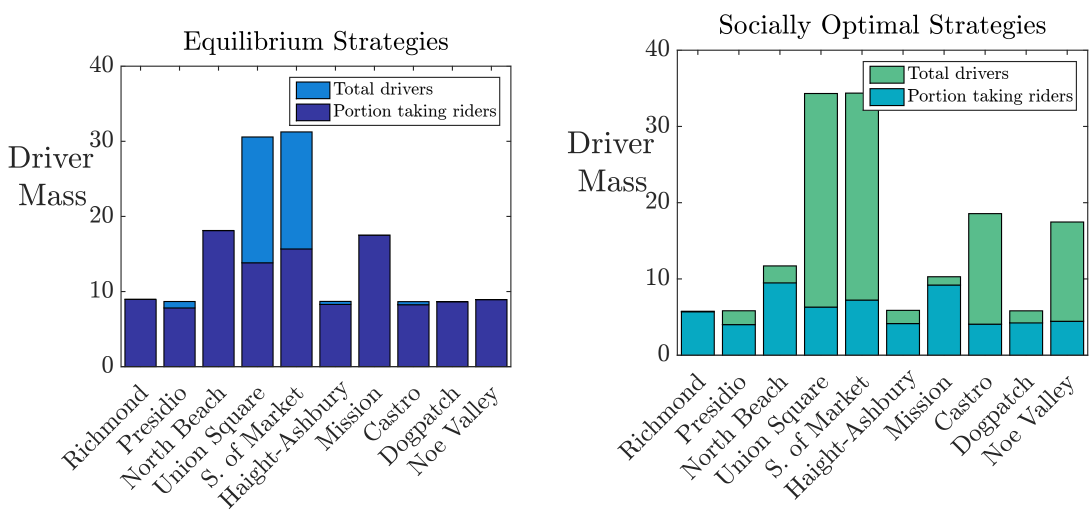

(a) (b)

Figure 4.8: Steady state distribution of drivers at each node under the (a) equilibrium strategies and (b) socially optimal strategies showing the portion of drivers that take riders and the portion that do not take riders.

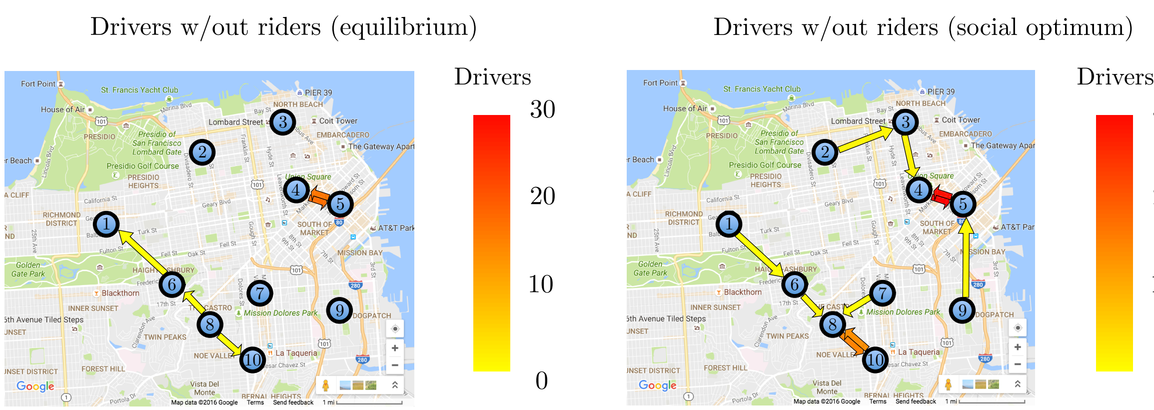

(a) (b)

Figure 4.9: Drivers transitioning between nodes without riders in the (a) equilibrium case and (b) under the socially optimal strategies.

individual driver will get it which depends on congestion. For the probability of finding an empty spot we draw inspiration from a queuing model where customers are selected randomly from a queue. We assume that block face j has c_j spots and each spot becomes available at a rate of γ_j (in units of car/min). The total cost of driving by each block face is given by

$$q_j(x) = M_j p_j + \Big((C_j)_{\text{wait}} + M_j\Big)(1 - p_j) \tag{4.83}$$

Here, M_j is the cost of parking and $(C_j)_{\text{wait}}$ is the cost of waiting for a space which we take to be the $(C_j)_{\text{wait}} = \tau \Delta t_j$ where τ is a time vs. money tradeoff parameter and Δt_j is the average amount of time spent on block face j. The probability of an individual driver getting a parking spot, p_j, is given by

$$p_j = \Big(1 - e^{-c_j \gamma_j \Delta t_j}\Big) \frac{1}{1 + x_j} \tag{4.84}$$

This equation consists of two parts. The first is the probability that a space opens up in the time the driver spends on the block face. The second part is the probability that an individual driver of the population on the edge gets that space. This second term is 1 when there is 0 mass on the edge, $\frac{1}{2}$ when there is a mass of 1 other driver on the edge, $\frac{1}{3}$ when there is a mass of 2 other drivers on the edge, etc.

The full loss function can be rewritten as

$$q_j(x) = \underbrace{\Big(M_j + (C_j)_{\text{wait}}\Big)}_{a_j} \underbrace{-(C_j)_{\text{wait}}\Big(1 - e^{-c_j \gamma_j \Delta t_j}\Big)}_{b_j} \frac{1}{1 + x_j} \tag{4.85}$$

The potential function is then given by

$$F(x) = \sum_j \int_0^{x_j} q_j(u)\ du = \sum_j \Big[a_j x_j + b_i \ln\big(1 + x_j\big)\Big] \tag{4.86}$$

We use the following parameter values

Δt_j (min)	γ_j (1/hr)	M_j (\$)	τ (\$/min)
0.5	1/120	5	0.5

for each street but we look at a scenario where there are different numbers of parking spaces on each street according to the values shown in the table below.

Street	1-6	7-8	9-10	11-14	15-16	17-18	19 -24
Num. spaces (c_j)	20	10	20	10	20	10	20

The streets are numbered in Figure 4.10a.

We consider this game in the infinite horizon case with deterministic transitions. In Figure 4.10, we look at the traffic distribution at the Wardrop equilibrium, at the social

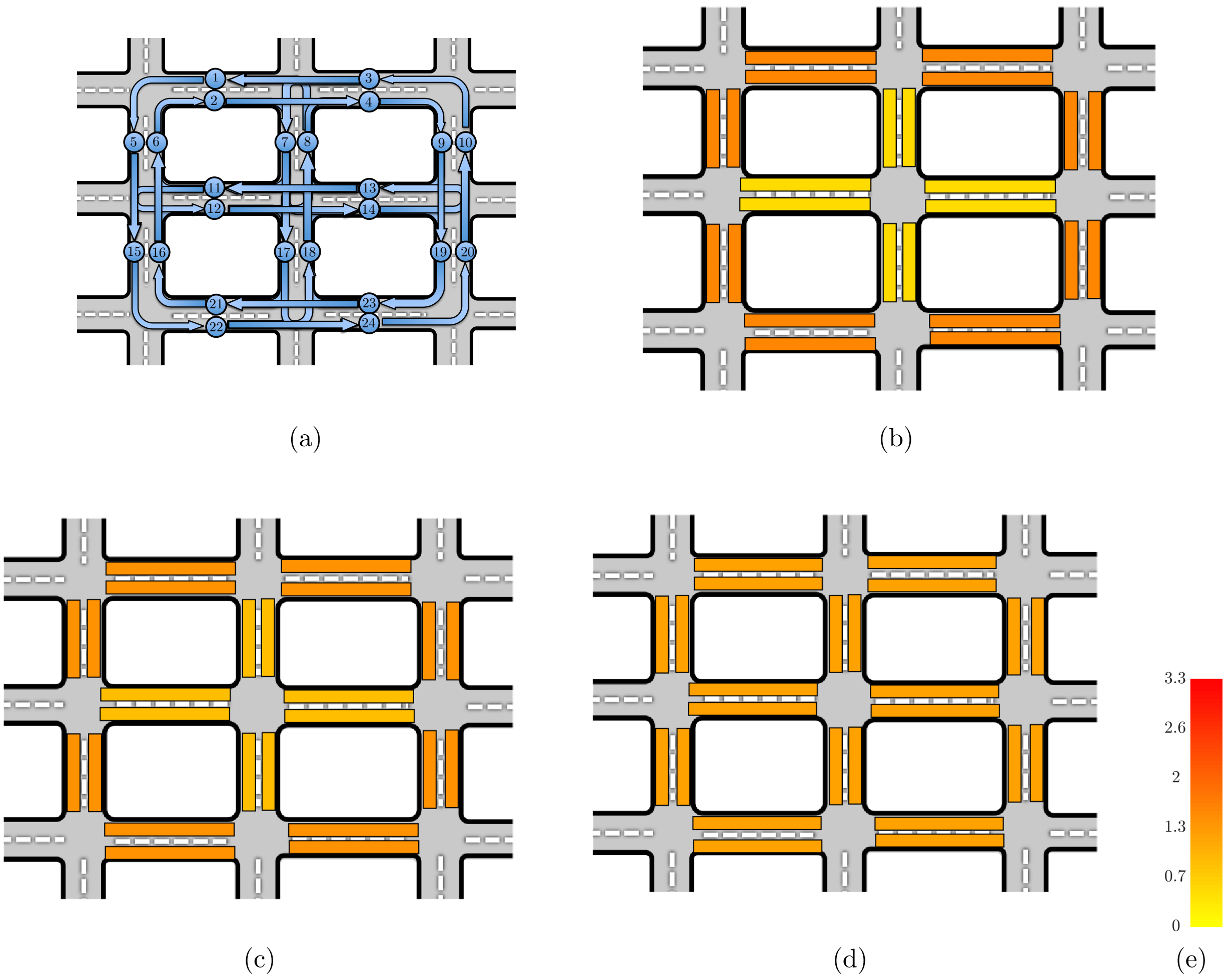

Figure 4.10: (a) Block faces and allowed transitions. (b) Wardrop equilibrium. (c) Social optimum. (d) Uniform turns. (e) Scale (cars).

optimum, and when the population members use a uniform turning strategy. We note that one application of this model could be determining how to vary parking prices between streets in order to improve congestion since even in this simple scenario the Nash solution is suboptimal. We also note that a model like this could be integrated with the queue-routing game developed in Chapter 3 to provide a more sophisticated model of parking behavior.

Chapter 5

External-Cost Routing Game

In traditional routing games, transportation users only care about their travel time. In real world scenarios, however, there are likely to be other factors that population members consider such as monetary cost, convenience, etc. The tradeoffs between these different factors are also likely to vary among different members of the population.

Multi-criteria routing equilibria were considered early on by Dial [95] extending the work of Quandt [96] and Schneider [97]. These early formulations considered multiple decision criteria (time and money) but ignored the effects of congestion. Congestion effects were added by Dafermos [98] in the deterministic case and Daganzo [99] in the stochastic case. In general, multi-criteria models either divide up the population into a finite classes of users each with distinct preferences or they represent the population by some distribution (infinite classes of users). Nagurney and Dong [100, 101], and Li and Chen [102] consider finite classes of users using variational inequality approaches. Yang [103] also considers finite user classes and studies the comparison between the socially optimal flows and equilibrium flows.

Leurent [104, 105] was one of the earliest to represent the population's value for time vs. money as a continuous distribution over some positive interval and formulated an optimization problem for finding the equilibrium. Dial [106, 107] considered a more general scenario where both travel time and monetary cost can depend on congestion and framed the problem as a variational inequality. Marcotte, et al. [108–111] generalized the work of Leurent and Dial presenting a general variational inequality formulation. Much of the strength of these formulations has been methods to turn infinite dimensional variational inequalities into finite dimensional problems [111, 112]. Our work follows in the mode of these models. Much of the focus since then has been using these models to devise tolling schemes, see [113–115] in the finite user class case and [116–122] in the continuous distribution case. Another recent branch of research has applied multi-objective optimization techniques to the external factor problem. Raith, Wang, et al. [123–127] consider an equilibrium they call a *bi-objective user equilibrium* where no transportation users can improve their travel time or toll cost without worsening the other criteria. They also consider similar equilibrium definition for more than two objectives [128, 129].

In this chapter, we present an optimization formulation that revisits the formulations of Marcotte, Leurent, and Dial. Population members consider an external factor (money, convenience,etc) as well as well as travel time in their routing decisions. Population preference for this external factor in relation to travel time is represented by an arbitrary distribution over a parameter θ. We model the whole set of routes as being divided into subsets $\{\mathcal{R}_o\}_{o\in\mathcal{O}}$ that each come with an external cost α_o. Drivers in the population select which subset of routes they want to use and then within that subset, they select which route they want to take. The total cost for selecting subset $\mathcal{R}_o$ and route $r \in \mathcal{R}_o$ is given by

$$\ell_r(z) + \alpha_o\theta \tag{5.1}$$

This cost depends on the total mass distribution z and the travel time $\ell_r(z)$ but it also depends on θ, the individual population member's value of the external cost α_o. We present the appropriate equilibrium definition for agents who consider this cost; and then from this definition, we derive how the population will divide itself up among the various transportation options. We also give a potential function and the appropriate optimization problem that can be used to compute the population mass distribution associated with the equilibrium.

Our approach revisits the formulations of Leurent and Marcotte. It is a cleaner version of Leurent's optimization formulation [104] and a subcase of Marcotte's variational inequality formulation [111] where the external factor cost does not depend on traffic flow. We present our own simple proofs of the form of the equilibrium and equivalence between the minimum of our potential function optimization problem and the equilibrium. A main distinction between our presentation and these previous works is that Leurent, Dial, and Marcotte assume the distribution over the external cost parameter (which they consider to be the monetary value of time) is supported on the non-negative reals ($\theta \in [0, \infty)$ in our notation). Our proofs highlight the fact that this is not necessary. In addition, the external costs α_o can be positive or negative. Either θ or α_o being negative does not make sense when the external cost is money, but it can be useful when the external cost represents something else. At the end of this chapter, we specifically comment on how this framework could be used to model arbitrary preference for one mode of transportation over another.

Along with presenting our formulation, we discuss several contexts where this framework could be applied to analyze non-monetary tradeoffs. An interesting example is understanding the impact of a population's concern for their location privacy. Recent work has explored privacy in routing games from the view point of differential privacy [130, 131]. Using our framework, we can analyze how drivers' concern for their location privacy will shift the traffic equilibrium. We can think of some privacy price, α_{priv} that drivers pay whenever they use navigation services that monitor their location. Drivers who choose not to use these services only have access to a limited set of routes without congestion information.

As mentioned above, we can use also use this framework to compare different modes of transportation as well where the parameter θ represents the population's preference for one form over the other. Our work differs from previous work using a continuous distribution to represent transportation mode preferences [132] in that we take advantage of the fact that θ and each α_i can be either negative or positive in order to model the fact that some members of the population may prefer different travel modes even when the travel time is equal. We go into further detail in Section 5.2.

The rest of this chapter is organized as follows. In Section 5.1, we present our equilibrium concept and derive some of its properties. We then give a potential function and the corresponding optimization problem for computing the equilibrium as well as several simple examples. In Section 5.2, we discuss several new applications for this framework.

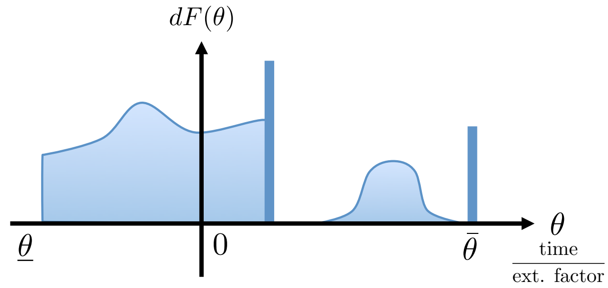

Figure 5.1: Example population distribution of time vs. external factor tradeoff.

5.1 External-Cost Wardrop Equilibrium

Setup

Let $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ be a graph with nodes and edges and let $\mathcal{R}$ be the set of all routes through a network from an origin node to a destination node. Let $\{\mathcal{R}_o\}_{o \in \mathcal{O}}$ be a collection of subsets of routes. Each subset of routes, $\mathcal{R}_o$ has a price, α_o that users have to pay in order to use those particular routes. Without loss of generality, we will assume that the route subsets are ordered by price,

$$\alpha_{o-1} > \alpha_o \qquad \forall o \in \mathcal{O} \tag{5.2}$$

Once they pay to use a certain set of routes, drivers then play a standard routing game on those routes (though their congestion costs may depend on members of the population who choose a different set of routes.) We note that we do not assume any specific structure on the subsets $\{\mathcal{R}_o\}_{o \in \mathcal{O}}$. We also note that it is not a restriction to assume that α_{o-1} is strictly greater than α_o. If we have two groups, $\mathcal{R}_o$ and $\mathcal{R}_{o'}$ such that $\alpha_o = \alpha_{o'}$, we can just consider $\mathcal{R}_o \cup \mathcal{R}_{o'}$ as one group.

As previously mentioned, each member of the population has some type θ that represents their tradeoff between time and the external factor, i.e. how much external cost they are willing to incur for access to quicker routes through the network. We assume we are given some population distribution for this parameter, $dF(\theta)$. A sample distribution is illustrated in Figure 5.1. Note that this distribution can be supported above and below 0. $\bar{\theta}$ is the maximum time-money tradeoff for anyone in the population and $\underline{\theta}$ is the minimum tradeoff.

The overall population's decision to pay for each set of routes is encoded by a vector valued indicator function $I : [\underline{\theta}, \bar{\theta}] \to \Delta_{|\mathcal{O}|}$ where $\Delta_{|\mathcal{O}|}$ is the simplex of dimension $|\mathcal{O}|$. The oth element $I_o(\theta)$ represents the fraction of users with type θ that select subset $\mathcal{R}_o$. Here, we have implicitly made the assumption that every member of the population chooses some option. We can compute the total mass of users that pay for routes in $\mathcal{R}_o$ as

$$m_o(I_o) = \int_{[\underline{\theta},\bar{\theta}]} I_o(\theta) dF(\theta) \tag{5.3}$$

This mass is then divided up over the various routes in $\mathcal{R}_o$. Let $z_o \in \mathbb{R}_+^{|\mathcal{R}_o|}$ be the vector of masses assigned to each of these routes. We have that

$$\sum_{r \in \mathcal{R}_o} (z_o)_r = m_o \tag{5.4}$$

The total population mass m is given by

$$m = \sum_o m_o = \int_{[\underline{\theta},\bar{\theta}]} \sum_o I_o(\theta) dF(\theta) \tag{5.5}$$

We will also use $z = (z_o)_{o \in \mathcal{O}}$ as a short hand for the entire set of mass distributions.

Given $I = (I_o)_{o \in \mathcal{O}}$ and the corresponding mass distributions $z = (z_o)_{o \in \mathcal{O}}$, we can compute the total flows on each edge of the network $x \in \mathbb{R}_+^{|\mathcal{E}|}$ as

$$x = \sum_o \mathbf{E}_{\mathcal{R}_o} z_o \tag{5.6}$$

where $\mathbf{E}_{\mathcal{R}_o} \in \{0,1\}^{|\mathcal{E}| \times |\mathcal{R}_o|}$ is the routing matrix for the routes in $\mathcal{R}_o$.

Equilibrium

We now define the appropriate equilibrium concept.

Definition 10 (External-Cost Wardrop Equilibrium) *An* external-cost Wardrop equilibrium *is a set of measurable functions $I = (I_o)_{o \in \mathcal{O}} : [\underline{\theta}, \bar{\theta}] \to \Delta_{|\mathcal{O}|}$ and corresponding mass distributions $z = (z_o)_{o \in \mathcal{O}}$ satisfying Equations* (5.3)*,* (5.4)*, and* (5.5) *that satisfies the following. For any $\theta \in [\underline{\theta}, \bar{\theta}]$ and $o \in \mathcal{O}$ such that $I_o(\theta) > 0$ and any $r \in \mathcal{R}_o$ such that $(z_o)_r > 0$*

$$\ell_r(z) + \alpha_o \theta \leq \ell_{r'}(z) + \alpha_{o'} \theta \tag{5.7}$$

for any $o' \in \mathcal{O}$ and $r' \in \mathcal{R}_{o'}$.

Intuitively, this states that any population member of type θ who pays for $\mathcal{R}_o$ and drives a specific route could not have done better by selecting a different subset of routes and/or a different route. We note that whenever $\alpha_o = \alpha_{o'}$ this reduces to the standard Wardrop equilibrium condition. Thus, all population members who choose a specific subset of routes are playing a Wardrop equilibrium within that subset. We also note that like in the traditional routing game, the equilibrium condition only places restrictions on strategies with positive mass ($m_o > 0$); however, Equation (5.3) requires that $I_o(\theta) = 0$ almost everywhere whenever $m_o = 0$.

We now deduce several properties about the choice functions $I = (I_o)_{o \in \mathcal{O}}$.

Lemma 1 *Assume an ordering on* $\{\alpha_o\}_{o\in\mathcal{O}}$ *such that*

$$\alpha_1 > \cdots > \alpha_{|\mathcal{O}|} \tag{5.8}$$

Suppose I *and* z *form an external-cost Wardrop equilibrium. Then there exists* $\{\bar{\theta}_o\}_{o\in\mathcal{O}}$ *such that*

$$\underline{\theta} \le \bar{\theta}_1 \le \cdots \le \bar{\theta}_{|\mathcal{O}|} \le \bar{\theta} \tag{5.9}$$

and $\{\underline{\gamma}_o\}_{o\in\mathcal{O}} \in [0,1]$ *and* $\{\bar{\gamma}_o\}_{o\in\mathcal{O}} \in [0,1]$ *such that* I *satisfies*

$$I_o(\theta) = \begin{cases} \underline{\gamma}_o & ; \text{ if } \theta = \bar{\theta}_{o-1} \\ 1 & ; \text{ if } \bar{\theta}_{o-1} < \theta < \bar{\theta}_o \\ \bar{\gamma}_o & ; \text{ if } \theta = \bar{\theta}_o \\ 0 & ; \text{ otherwise} \end{cases} \tag{5.10}$$

almost everywhere (where we define $\bar{\theta}_0 = \underline{\theta}$*). Whenever* $\bar{\theta}_{o-1} = \bar{\theta}_o$*,* $\underline{\gamma}_o = \bar{\gamma}_o$*.*

The set $\{I_o\}_{o\in\mathcal{O}}$ is illustrated in Figure 5.2. The set $\{\bar{\theta}_o\}_{o\in\mathcal{O}}$ are the critical points detailed by Marcotte [111, 112].

Proof 11 *First we note that* (5.3) *guarantees that* $I_o(\theta) = 0$ *almost everywhere for any* o *such that* $m_o = 0$*. For every* $o \in \mathcal{O}$ *such that* $m_o > 0$*, define*

$$\underline{\theta}_o = \inf_\theta \{\theta : I_o(\theta) > 0\} \qquad \bar{\theta}_o = \sup_\theta \{\theta : I_o(\theta) > 0\} \tag{5.11}$$

First, we show that for any two options o *and* o' *such that* $\alpha_{o'} > \alpha_o$ *and both* $m_o > 0$ *and* $m_{o'} > 0$*, then* $\bar{\theta}_{o'} \le \underline{\theta}_o$*. Assume not, i.e. that* $\underline{\theta}_o < \bar{\theta}_{o'}$*. Select* $\theta, \theta' \in [\underline{\theta}_o, \bar{\theta}_{o'}]$ *such that* $\theta < \theta'$*,* $I_o(\theta) > 0$*, and* $I_{o'}(\theta') > 0$*. Select routes* $r \in \mathcal{R}_o$ *and* $r' \in \mathcal{R}_{o'}$ *such that both* $(z_o)_r > 0$ *and* $(z_{o'})_{r'} > 0$*. Applying* (5.7) *at* θ *and* θ' *respectively gives.*

$$(\alpha_{o'} - \alpha_o)\theta \ge \ell_r - \ell_{r'} \ge (\alpha_{o'} - \alpha_o)\theta' \tag{5.12}$$

Since $\alpha_{o'} - \alpha_o > 0$*, it follows that* $\theta \ge \theta'$ *which is a contradiction. Thus we have that* $\bar{\theta}_{o'} \le \underline{\theta}_o$ *for any options with positive mass. For any option* o *with* $m_o = 0$*, let* $\bar{\theta}_o = \bar{\theta}_{o-1}$*. Since* $I(\theta)$ *maps to the simplex, we have that* $I_o(\theta) = 1$ *almost everywhere for* $\theta \in (\bar{\theta}_{o-1}, \bar{\theta}_o)$*.*

We note that given the form of $\{I_o\}_{o\in\mathcal{O}}$ expounded in Lemma 1, we can compute $\{\bar{\theta}_o\}_{o\in\mathcal{O}}$ using the cumulative distribution function of $dF(\theta)$. Let CDF : $\theta \mapsto m$ be the cumulative distribution function. Define a function $\Theta : m \mapsto \theta$ as $\Theta(\cdot) = \text{CDF}^{-1}(\cdot)$. We can then compute $\bar{\theta}_o$ as

$$\bar{\theta}_o = \Theta\Big(\textstyle\sum_{i\le o} m_i\Big) \tag{5.13}$$

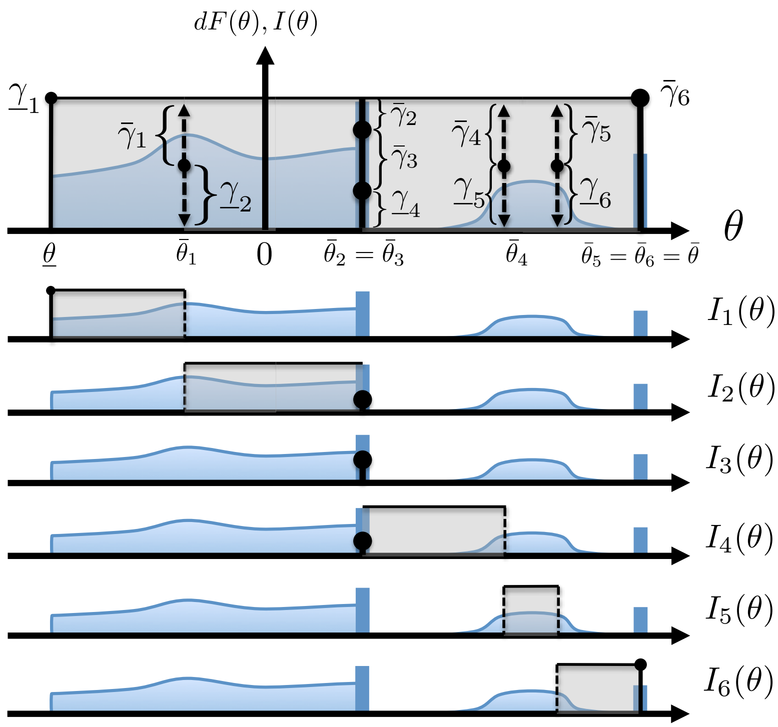

Figure 5.2: Illustration of a possible set of indicator functions, $\{I_o\}_{o\in\mathcal{O}}$, at equilibrium for the distribution in Figure 5.1.

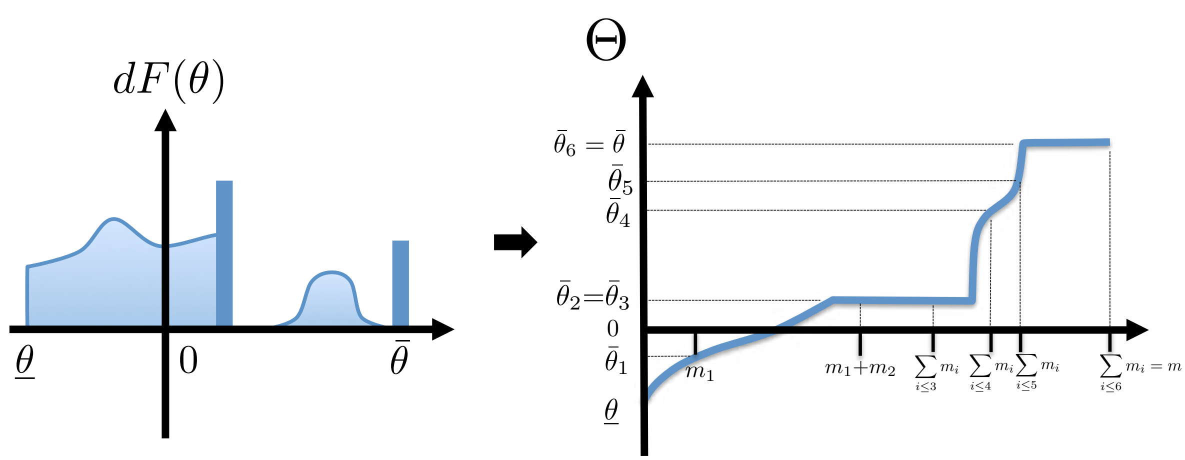

Figure 5.3: Inverse cummulative distribution function, $\Theta(m)$, for the distribution in 5.1.

This inverse cumulative distribution is illustrated in Figure 5.3. We can also compute the values of $\{\underline{\gamma}_o\}_{o\in\mathcal{O}}$ and $\{\bar{\gamma}_o\}_{o\in\mathcal{O}}$ using the masses $\{m_o\}_{o\in\mathcal{O}}$. Whenever $m_o = 0$, then $\underline{\gamma}_o = \bar{\gamma}_o = 0$. Whenever $m_o > 0$ and $\bar{\theta}_{o-1} = \bar{\theta}_o$, then

$$\underline{\gamma}_o = \bar{\gamma}_o = \frac{m_o}{F(\{\bar{\theta}_o\})} \tag{5.14}$$

The remaining values are then computed inductively. Let $1, \ldots, o'$ be the set of options such that $\underline{\theta} = \bar{\theta}_o$ for $o \leq o'$. We can compute

$$\underline{\gamma}_{o'+1} = 1 - \sum_{i=1}^{o'} \bar{\gamma}_i \tag{5.15}$$

and we can use

$$m_{o'+1} = \underline{\gamma}_{o'+1} F(\{\bar{\theta}_{o'}\}) + F((\bar{\theta}_{o'}, \bar{\theta}_{o'+1})) + \bar{\gamma}_{o'+1} F(\{\bar{\theta}_{o'+1}\}) \tag{5.16}$$

to solve for $\bar{\gamma}_{o'+1}$. Similarly, we can repeat this procedure with $\bar{\theta}_{o'+1}$ instead of $\underline{\theta}$. The rest of the $\gamma' s$ follow by induction.

We note that whenever $\bar{\theta}_o$ falls on a set of measure zero (as is the case with $\bar{\theta}_1$, $\bar{\theta}_4$, and $\bar{\theta}_5$ in Figure 5.2), there is some ambiguity in the choice of $\bar{\gamma}_o$ and $\underline{\gamma}_{o-1}$. In this case, the choice of γ does not affect the equilibrium mass distribution.

Potential function and optimization problem

Lemma 1 indicates that given an equilibrium mass distribution $\{z_o\}_{o\in\mathcal{O}}$, we can determine the choice functions $\{I_o(\theta)\}_{o\in\mathcal{O}}$ of the form (5.10) using Equations (5.13) and solving for gamma inductively as discussed above. These arguments allow us to solve for the equilibrium mass directly. We define the appropriate potential function and show that the KKT necessary conditions for minimizing this function with respect to the appropriate constraints gives a mass distribution satisfying the external-cost Wardrop equilibrium conditions.

The potential function is given by

$$F(z) = \sum_e \int_0^{x_e} l_e(u)\, du + \sum_i \int_{\sum_{o<i} m_o}^{\sum_{o\leq i} m_o} \alpha_i \Theta(u)\, du \tag{5.17}$$

We note that we could also write this potential as

$$F(z) = \sum_e \int_0^{x_e} l_e(u)\, du + \sum_{i=1}^{|\mathcal{O}|-1} \int_0^{\sum_{o\leq i} m_o} (\alpha_i - \alpha_{i+1})\Theta(u)\, du + \int_0^{\sum_{o\leq|\mathcal{O}|} m_o} \alpha_{|\mathcal{O}|}\Theta(u)\, du \tag{5.18}$$

On the set defined by the conservation of mass constraint (Equation (5.5)), the last term is constant. Indeed, we can write

$$F(z) = \sum_e \int_0^{x_e} l_e(u)\ du + \sum_{i=1}^{|\mathcal{O}|-1} \int_0^{\sum_{o\leq i} m_o} (\alpha_i - \alpha_{i+1})\Theta(u)\ du + \underbrace{\int_0^m \alpha_{|\mathcal{O}|}\Theta(u)\ du}_{\text{constant}} \tag{5.19}$$

In this form, we can see that the potential function is convex on the set defined by Equation (5.5). This follows from the fact that $\Theta(\cdot)$ is increasing for any distribution $dF(\theta)$ as illustrated in Figure 5.3 and $\alpha_i - \alpha_{i+1} > 0$ for any set of α_i's that satisfy the ordering convention (5.2). Note that we can write $F(z)$ as a function of z only since $\{m_o\}_{o\in\mathcal{O}}$ and x are both functions of z.

Remark 19 *This potential function is closely related to Leurent's objective function, Equation (8) in [104], and is equivalent to the potential function given in Equation (54) in [111] and Equation (40) in [112] when θ is supported only on $\mathbb{R}^+$.*

We can write the following optimization problem for finding the equilibrium.

Theorem 8 *Let $z = (z_o)_{o\in\mathcal{O}}$ be a mass distribution that solves the following optimization problem and $I = (I_o)_{o\in\mathcal{O}}$ be a set of choice functions with form determined by* (5.10), (5.13), *and the inductive procedure outlined above.*

$$\min_z \quad F(z) \tag{5.20a}$$

$$s.t. \quad m = \sum_o \sum_{r\in\mathcal{R}_o} (z_o)_r, \quad z_o \geq 0, \qquad \forall o \in \mathcal{O} \tag{5.20b}$$

$$m_o = \sum_{r\in\mathcal{R}_o} (z_o)_r \tag{5.20c}$$

$$x = \sum_{o\in\mathcal{O}} \mathbf{E}_{\mathcal{R}_o} z_o \tag{5.20d}$$

It follows that $z = (z_o)_{o\in\mathcal{O}}$ and $I = (I_o)_{o\in\mathcal{O}}$ are an external-cost Wardrop equilibrium.

Proof 12 *The Lagrangian is given by*

$$\mathcal{L}(z,\lambda,\mu) = F(z) - \lambda\left(m - \sum_o \sum_{r\in\mathcal{R}_o} (z_o)_r\right) - \sum_o \mu_o^T z_o \tag{5.21}$$

where $\lambda \in \mathbb{R}$ and $\mu_o \in \mathbb{R}_+^{|\mathcal{R}_o|}$ and where we have substituted in (5.20c) *and* (5.20d). *For a given $r \in \mathcal{R}_o$, the first order optimality conditions give*

$$\ell_r + \sum_{o\leq i} \alpha_i \bar{\theta}_i - \sum_{o<i} \alpha_i \bar{\theta}_{i-1} = \lambda + (\mu_o)_r \tag{5.22}$$

where $(\mu_o)_r \geq 0$ *with equality achieved whenever* $(z_o)_r > 0$ *by complementary slackness. We can rewrite Equation* (5.22) *in two different ways*

$$\ell_r + \sum_{o \leq i} (\alpha_i - \alpha_{i+1})\bar{\theta}_i = \lambda + (\mu_o)_r \tag{5.23a}$$

$$\ell_r + \sum_{o < i} (\alpha_{i-1} - \alpha_i)\bar{\theta}_{i-1} = \lambda + (\mu_o)_r \tag{5.23b}$$

We now consider a specific choice function $I_o(\theta)$ *and a route* $r \in \mathcal{R}_o$ *such that* $(z_o)_r > 0$. *By Lemma 1, we simply need to show that* (5.7) *is satisfied for any* $\theta \in [\bar{\theta}_{o-1}, \bar{\theta}_o]$. *Take any* $\theta \in [\bar{\theta}_{o-1}, \bar{\theta}_o]$ *and two subsets* $o, o' \in \mathcal{O}$ *such that* $o' > o$ *and corresponding routes* $r \in \mathcal{R}_o$ *and* $r' \in \mathcal{R}_{o'}$ *such that* $(z_o)_r > 0$. *From* (5.23a), *we have that*

$$\ell_r - \ell_{r'} + \sum_{o \leq i < o'} (\alpha_i - \alpha_{i+1})\bar{\theta}_i = (\mu_o)_r - (\mu_{o'})_{r'} \tag{5.24a}$$

$$\ell_r - \ell_{r'} + \sum_{o \leq i < o'} (\alpha_i - \alpha_{i+1})\theta \leq (\mu_o)_r - (\mu_{o'})_{r'} \tag{5.24b}$$

$$\ell_r - \ell_{r'} + (\alpha_o - \alpha_{o'})\theta \leq (\mu_o)_r - (\mu_{o'})_{r'} \tag{5.24c}$$

where, in Equation (5.24b), *we have used that* $\alpha_i - \alpha_{i+1} > 0$ *for all* i *and* $\theta \leq \bar{\theta}_i$ *for all* $i \geq o$. *Since* $(z_o)_r > 0$ *(and thus by complementary slackness* $(\mu_o)_r = 0$*) and* $(\mu_{o'})_{r'} \geq 0$, *it follows that*

$$\ell_r + \alpha_o\theta \leq \ell_{r'} + \alpha_{o'}\theta \tag{5.25}$$

Now if $o' < o$, *we have that*

$$\ell_{r'} - \ell_r + \sum_{o' < i \leq o} (\alpha_{i-1} - \alpha_i)\bar{\theta}_{i-1} = (\mu_o)_{r'} - (\mu_{o'})_r \tag{5.26a}$$

$$\ell_{r'} - \ell_r + \sum_{o' < i \leq o} (\alpha_{i-1} - \alpha_i)\theta \geq (\mu_o)_{r'} - (\mu_{o'})_r \tag{5.26b}$$

$$\ell_{r'} - \ell_r + (\alpha_{o'} - \alpha_o)\theta \geq (\mu_o)_{r'} - (\mu_{o'})_r \tag{5.26c}$$

since $\alpha_{i-1} - \alpha_i > 0$ *for all* i *and* $\theta \geq \bar{\theta}_{i-1}$ *for* $i \leq o$. *This yields*

$$\ell_r + \alpha_o\theta \leq \ell_{r'} + \alpha_{o'}\theta \tag{5.27}$$

again by complementary slackness which proves the result.

Remark 20 *We point out again that in the above proofs, the signs of* θ *and the* α_o*'s were not important. Only the fact that* $\alpha_{o-1} - \alpha_o > 0$ *and the fact that* θ *was contained in the interval* $[\bar{\theta}_{o-1}, \bar{\theta}_o]$ *was used.*

Figure 5.4 gives a graphical illustration of the equilibrium condition for the sample distribution in Figure 5.1. We note that the familiar balance Wardrop balance condition is maintained at the various transition points $\bar{\theta}_1, \bar{\theta}_2$, etc. This balance at the transition points is actually the way Leurent defines the equilibrium in [104]. In this particular case, we have assumed that the route groups are disjoint. Even so, we note that this illustration is fairly complicated involving a complex distribution and both positive and negative prices. For simpler illustrations with θ only positive and positive α_o's, consider the examples in the next section. For simpler examples with positive and negative values of θ as well as positive and negative values of α_o, consider the multi-modal example in Section 5.2.

Examples

We now compute the external-cost equilibrium for the parallel graph shown in Figure 5.5a and the graph shown in Figure 5.5b assuming the population has a uniform distribution on the value of time vs. the external factor with $\underline{\theta} = 0$, $\bar{\theta} = 10$, and a total mass of $M = 20$. The inverse cumulative distribution function is given by $\Theta(u) = 0.5u$.

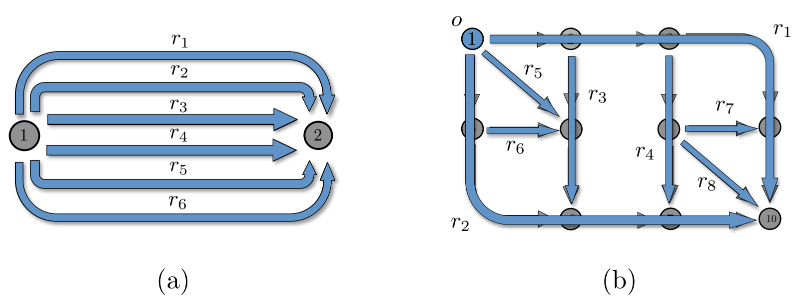

Figure 5.5: (a) Parallel graph and (b) more general graph with routes labeled.

For both graphs we assume linear edge latencies of the form

$$l_i(x) = x_i + 1 \tag{5.28}$$

We assume the following route groupings and prices.

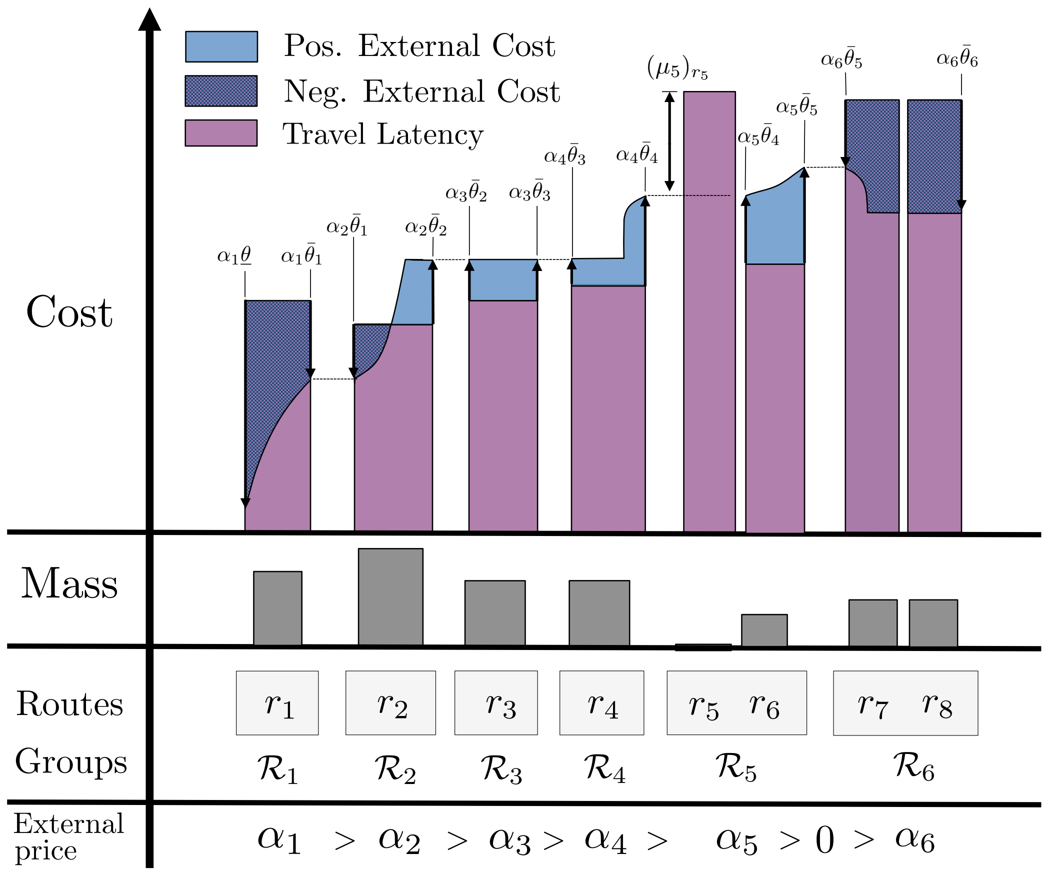

Figure 5.4: Visualization of the external-cost Wardrop equilibrium condition for the distribution shown in Figure 5.1 and the route subset structure given. Purple represents the travel latency for specific routes. Blue represents the perceived external cost which depends on the price of each routing option α_o and individual population members type θ. Note that the perceived external cost can be negative either because θ is negative (as is the case for $\mathcal{R}_1$ in the figure) or because α_o is negative (as is the case for $\mathcal{R}_6$ in the figure). The familiar Wardrop balance condition is preserved at the transitions between various routing options. We note that the distribution of mass on the various routes could be significantly more complicated for more complex route groupings.

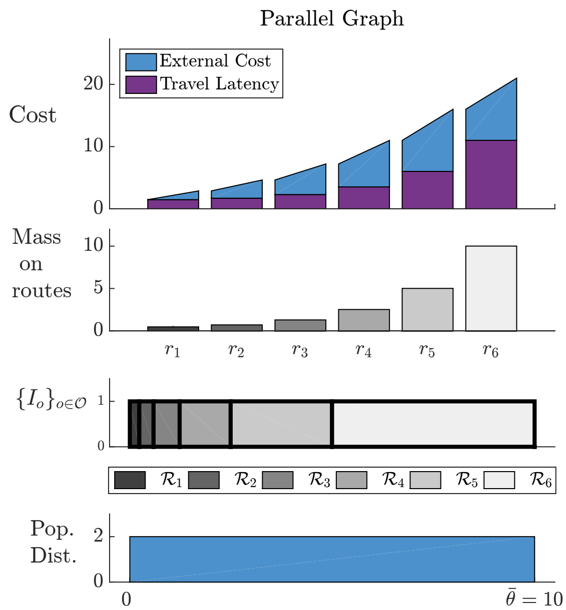

(a)

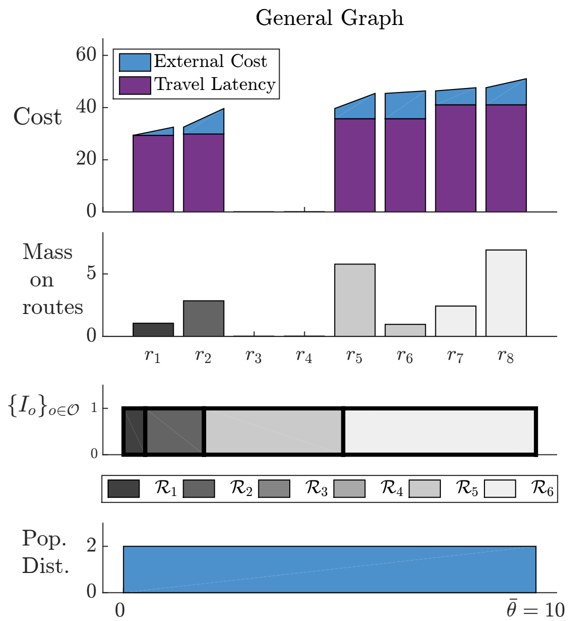

(b)

Figure 5.6: External-cost Wardrop equilibrium conditions for (a) a parallel graph (Figure 5.5a) and (b) a more general graph (Figure 5.5b) with a uniform distribution on $\theta \in [0, 10]$ and the route groupings and prices shown in Table 5.1.

	α_i	Parallel groups	General graph groups
$\mathcal{R}_1$	6	$\{r_1, r_2, r_3, r_4, r_5, r_6\}$	$\{r_1, r_2, r_3, r_4, r_5, r_6, r_7, r_8\}$
$\mathcal{R}_2$	5	$\{r_2, r_3, r_4, r_5, r_6\}$	$\{r_2, r_3, r_4, r_5, r_6, r_7, r_8\}$
$\mathcal{R}_3$	4	$\{r_3, r_4, r_5, r_6\}$	$\{r_3, r_4, r_5, r_6, r_7, r_8\}$
$\mathcal{R}_4$	3	$\{r_4, r_5, r_6\}$	$\{r_4, r_5, r_6, r_7, r_8\}$
$\mathcal{R}_5$	2	$\{r_5, r_6\}$	$\{r_5, r_6, r_7, r_8\}$
$\mathcal{R}_6$	1	$\{r_6\}$	$\{r_7, r_8\}$

Table 5.1: Route groupings and prices for Figure 5.6.

The resulting equilibria balance conditions, the choice functions $\{I_o\}_{o\in\mathcal{O}}$, and the route masses z are are shown in Figure 5.6. Note that in this case, the route groupings are nested with more expensive options having more routes and $\theta \geq 0$ for the whole population. When both of these things are true, clearly the mass choosing a particular route will choose the cheapest grouping in which that route is an option. This can be seen in Figure 5.6.

5.2 Applications: Information Pricing, Privacy, and Multi-Modal Routing

Classical routing and variable demand

It is straightforward to see that this framework reduces to the classical routing game whenever the population distribution is a delta function at zero. The variable demand routing game can also be thought of as a special case. Consider the route groupings,

$$\{\mathcal{R}_1, \mathcal{R}_2\} = \{\mathcal{R}, \emptyset\}, \tag{5.29a}$$

$$\{\alpha_1, \alpha_2\} = \{\alpha, 0\} \quad \text{with} \quad \alpha > 0 \tag{5.29b}$$

with $dF(\theta)$ supported on $\mathbb{R}_-$. The latency for taking a route in the empty set (not driving) is considered 0. The equilibrium condition for any route $r \in \mathcal{R}$ such that $z_r > 0$ is given by

$$\ell_r + \alpha\theta \leq 0 \tag{5.30}$$

Drivers with more negative values of θ are okay with longer travel times before they decide not to drive. The potential function is given by

$$F(z) = \sum_e \int_0^{x_e} l_e(u)\ du + \int_0^{m_1} \alpha_1 \Theta(u)\ du + \int_{m_1}^{m_1+m_2} \alpha_2 \Theta(u)\ du \tag{5.31a}$$

$$= \sum_e \int_0^{x_e} l_e(u)\ du + \int_0^{m_1} \alpha \Theta(u)\ du \tag{5.31b}$$

Here the demand curve is given by

$$d(\cdot) = -\frac{1}{\alpha}\Theta^{-1}(\cdot) \tag{5.32}$$

Traveler information systems market

One clear application of this framework would be modeling the market for traveler information systems. Various routing apps such as Google maps or Waze each provide users with a group of routes to choose from. Better apps would provide more routing options, shortcuts, etc. The price α_o would be the amount users pay for each service. This modeling framework would allow us to compute how users would decide between various services given their time-money tradeoff distribution.

Privacy

Another interesting application of this framework would be to model drivers' interest in their location privacy. Consider a scenario where users of a navigation service such as Google maps only receive congestion information about the network if they allow the navigation service to monitor their location. Any user who "opts out", decides not to share their location and not receive congestion information, takes the nominally shortest route, the shortest route without congestion. Any users that "opt in", decide to share their location information, receive information about all possible routes. The routing groups are $\mathcal{R}_{\text{out}} = \{r_1\}$ where r_1 is the shortest uncongested route and $\mathcal{R}_{\text{in}} = \mathcal{R}$. Members who opt out do not pay any additional cost ($\alpha_{\text{out}} = 0$) and members who opt in pay an additional cost of $\alpha_{\text{in}} > 0$ times θ. The parameter θ here represents each population member's value of privacy versus their value of travel time.

Multi-modal routing

This framework also provides a way to study the choice commuters make between different modes of transportation such as taking the subway or driving. We assume there are two commuting options with sets of routes $\mathcal{R}_1$ and $\mathcal{R}_2$. Depending on the commuting options these routes may be disjoint. For example, if commuters choose between driving or taking the subway, the subway routes will be separate from the driving routes and have different congestion effects. Congestion on driving routes will primarily result in increased travel time. Congestion on the subway routes could result in greater inconvenience, more crowded platforms, less available seats, etc. This framework is similar to the variable demand case except rather than choosing between driving and not driving at all, commuters choose between two different commuting options.

The population distribution $dF(\theta)$ is supported from the $\underline{\theta} \leq 0$ to $\bar{\theta} \geq 0$ and models the populations preference for one option over the other. If we set the price for option 1 as $\alpha_1 > 0$ and the price for option 2 as $\alpha_2 < 0$, then population mass below 0 has a preference for option 1 and population mass above 0 has a preference for option 2. We illustrate a sample preference distribution for two options in Figure 5.2. If in particular, we set $\alpha_1 = \frac{1}{2}$ and $\alpha_2 = -\frac{1}{2}$, then at $\bar{\theta}_1$ at equilibrium we have that

$$\min_{r \in \mathcal{R}_1} \ell_r(z) + \frac{1}{2}\bar{\theta}_1 = \min_{r \in \mathcal{R}_2} \ell_r(z) - \frac{1}{2}\bar{\theta}_1 \tag{5.33a}$$

$$\Rightarrow \qquad \bar{\theta}_1 = \min_{r \in \mathcal{R}_2} \ell_r(z) - \min_{r \in \mathcal{R}_1} \ell_r(z) \tag{5.33b}$$

In other words, the value of θ for each member of the population indicates how much faster option 1 must be than option 2 before they will switch to option 1.

We note that this framework can be used to model two different commuting options; however, more than two is problematic. Since θ is one dimensional, it can represent relative preference between two options; however, it would only make sense to compare three or more commuting options if there was a clear preference ordering for these options that *all* population members agreed upon. This is unlikely. It would be interesting to consider expanding this framework to a multi-dimensional preference parameter θ that compares multiple options.

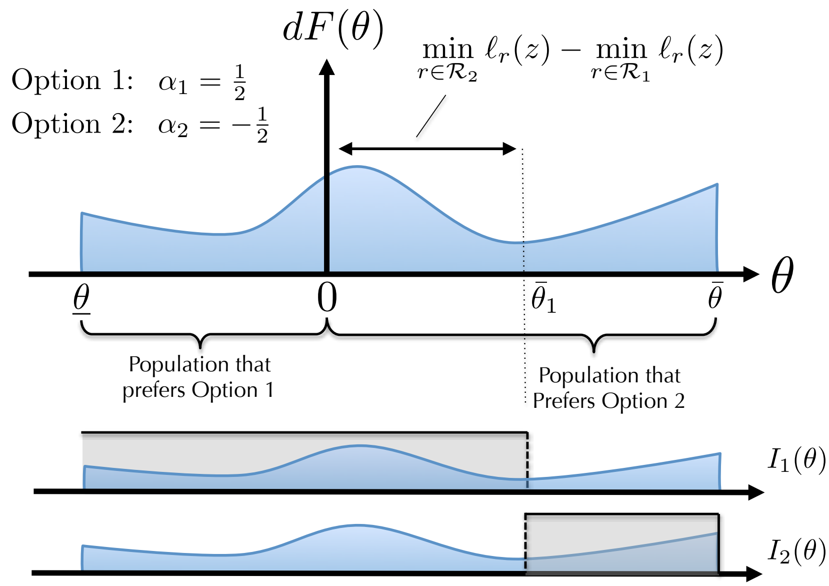

Figure 5.7: Preference distribution for two different commuting options. If the routing costs of each option are equal, population mass at $\theta < 0$ prefers Option 1 and population mass at $\theta > 0$ prefers Option 2.

The two commuting options in this framework can be varied and it is not required that the two sets of routes be disjoint. For example, we might seek to compare the demand for driving vs the demand for taxis. If the two options share routes, however, it is important that mass from either group have the same effect on congestion on the shared routes. If not, the game is no longer a potential game. For example, this framework would have problems comparing the demand for driving vs. the demand for carpool services since the portion of the population that chooses to carpool has less effect on road congestion than the portion that chooses to drive. Comparing driving vs. ride-sharing would be possible in situations where it was a reasonable to assume the average number of commuters in a regular car and a ride-sharing car were the same. Comparing driving vs. taking the bus would also be problematic if the cost of taking the bus depends on how congested the roads are along the bus route. Taking the bus does not add to congestion the same way driving a car does since the bus will be driving that route anyway regardless of how many commuters use it. If however, it was reasonable to assume that bus commuters do not consider congestion along the bus route but only how crowded the bus is as well as their personal preference for not driving, then this framework could be used.

To further illustrate the use of the model in this way, we consider a simple parallel network with two commuting options each with two routes (shown in Figure 5.8a). We assume a total mass of $M = 8$ and $\alpha_1 = \frac{1}{2}$ and $\alpha_2 = -\frac{1}{2}$. The population distribution is uniform from $\underline{\theta} = -4$ to $\bar{\theta} = 4$. We consider three different sets of latencies of the form

$$l_i(z_i) = a_i z_i + b_i \tag{5.34}$$

with coefficients shown in the table below.

	a_i				b_i			
Routes	r_1	r_2	r_3	r_4	r_1	r_2	r_3	r_4
Case 1	2	2	2	2	1	1	1	1
Case 2	1	8	1	1	0	0	8	8
Case 3	4	2	1	1	2	2	1	1

The resulting equilibria and balance conditions are illustrated in Figures 5.8b, 5.8c, and 5.8d. In case 1, the latencies on all the routes are the same and thus since the population distribution is symmetric about zero, population members with $\theta < 0$ choose option 1 and members with $\theta > 0$ choose option 2. In cases 2 and 3, this symmetry is broken and members who would prefer option 1 or 2 switch to the other option.

Ridesharing: MDP routing game

Finally, we note briefly that we could apply this framework in the ridesharing MDP routing game scenario to model how drivers make choices to drive for one ride-sharing company or another. The price for driving for a particular ridesharing company, α_o, would depend on the percentage of the fares that the drivers get to keep as well as any other overhead costs. The equilibrium mass distribution would reflect the fact that drivers would prefer to drive for a company that gives them more money but if too many drivers chose that company they would have to compete more for the specific rides they want. Note that this model would not currently capture the fact that rider demand also fluctuates with the number of drivers a particular company has. Extending the model to encompass this interaction would be interesting.

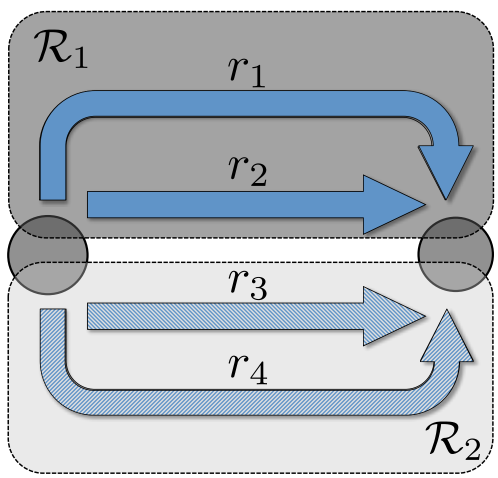

(a)

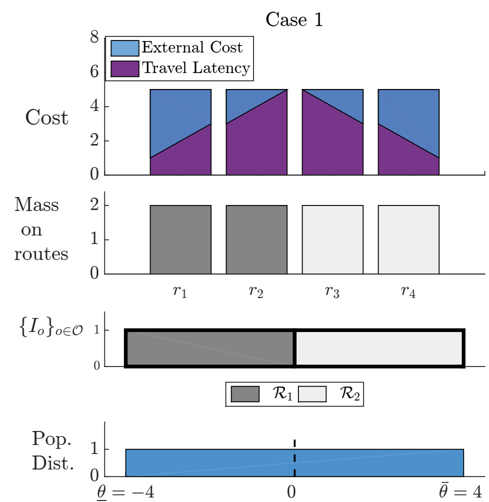

(b)

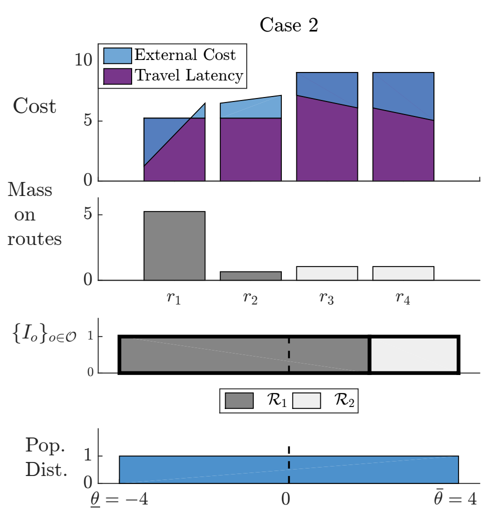

(c)

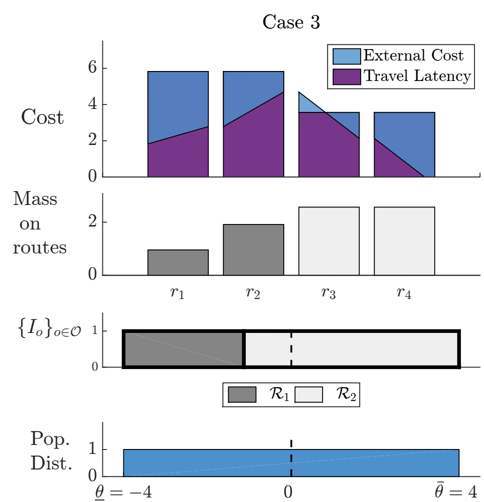

(d)

Figure 5.8: Equilibrium and balance conditions for the simple network shown in (a) for several different cases. (b) Case 1: the two routing options are symmetric so the population mass chooses their preferred mode of transportation and the population mass is split around $\theta = 0$. (c) Case 2: option 2 becomes more congested so some mass that prefers option 2 switches to option 1. (d) Case 3: option 1 becomes more congested so some mass switches from option 2 to option 1.

Chapter 6

Conclusion

In conclusion, we comment on modeling considerations that are important in applications, connections between the various aspects of our work, and future directions for research.

6.1 Modeling Considerations and Connections

A major part of the art of modeling is choosing a framework that is powerful enough to capture the desired phenomena but mathematically simple enough to solve. In our work, we hope to strike this balance. In the MDP routing game formulation in particular, we have provided a stochastic population game with a straight forward discrete time formulation that is fairly simple as compared to traditional mean-field games that require solving either coupled partial differential equations (in the general setting) or an optimal control problem (in the potential game setting). In the infinite-horizon MDP routing game, we have provided a low dimensional, tractable optimization problem that captures many of the important competitive features of various scenarios such as circling for parking or ridesharing. These simple types of models could prove useful in price setting or tolling problems providing a tractable equilibrium computation scheme to be used in the lower level of traditional tolling bilevel programs.

As previously mentioned at the end of Section 4.2, one of the weaknesses of the MDP routing game is the fundamentally discrete time nature of the formulation. Since all mass transitions from one node to another at each time step, the model is somewhat inconsistent with tradition routing game loss functions that model latency on an edge or waiting time at a node. It would be a useful direction of research to figure out *when* and *how much* simple models like the MDP routing game diverge from real world phenomena they are attempting to model. In situations where the correspondence is close, the simple MDP routing game model could be used, and in situations where the model does not capture the phenomena well, one could switch to a more complex, more faithful approach such as a full mean-field game model.

As mentioned several times, the infinite horizon MDP routing game model of circling for parking presented in Section 4.4 could be incorporated into the queue-routing game. The strategy space would involve choosing a parking area, an entrance node, a route to that node, and a circling policy to be applied within the parking area. The edge latencies within the parking area would depend on both thru traffic congestion as well as circling congestion and the potential function would be a combination of (3.15) and (4.28) (with the edge integral terms only appearing once).

The external-cost framework for value of time vs. money could be incorporated into the queue-routing game in order to provide a more complex look at how parking prices would affect the equilibrium flow distributions. In this case, each parking area would correspond to an option and we would be setting the price of that option to the price of parking in that area, $\frac{C^p_{\mathrm{p}}}{\gamma^p}$. The distribution of θ would represent the population's time-money tradeoff. The external-cost framework could also be used to model a parking population's preference for one parking area over another. In the case where there were only two parking areas the comparison framework detailed in Section 5.2 for multi-modal routing could be used. This framework allows us to compare arbitrary population preference for one of two parking areas. If we could assume that the population chooses parking areas based on one objective factor alone, proximity to the attraction for example, we could use the external-cost framework to rank more than two parking areas. Here the price of each option would be R^{ap} and θ would represent the population's proximity vs. travel time tradeoff.

It would be interesting also to see how the continuous distribution external-factor consideration could be combined with the MDP routing game. At a basic level, the external-factor framework could be used to compare drivers choosing between driving for two different ridesharing services such as Uber or Lyft. It would also be interesting to consider situations where two of the different loss types (action losses, edge losses, or node losses) are fundamentally different from each other and θ represents the tradeoff between the various types of losses. With our current formulation this would only be possible if one type of loss was constant. An example might be in the traffic circling scenario where edge or action losses representing the inconvenience of making specific turns at an intersection (left turns are more inconvenient than right turns, etc) are constant and where the node losses representing the probability of not getting a parking space are congestion dependent. θ in this example would represent the tradeoff between ease of driving vs. desire to find a parking space.

6.2 Future Directions

There are multiple future research directions to be explored as well. For the queue-routing game, a more realistic model of circling, such as the MDP routing game model, should be incorporated. It would also be very interesting to incorporate the model into a bilevel program to design parking prices as well as to compare the models predictions to real world data in various scenarios. In particular, it would be interesting to compare scenarios where drivers have more or less information about parking availability (through phone apps, etc) to see if the model accurate predicts how parking traffic will shift as the prices are adjusted.

The MDP routing game framework could be incorporated into a bilevel programming framework to design surge prices and trip rates. It would be a interesting to compare the results of the model with actual driving patterns of ridesharing drivers. Along these same lines from a more theoretical perspective, it would be very interesting to compare long term optimization behavior (the MDP Wardrop equilibrium) with another equilibrium concept where agents optimize their immediate rewards or losses. In some cases, this immediate reward optimization, which we might call a *myopic Wardrop equilibrium* could be more consistent with drivers' actual behavior. Designing a way to compute such an equilibrium as well as to compute a combined equilibrium where some agents optimize over the entire time horizon and some only care about immediate rewards could have wide spread applications. It would also be interesting to study the discounted-loss MDP routing game case and draw connections with the myopic equilibrium as the limit of the discounted problem as the discount factor shrinks to zero. Lower bounds on the price of anarchy for MDP routing games should be studied. Inspiration from price of anarchy results in MDP routing games could also lead to interesting results in the general mean-field game context. It would also be very interesting to explore Braess paradox further in the MDP routing game setting. In some ways, studying Braess paradox in this setting is more natural than the classical routing game setting since any results would depend only on the connectivity structure of the graph (the incidence matrix G, the space of cycles, etc) as opposed to also depending on where the mass enters and exits the network (the source-sink vector S).

Finally, jt would be interesting to revisit Marcotte's variational inequality formulation of the external-cost framework and extend the results about θ being supported above and below zero to this context. Making this extension would allow for direct comparison of travel options that have more than one type of congestion effect. It would also be very interesting to extend this comparison framework to more than two decision criteria, i.e. allowing θ to be multi-dimensional. This would provide a framework to compare transportation options on multiple criteria or a framework to directly compare multiple transportation options pairwise.

As intelligent transportation systems, both advanced traffic information systems and autonomous driving systems, become more ubiquitous, complex models of competition such as the ones presented in this work will become more and more useful. Given the right models of competition, urban planners will be able to design new systems that leverage the incentives users face to provide efficient, robust transportation and maximize the quality of life in urban areas.

Bibliography

[1] G. D. Squires, *Urban sprawl: Causes, consequences, & policy responses.* The Urban Insitute, 2002.

[2] T. J. Nechyba and R. P. Walsh, "Urban sprawl," *The Journal of Economic Perspectives*, vol. 18, no. 4, pp. 177–200, 2004.

[3] J. K. Brueckner, "Urban sprawl: Lessons from urban economics," *Brookings-Wharton papers on urban affairs*, vol. 2001, no. 1, pp. 65–97, 2001.

[4] R. Camagni, M. C. Gibelli, and P. Rigamonti, "Urban mobility and urban form: the social and environmental costs of different patterns of urban expansion," *Ecological economics*, vol. 40, no. 2, pp. 199–216, 2002.

[5] D. Schrank, B. Eisele, T. Lomax, and J. Bak, "2015 urban mobility scorecard," Texas A&M Transportation Institute and INRIX, Tech. Rep., 2015.

[6] J. I. Levy, J. J. Buonocore, and K. von Stackelberg, "Evaluation of the public health impacts of traffic congestion: a health risk assessment," *Environ Health*, vol. 9, no. 1, p. 65, 2010.

[7] K. Zhang and S. Batterman, "Air pollution and health risks due to vehicle traffic," *Science of The Total Environment*, vol. 450–451, pp. 307–316, Apr 2013.

[8] E. Johnson, "6 parking apps that save you time and money," December 2014, [Online; posted 18-December-2014]. [Online]. Available: http://www.forbes.com/sites/emmajohnson/2014/12/18/5-parking-apps-that-help-you-save-time-and-money/

[9] B. Jinks, "Icahn says lyft worth more than $2 billion based on uber value," May 2015, [Online; posted 19-May-2015]. [Online]. Available: https://www.bloomberg.com/news/articles/2015-05-19/icahn-says-lyft-worth-more-than-2-billion-based-on-uber-value

[10] M. Chafkin, "Autonomous cars: The future of mobility," September 2016, [Online; posted 27-September-2016]. [Online]. Available: http://www.forbes.com/sites/danielnewman/2016/09/27/autonomous-cars-the-future-of-mobility/

[11] T. T. Team, "All tesla cars being produced now have full self-driving hardward," October 2016, [Online; posted 19-October-2016]. [Online]. Available: https://www.tesla.com/blog/all-tesla-cars-being-produced-now-have-full-self-driving-hardware

[12] M. Chafkin, "Uber's first self-driving fleet arrives in pittsburgh this month," August 2016, [Online; posted 18-August-2016]. [Online]. Available: https://www.bloomberg.com/news/features/2016-08-18/uber-s-first-self-driving-fleet-arrives-in-pittsburgh-this-month-is06r7on

[13] J. G. Wardrop, "Road paper. some theoretical aspects of road traffic research." in *ICE Proceedings: Engineering Divisions*, vol. 1, no. 3. Thomas Telford, 1952, pp. 325–362.

[14] J. Lofberg, "Yalmip : A toolbox for modeling and optimization in MATLAB," in *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004. [Online]. Available: http://control.ee.ethz.ch/~{}joloef/yalmip.php

[15] M. Patriksson, *The traffic assignment problem: models and methods.* Courier Dover Publications, 2015.

[16] W. H. Sandholm, "Potential games with continuous player sets," *Journal of Economic Theory*, vol. 97, no. 1, pp. 81–108, 2001.

[17] M. Beckmann, C. McGuire, and C. B. Winsten, "Studies in the economics of transportation," Tech. Rep., 1956.

[18] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, "Network flows: theory, algorithms, and applications," 1993.

[19] G. F. Newell, *Traffic flow on transportation networks*, 1980, no. Monograph.

[20] D. Kulash, "A transportation equilibrium model," Urban Institute Paper, pp. 708–45, 1971.

[21] M. R. Wigan, "Benefit assessment for network traffic models and application to road pricing," 1971.

[22] T. Roughgarden, *Selfish routing and the price of anarchy.* MIT press Cambridge, 2005, vol. 174.

[23] ——, "The price of anarchy is independent of the network topology," *Journal of Computer and System Sciences*, vol. 67, no. 2, pp. 341–364, 2003.

[24] T. Roughgarden and É. Tardos, "Bounding the inefficiency of equilibria in nonatomic congestion games," *Games and Economic Behavior*, vol. 47, no. 2, pp. 389–403, 2004.

[25] M. J. Smith, "The existence, uniqueness and stability of traffic equilibria," *Transportation Research Part B: Methodological*, vol. 13, no. 4, pp. 295–304, 1979.

[26] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses, "Selfish routing in capacitated networks," *Mathematics of Operations Research*, vol. 29, no. 4, pp. 961–976, 2004.

[27] T. Roughgarden and É. Tardos, "How bad is selfish routing?" *Journal of the ACM (JACM)*, vol. 49, no. 2, pp. 236–259, 2002.

[28] S. P. Anderson and A. de Palma, "The economics of pricing parking," *Journal of Urban Economics*, vol. 55, no. 1, pp. 1 – 20, 2004.

[29] D. Shoup, "Cruising for parking," *Transport Policy*, vol. 13, pp. 479–486, 2006.

[30] T. Brown, "Five cities with congestion pricing," 2011. [Online]. Available: http://thisbigcity.net/five-cities-with-congestion-pricing/

[31] G. Topham, "Congestion charges around the world," 2011. [Online]. Available: http://www.theguardian.com/world/2011/dec/08/congestion-charges-around-the-world

[32] B. Schaller, "New york city's congestion pricing experience and implications for road pricing acceptance in the united states," *Transport Policy*, vol. 17, no. 4, pp. 266 – 273, 2010.

[33] D. King, M. Manville, and D. Shoup, "The political calculus of congestion pricing," *Transport Policy*, vol. 14, no. 2, pp. 111 – 123, 2007.

[34] W. Harrington, A. J. Krupnick, and A. Alberini, "Overcoming public aversion to congestion pricing," *Transportation Research Part A: Policy and Practice*, vol. 35, no. 2, pp. 87–105, 2001.

[35] San Francisco parking pilot evaluation, SFpark, conducted by San Francisco Municipal Transportation Authority, [Online:] http://sfpark.org/.

[36] Washington D.C. performance-based parking pilot conduced by the District Department of Transportation, [Online:] http://ddot.dc.gov/service/performance-based-parking-pilots.

[37] Seattle performance-based parking pricing study conducted by the Seattle Department of Transportation, [Online:] http://www.seattle.gov/transportation/parking/performancepricing.htm.

[38] SFMTA, "Sfpark: Pilot project evaluation," San Francisco Municipal Transportation Agency, Tech. Rep., 2014.

[39] D. Calderone, E. Mazumdar, L. J. Ratliff, and S. S. Sastry, "Understanding the impact of parking on urban mobility via routing games on queue-flow networks," in *Decision and Control (CDC), 2016 IEEE 55th Conference on*. IEEE, 2016, pp. 7605–7610.

[40] P. Naor, "The regulation of queue size by levying tolls," *Econometrica: journal of the Econometric Society*, pp. 15–24, 1969.

[41] National Research Council, *Special Report 209: Highway Capacity Manual.* Transportation Research Board, 1985.

[42] L. S. Shapley, "Stochastic games," *Proceedings of the national academy of sciences*, vol. 39, no. 10, pp. 1095–1100, 1953.

[43] M. J. Sobel *et al.*, "Noncooperative stochastic games," *The Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 1930–1935, 1971.

[44] J.-F. Mertens and A. Neyman, "Stochastic games," *International Journal of Game Theory*, vol. 10, no. 2, pp. 53–66, 1981.

[45] A. Condon, "The complexity of stochastic games," *Information and Computation*, vol. 96, no. 2, pp. 203–224, 1992.

[46] A. Neyman and S. Sorin, *Stochastic games and applications.* Springer Science & Business Media, 2003, vol. 570.

[47] J. A. Potters, T. Raghavan, and S. H. Tijs, "Pure equilibrium strategies for stochastic games via potential functions," in *Advances in Dynamic Games and their Applications.* Springer, 2009, pp. 1–12.

[48] D. González-Sánchez and O. Hernández-Lerma, "A survey of static and dynamic potential games," *Science China Mathematics*, vol. 59, no. 11, pp. 2075–2102, 2016.

[49] G. Y. Weintraub, C. L. Benkard, and B. Van Roy, "Oblivious equilibrium: A mean field approximation for large-scale dynamic games," in *NIPS*, 2005, pp. 1489–1496.

[50] S. Adlakha, R. Johari, G. Weintraub, and A. Goldsmith, "Oblivious equilibrium for large-scale stochastic games with unbounded costs," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on.* IEEE, 2008, pp. 5531–5538.

[51] G. Y. Weintraub, C. L. Benkard, P. Jeziorski, and B. Van Roy, "Nonstationary oblivious equilibrium," *Manuscript, Columbia University*, pp. 1375–1411, 2008.

[52] G. Y. Weintraub, C. L. Benkard, and B. Van Roy, "Computational methods for oblivious equilibrium," *Operations research*, vol. 58, no. 4-part-2, pp. 1247–1265, 2010.

[53] B. Jovanovic and R. W. Rosenthal, "Anonymous sequential games," *Journal of Mathematical Economics*, vol. 17, no. 1, pp. 77–87, 1988.

[54] J. Bergin, D. Bernhardt, *et al.*, "Anonymous sequential games with general state space," Tech. Rep., 1991.

[55] J. Bergin and D. Bernhardt, "Anonymous sequential games with aggregate uncertainty," *Journal of Mathematical Economics*, vol. 21, no. 6, pp. 543–562, 1992.

[56] ——, "Anonymous sequential games: existence and characterization of equilibria," *Economic Theory*, vol. 5, no. 3, pp. 461–489, 1995.

[57] P. Wiecek and E. Altman, "Stationary anonymous sequential games with undiscounted rewards," *Journal of optimization theory and applications*, vol. 166, no. 2, pp. 686–710, 2015.

[58] P. Wiecek, E. Altman, and Y. Hayel, "An anonymous sequential game approach for battery state dependent power control," in *International Conference on Network Control and Optimization*. Springer, 2009, pp. 121–136.

[59] ——, "Stochastic state dependent population games in wireless communication," *IEEE Transactions on Automatic Control*, vol. 56, no. 3, pp. 492–505, 2011.

[60] J.-M. Lasry and P.-L. Lions, "Jeux à champ moyen. i–le cas stationnaire," *Comptes Rendus Mathématique*, vol. 343, no. 9, pp. 619–625, 2006.

[61] ——, "Jeux à champ moyen. ii–horizon fini et contrôle optimal," *Comptes Rendus Mathématique*, vol. 343, no. 10, pp. 679–684, 2006.

[62] ——, "Mean field games," *Japanese Journal of Mathematics*, vol. 2, no. 1, pp. 229–260, 2007.

[63] M. Huang, R. P. Malhamé, P. E. Caines, *et al.*, "Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the nash certainty equivalence principle," *Communications in Information & Systems*, vol. 6, no. 3, pp. 221–252, 2006.

[64] M. Huang, P. E. Caines, and R. P. Malhamé, "Large-population cost-coupled lqg problems with nonuniform agents: individual-mass behavior and decentralized *varepsilon*-nash equilibria," *IEEE transactions on automatic control*, vol. 52, no. 9, pp. 1560–1571, 2007.

[65] A. Lachapelle, J. Salomon, and G. Turinici, "Computation of mean field equilibria in economics," *Mathematical Models and Methods in Applied Sciences*, vol. 20, no. 04, pp. 567–588, 2010.

[66] Y. Achdou and I. Capuzzo-Dolcetta, "Mean field games: Numerical methods," *SIAM Journal on Numerical Analysis*, vol. 48, no. 3, pp. 1136–1162, 2010.

[67] Y. Achdou, F. Camilli, and I. Capuzzo-Dolcetta, "Mean field games: numerical methods for the planning problem," *SIAM Journal on Control and Optimization*, vol. 50, no. 1, pp. 77–109, 2012.

[68] ——, "Mean field games: convergence of a finite difference method," *SIAM Journal on Numerical Analysis*, vol. 51, no. 5, pp. 2585–2612, 2013.

[69] O. Guéant, "New numerical methods for mean field games with quadratic costs." *Networks & Heterogeneous Media*, vol. 7, no. 2, 2012.

[70] S. Cacace, F. Camilli, and C. Marchi, "A numerical method for mean field games on networks," *ESAIM: Mathematical Modelling and Numerical Analysis*, vol. 51, no. 1, pp. 63–88, 2017.

[71] F. Mériaux, V. Varma, and S. Lasaulce, "Mean field energy games in wireless networks," in *Signals, Systems and Computers (ASILOMAR), 2012 Conference Record of the Forty Sixth Asilomar Conference on*. IEEE, 2012, pp. 671–675.

[72] F. Mériaux, S. Lasaulce, and H. Tembine, "Stochastic differential games and energy-efficient power control," *Dynamic Games and Applications*, vol. 3, no. 1, pp. 3–23, 2013.

[73] H. Tembine, P. Vilanova, M. Assaad, and M. Debbah, "Mean field stochastic games for sinr-based medium access control," in *Proceedings of the 5th International ICST Conference on Performance Evaluation Methodologies and Tools*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2011, pp. 572–581.

[74] H. Tembine, E. Altman, R. El-Azouzi, and Y. Hayel, "Evolutionary games in wireless networks," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 40, no. 3, pp. 634–646, 2010.

[75] H. Tembine *et al.*, "Energy-constrained mean field games in wireless networks," *Strategic Behavior and the Environment*, vol. 4, no. 1, pp. 99–123, 2014.

[76] M. Manjrekar, V. Ramaswamy, and S. Shakkottai, "A mean field game approach to scheduling in cellular systems," in *INFOCOM, 2014 Proceedings IEEE*. IEEE, 2014, pp. 1554–1562.

[77] M. Aziz and P. E. Caines, "Computational investigations of decentralized cellular network optimization via mean field control," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*. IEEE, 2014, pp. 5560–5567.

[78] O. Guéant, J. M. Lasry, P. L. Lions, *et al.*, "Mean field games and oil production," Tech. Rep., 2010.

[79] R. Couillet, S. M. Perlaza, H. Tembine, and M. Debbah, "A mean field game analysis of electric vehicles in the smart grid," in *Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on*. IEEE, 2012, pp. 79–84.

[80] M. Kamgarpour and H. Tembine, "A bayesian mean field game approach to supply demand analysis of the smart grid," in *Communications and Networking (BlackSeaCom), 2013 First International Black Sea Conference on.* IEEE, 2013, pp. 211–215.

[81] F. Bagagiolo and D. Bauso, "Mean-field games and dynamic demand management in power grids," *Dynamic Games and Applications*, vol. 4, no. 2, pp. 155–176, 2014.

[82] A. T. Siwe and H. Tembine, "Mean-field-type games on airline networks and airport queues: Braess paradox, its negation, and crowd effect," in *Systems, Signals & Devices (SSD), 2016 13th International Multi-Conference on.* IEEE, 2016, pp. 595–600.

[83] D. A. Gomes, J. Mohr, and R. R. Souza, "Discrete time, finite state space mean field games," *Journal de mathématiques pures et appliquées*, vol. 93, no. 3, pp. 308–328, 2010.

[84] O. Guéant, "From infinity to one: The reduction of some mean field games to a global control problem," *arXiv preprint arXiv:1110.3441*, 2011.

[85] O. Gueant, "Existence and uniqueness result for mean field games with congestion effect on graphs," *Applied Mathematics & Optimization*, vol. 72, no. 2, pp. 291–303, 2015.

[86] D. Gomes and V. Voskanyan, "Extended mean field games-formulation, existence, uniqueness and examples," *Preprint*, 2013.

[87] D. A. Gomes, S. Patrizi, and V. Voskanyan, "On the existence of classical solutions for stationary extended mean field games," *Nonlinear Analysis: Theory, Methods & Applications*, vol. 99, pp. 49–79, 2014.

[88] M. Huang, P. E. Caines, and R. P. Malhamé, "Social optima in mean field lqg control: centralized and decentralized strategies," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1736–1751, 2012.

[89] H. Yin, P. G. Mehta, S. P. Meyn, and U. V. Shanbhag, "On the efficiency of equilibria in mean-field oscillator games," *Dynamic Games and Applications*, vol. 4, no. 2, pp. 177–207, 2014.

[90] M. Balandat and C. J. Tomlin, "On efficiency in mean field differential games," in *American Control Conference (ACC), 2013.* IEEE, 2013, pp. 2527–2532.

[91] D. Calderone and S. S. Sastry, "Markov decision process routing games," in *Proceedings of the 8th International Conference on Cyber-Physical Systems.* ACM, 2017, pp. 273–279.

[92] A. S. Manne, "Linear programming and sequential decisions," *Management Science*, vol. 6, no. 3, pp. 259–267, 1960.

[93] J. L. Nazareth and R. B. Kulkarni, "Linear programming formulations of markov decision processes," *Operations research letters*, vol. 5, no. 1, pp. 13–16, 1986.

[94] D. Bello and G. Riano, "Linear programming solvers for markov decision processes," in *Systems and Information Engineering Design Symposium*, 2006, pp. 90–95.

[95] R. B. Dial, "A model and algorithm for multicriteria route-mode choice," *Transportation Research Part B: Methodological*, vol. 13, no. 4, pp. 311–316, 1979.

[96] R. E. Quandt, "A probabilistic abstract mode model," *Studies in Travel Demand VIII, Mathematica, Inc., Princeton, New Jersey*, pp. 127–149, 1967.

[97] M. Schneider, "Access and land development," *Highway Research Board Special Report*, no. 97, 1968.

[98] S. Dafermos, "The general multimodal network equilibrium problem with elastic demand," *Networks*, vol. 12, no. 1, pp. 57–72, 1982.

[99] C. F. Daganzo, "Stochastic network equilibrium with multiple vehicle types and asymmetric, indefinite link cost jacobians," *Transportation Science*, vol. 17, no. 3, pp. 282–300, 1983.

[100] A. Nagurney, "A multiclass, multicriteria traffic network equilibrium model," *Mathematical and Computer Modelling*, vol. 32, no. 3-4, pp. 393–411, 2000.

[101] A. Nagurney and J. Dong, "A multiclass, multicriteria traffic network equilibrium model with elastic demand," *Transportation Research Part B: Methodological*, vol. 36, no. 5, pp. 445–469, 2002.

[102] S. jie LI and G. ya CHEN, "Multiclass, two-criteria traffic network equilibrium models and vector variational inequalities," *Systems Engineering - Theory Practice*, vol. 28, no. 1, pp. 141 – 145, 2008.

[103] H. Yang and H.-J. Huang, "The multi-class, multi-criteria traffic network equilibrium and systems optimum problem," *Transportation Research Part B: Methodological*, vol. 38, no. 1, pp. 1 – 15, 2004.

[104] F. Leurent, "Cost versus time equilibrium over a network," *European Journal of Operational Research*, vol. 71, no. 2, pp. 205–221, 1993.

[105] ——, "The theory and practice of a dual criteria assignment model with a continuously distributed value-of-time," in *ISTTT*. Pergamon, 1996, pp. pp–455.

[106] R. B. Dial, "Bicriterion traffic assignment: basic theory and elementary algorithms," *Transportation Science*, vol. 30, no. 2, pp. 93–111, 1996.

[107] ——, "Bicriterion traffic assignment: efficient algorithms plus examples," *Transportation Research Part B: Methodological*, vol. 31, no. 5, pp. 357–379, 1997.

[108] P. Marcotte and D. Zhu, "An efficient algorithm for a bicriterion traffic assignment problem," in *Advanced methods in transportation analysis*. Springer, 1996, pp. 63–73.

[109] P. Marcotte, S. Nguyen, and K. Tanguay, "Implementation of an efficient algorithm for the multiclass traffic assignment problem," in *Transportation and Tra c Theory. Proceedings of the 13th International Symposium on Transportation and Tra c Theory, Pergamon, Oxford*. Citeseer, 1996, pp. 217–236.

[110] P. Marcotte and D. L. Zhu, "Equilibria with infinitely many differentiated classes of customers," 1997.

[111] P. Marcotte, "Reformulations of a bicriterion equilibrium model," in *Reformulation: Nonsmooth, Piecewise Smooth, Semismooth and Smoothing Methods*. Springer, 1998, pp. 269–291.

[112] P. Marcotte and M. Patriksson, "Traffic equilibrium," *Handbooks in Operations Research and Management Science*, vol. 14, pp. 623–713, 2007.

[113] H. Yang and X. Zhang, "Multiclass network toll design problem with social and spatial equity constraints," *Journal of Transportation Engineering*, vol. 128, no. 5, pp. 420–428, 2002.

[114] M. Florian, "Network equilibrium models for analyzing toll highways," in *Mathematical and Computational Models for Congestion Charging*. Springer, 2006, pp. 105–115.

[115] D. Han and H. Yang, "The multi-class, multi-criterion traffic equilibrium and the efficiency of congestion pricing," *Transportation Research Part E: Logistics and Transportation Review*, vol. 44, no. 5, pp. 753–773, 2008.

[116] R. B. Dial, "Network-optimized road pricing: Part i: A parable and a model," *Operations Research*, vol. 47, no. 1, pp. 54–64, 1999.

[117] ——, "Network-optimized road pricing: Part ii: Algorithms and examples," *Operations Research*, vol. 47, no. 2, pp. 327–336, 1999.

[118] C.-C. Lu, X. Zhou, and H. Mahmassani, "Variable toll pricing and heterogeneous users: Model and solution algorithm for bicriterion dynamic traffic assignment problem," *Transportation Research Record: Journal of the Transportation Research Board*, no. 1964, pp. 19–26, 2006.

[119] C.-C. J. Lu, "Multi-criterion dynamic traffic assignment models and algorithms for road pricing applications with heterogeneous users," Ph.D. dissertation, 2007.

[120] C.-C. Lu, H. S. Mahmassani, and X. Zhou, "A bi-criterion dynamic user equilibrium traffic assignment model and solution algorithm for evaluating dynamic road pricing strategies," *Transportation Research Part C: Emerging Technologies*, vol. 16, no. 4, pp. 371–389, 2008.

[121] P. Marcotte and D. Zhu, "Existence and computation of optimal tolls in multiclass network equilibrium problems," *Operations Research Letters*, vol. 37, no. 3, pp. 211–214, 2009.

[122] W.-X. Wu and H.-J. Huang, "Finding anonymous tolls to realize target flow pattern in networks with continuously distributed value of time," *Transportation Research Part B: Methodological*, vol. 65, pp. 31–46, 2014.

[123] J. Y. Wang, A. Raith, and M. Ehrgott, "Tolling analysis with bi-objective traffic assignment," in *Multiple Criteria Decision Making for Sustainable Energy and Transportation Systems*. Springer, 2010, pp. 117–129.

[124] X. Wang, H.-J. Huang, and L.-J. Tian, "Bi-criteria system optimum with fixed demand and continuously distributed value of time," in *CICTP 2012: Multimodal Transportation Systems—Convenient, Safe, Cost-Effective, Efficient*, 2012, pp. 90–98.

[125] X. Wang and H.-J. Huang, "Bi-criteria system optimum traffic assignment in networks with continuous value of time," *PROMET-Traffic&Transportation*, vol. 25, no. 2, pp. 119–125, 2013.

[126] J. Y. Wang and M. Ehrgott, "Modelling route choice behaviour in a tolled road network with a time surplus maximisation bi-objective user equilibrium model," *Procedia-Social and Behavioral Sciences*, vol. 80, pp. 266–288, 2013.

[127] ——, "Bi-objective network equilibrium, traffic assignment and road pricing," in *Operations Research Proceedings 2012*. Springer, 2014, pp. 563–568.

[128] A. Raith, J. Y. Wang, M. Ehrgott, and S. A. Mitchell, "Solving multi-objective traffic assignment," *Annals of Operations Research*, vol. 222, no. 1, pp. 483–516, 2014.

[129] A. Raith *et al.*, "Multiobjective routing and transportation problems," Ph.D. dissertation, ResearchSpace@ Auckland, 2009.

[130] M. Kearns, M. Pai, A. Roth, and J. Ullman, "Mechanism design in large games: Incentives and privacy," in *Proceedings of the 5th conference on Innovations in theoretical computer science*. ACM, 2014, pp. 403–410.

[131] R. Dong, W. Krichene, A. M. Bayen, and S. S. Sastry, "Differential privacy of populations in routing games," in *Decision and Control (CDC), 2015 IEEE 54th Annual Conference on*. IEEE, 2015, pp. 2798–2803.

[132] H. Yang, H. Y. Kong, and Q. Meng, "Value-of-time distributions and competitive bus services," *Transportation Research Part E: Logistics and Transportation Review*, vol. 37, no. 6, pp. 411–424, 2001.