



ROUTING GAMES, INFORMATION AND LEARNING

M.Sc. ECONOMICS AND FINANCE, MAJOR ECONOMICS

Chair of Mathematical Statistics

2021-2022

Supervisor:

Prof. Marco Scarsini

Co-supervisor:

Prof. Fabio Antonelli

Submitted by:

Francesco Giordano

733021

Ai miei genitori, che hanno reso tutto ciò possibile.

Routing Games, Information and Learning

LUISS Guido Carli

Department of Economics and Finance

by Francesco Giordano

June 10, 2022

Abstract

In this work we study some new aspects of selfish routing games. First, we survey the routing problem under complete information, in which an anonymous mass of traffic inflow has to be routed from a set of sources to a set of destination. Then, we study how uncertainty can affect the routing behavior: we consider a routing game with an unknown underlying state, in which players route some traffic demand in order to minimize their expected costs. Then, we consider the case where these games are repeatedly played with a random demand in each stage game, and clarify that this model opens to the possibility of learning the unknown state. Two types of social learning are relevant in non-atomic routing games: *strong learning*, when the true underlying state is identified, and *weak learning*, when players learn how to play as if the true state were known. Afterwards, we briefly discuss the central idea of routing under private information. Finally, we present the original contribution: considering strictly increasing, continuous and unbounded edge costs, we extend the results of strong learning to infinite capacity multi-commodity instances and to capacitated instances with one or more commodities. We define when a network state is identifiable, and show the network conditions sufficient to achieve almost sure learning. In capacitated instances, a sufficient condition to achieve learning is that the network respects a condition called *weak capacity conservation*, that is independent on the network topology. In infinite capacity instances, a sufficient condition is that the network is *series-parallel* in each commodity sub-network.

Keywords: Selfish Routing, Non-atomic Games, Incomplete Information, Social Learning, Congestion Games, Dynamic Games, Learning Failures, SP Networks, Capacity Conservation.

Contents

1	Introduction	1
1.1	Contribution	2
2	Literature review	4
3	Deterministic Routing Games	6
3.1	Non-Atomic Routing games	7
3.1.1	Examples: Pigou network and Braess network	8
3.2	Congestion games and Potential games	10
3.3	Equilibria and Optimality	11
3.4	Price of Anarchy	14
3.5	Atomic Routing games and convergence results	16
4	Routing Games under Uncertainty	19
4.1	Introducing uncertainty	20
4.2	Dynamic Routing games with unknown network state	21
4.2.1	Social Learning	22
4.3	Private Information	23
5	Social Learning in non-atomic Routing Games	25
5.1	Some structural properties	26
5.1.1	Load upper bound and capacity conditions	26
5.1.2	Feasibility condition and identifiability	27
5.2	Learning failures	29
5.2.1	Capacity conservation	33
5.2.2	SP networks and generalization	37
5.3	Learning in capacity constrained networks	38
5.4	Learning in networks with infinite capacities	40
6	Conclusion	42
6.1	Future Research	43
	Bibliography	43
	Appendix	48
	Proofs	49
	List of Symbols	56
	Acknowledgment	57

Chapter 1

Introduction

This thesis is aimed at presenting some modern aspects of *selfish routing* models. *Routing games* are the mathematical tools used to describe and model realistic instances of selfish routing, such as telecommunication and road traffic. We present the central concepts and results and we extend some of the existing results on social learning.

Routing models are substantially divided into two the categories: *non-atomic* models and *atomic* models. Atomic models consider a finite set of players, each of them controlling a non-negligible amount of the total traffic inflow; on the other side, non-atomic games are played by an anonymous mass of players and each player control a negligible amount of the total traffic.

This work will mainly cover results and contributions arising in the non-atomic setting, in which the set of players is defined in terms of infinitely divisible masses of traffic. The non-atomic routing model owes its popularity to at least two practical reasons: first, it is a good approximation of an atomic setup in which the number of players is large and the weight of each player is small; secondly, the assumption of an infinitely divisible mass of traffic demand helps the derivation of some analytical results. Nevertheless, the connection between the atomic and the non-atomic setup does not find its roots only in practical reasons: the formal derivation of convergence results from atomic to non-atomic congestion games is indeed a hot topic of research.

We first present the classic setup, which consists of a static game without uncertainty in which a traffic inflow has to be routed from a set of sources to a set of destinations: we mention the main results concerning equilibria, optimality, and efficiency.

Afterwards, we survey the latest proceedings in non-atomic games when some forms

of uncertainty are introduced – as in the level of traffic demands, or in the costs experienced by the players. Several sources of randomness emerge in real traffic situations and have been theoretically studied before. In the analysis of how uncertainty can affect routing behavior we focus on a form of incomplete information which can eventually be unveiled through the assumption of Bayesian rationality.

We consider routing instances with an unknown network state which influences edge costs, in which the traffic is routed by minimizing the expected cost of each path. Learning the true underlying state may be possible when this kind of incomplete information games are repeatedly played with a random level of demand in each stage game. Specifically, there are two possible kinds of social learning in non-atomic routing games: *strong learning*, when the true network state is identified, and *weak learning*, when players learn how to play as if the true state were known.

The original theoretical contribution is presented in the second part of the thesis. A recent result of [Macault et al. \(2022\)](#) clarified when *social learning* occurs in single commodity non-atomic routing games with infinite edge capacities: the authors showed that strong learning occurs almost surely if the network respects a topological condition called *series-parallel* and edge costs are continuous, strictly increasing and unbounded.

From this starting point, in the thesis we precisely define when a network state is *identifiable* and extend the network conditions required to achieve almost sure learning also in multi-commodity and capacitated instances.

The thesis is organized as follows: first, we review the fundamental literature ([Chapter 2](#)) and introduce the main [contribution](#). Then, in [Chapter 3](#) we report the central concepts and results of routing games played without uncertainty; afterwards, in [Chapter 4](#) we introduce uncertainty and we discuss a model in which a repeated routing game has an unknown network state that influences the costs experienced by the players. We also discuss the main ideas of Social Learning and routing under private information. Finally, in [Chapter 5](#) we study under which conditions Social Learning occurs in non-atomic routing games. The [conclusion](#) follows, while the proofs and the list of symbols are contained in the [Appendix](#).

1.1 Contribution

The main contribution of this work lies in the extension of some results of social learning in non-atomic routing games: more precisely, we try to understand under which circumstances social learning occurs in a dynamic non-atomic routing games both with and without capacities.

In doing so, we first study some structural properties arising in capacity constrained networks. Then, for the case with possibly infinite capacities, we study a class of networks called *series-parallel*, which is particularly important because excludes sub-network structures that may prevent learning. We generalize the desirable properties of series-parallel networks to the case of multiple source-destination pairs. Also, we provide some examples in which social learning may fail.

Specifically, we extend some known results about the conditions under which the public belief of the players converges almost surely to the true state of nature.

Considering strictly increasing, continuous and unbounded costs, what emerges is that the learning conditions differ between capacitated instances and instances with infinite edge capacities. In capacitated instances a sufficient condition to achieve learning is that, under full congestion, the load on each edge reaches the edge load upper bound. This condition is viewed as a weaker form of *capacity conservation*; under capacity conservation, the total capacity entering in each node equals to the total capacity exiting from that node. It is noteworthy to notice that this condition does not require any restriction on the network topology.

On the other hand, in infinite capacity instances a sufficient condition to achieve learning is that the sub-network available to each commodity is *series-parallel*.

Chapter 2

Literature review

Seminal ideas of traffic assignment models as instances of *selfish routing* date back to [Wardrop \(1952\)](#) and [Wardrop and Whitehead \(1952\)](#), which first introduced the notion of an equilibrium flow. The same problem was studied by [Beckmann et al. \(1956\)](#) and by [Ford and Fulkerson \(1962\)](#). Optimal flows and some efficiency measures of routing instances are first studied in [Megiddo \(1974\)](#), [Koutsoupias and Papadimitriou \(1999\)](#) and [Papadimitriou \(2001\)](#). More recent fundamental results on the efficiency of a routing instance are due to [Christodoulou and Koutsoupias \(2005\)](#) and [Roughgarden \(2002\)](#). [Roughgarden \(2007\)](#) and [Roughgarden \(2016\)](#) provide a general overview of both the atomic setup and the non-atomic setup. Also, the deterministic non-atomic setup is covered in depth in [Correa and Stier-Moses \(2011\)](#).

Some results on the convergence of atomic to non-atomic instances are found in [Haurie and Marcotte \(1985\)](#) and [Cominetti et al. \(2020\)](#). The efficiency of an equilibrium flow is studied in [Roughgarden \(2002\)](#); the same problem, but with a different approach, is addressed by [Correa et al. \(2008\)](#). [Colini-Baldeschi et al. \(2017\)](#) and [Cominetti et al. \(2021\)](#) studied the efficiency of routing games with variable traffic inflows. [Fisk \(1979\)](#) provides some examples of paradoxes that may emerge from equilibrium flows.

Routing games are part of the more general classes of *Congestion games* and *Potential games*. [Rosenthal \(1973\)](#) introduced models of congestion games; equilibria in congestion games are studied in [Holzman and Law-Yone \(1997\)](#), while [Scarsini and Tomala \(2012\)](#) provided an analysis of repeated Congestion games. [Monderer and Shapley \(1996\)](#) first described the class of potential games. [Sandholm \(2001\)](#) discussed the case of instances with continuous players' set.

Some extensions arise when considering underlying capacitated networks and some

sources of uncertainty. Capacitated non-atomic instances are first studied in [Correa et al. \(2004\)](#) - in these models, each edge of the underlying graph is endowed with a non-negative capacity. [Cominetti \(2015\)](#) provides a survey of selfish routing under uncertainty. Also, [Ordóñez and Stier-Moses \(2010\)](#) and [Cominetti and Torrico \(2016\)](#) study some models of risk-averse routing under uncertainty.

A source of uncertainty may be a lack of information about the game played. [Gairing et al. \(2008\)](#) considered a model of selfish routing under incomplete information. In the absence of common knowledge of all the aspects of a game, the provision public or private information may influence the routing behavior, as it is studied in [Tavafoghi and Teneketzis \(2017\)](#) and [Tavafoghi and Teneketzis \(2019\)](#). [Acemoglu et al. \(2018\)](#), [Wu et al. \(2017\)](#) and [Wu et al. \(2021\)](#) study some other informational aspects of Congestion games in the presence of private information. Finally, [Koessler et al. \(2021\)](#) deal with the information design problem in non-atomic congestion games.

Certain forms of uncertainty can eventually be learned thanks to the repetition of a game. A seminal work of Bayesian learning in infinitely repeated dynamic games is [Kalai and Lehrer \(1994\)](#). Social Learning is first studied in [Banerjee \(1992\)](#), [Bikhchandani et al. \(1992\)](#), and [Smith and Sørensen \(2000\)](#). [Arieli and Mueller-Frank \(2021\)](#) addresses the problem of social learning of an unknown state parameter. [Macault et al. \(2022\)](#) and [Wu and Amin \(2019\)](#) provide some results on social learning of an unknown network state. [Goeree et al. \(2006\)](#) deals with social learning in the presence of private and common values. Other aspects of learning in non-atomic games can be found in [Mertikopoulos \(2019\)](#) and [Hadikhanloo et al. \(2021\)](#).

In certain circumstances, the topology of the underlying network is key in order to guarantee learning. To this extent, the class of *series-parallel* networks is important because it excludes the possibility that a certain route is never used in equilibrium. [Duffin \(1965\)](#) studies the topology of series-parallel networks, while [Chen et al. \(2016\)](#) and [Milchtaich \(2006\)](#) apply this network topology to routing instances.

In our work, we build on concepts and results derived from more general sources other than the routing literature. For fundamental results of game theory, we rely on [Karlin and Peres \(2017\)](#). [Fudenberg et al. \(1998\)](#) is a reference for the theory of learning in general games. [Aumann \(1987\)](#) is a major source for what concerns Bayesian rationality; some results on Bayesian learning are derived from [Kamenica and Gentzkow \(2011\)](#), while concepts on incomplete information games and on information design are from [Forges \(2020\)](#), [Bergemann and Morris \(2019\)](#) and [Mathevet et al. \(2020\)](#). For other fundamental results we rely on [Gut \(2005\)](#), [Bondy and Murty \(2008\)](#), and [Williamson \(2019\)](#).

Chapter 3

Deterministic Routing Games

The preliminary setup required to define routing games consists of a directed multi-graph describing the paths available to the players; each edge of the graph is endowed with a cost function, which depends on the amount of traffic that routes through that edge. Players are selfish in the sense that they choose a path that minimizes their costs; in equilibrium, each player chooses a path of minimum cost.

Formally, we consider a directed multigraph

$$\mathcal{N} = (V, E)$$

where V is the set of vertices and E is the set of edges. The graph \mathcal{N} is endowed with a set of k *source-destination* vertex pairs,

$$\{(s_1, t_1), \dots, (s_i, t_i), \dots, (s_k, t_k)\}, \quad \text{where } s_i, t_i \in V, \quad \forall i \in \{1, \dots, k\}.$$

A mass of anonymous players - called *commodity* - has to route some units of traffic from a source to a destination. Each commodity (s_i, t_i) is characterized by a set of paths \mathcal{P}_i that go from s_i to t_i , with the assumption that $\mathcal{P}_i \neq \emptyset$ for each commodity i . Moreover, we define the set of all paths for all the commodities in the graph as

$$\mathcal{P} = \bigcup_{i=1}^k \mathcal{P}_i.$$

In addition, each commodity i can route through the network a positive amount of infinitely divisible traffic demand d_i ; thus, we define the vector of traffic demands $\mathbf{d} = \{d_i\}_{i \in I}$. Furthermore, each edge $e \in E$ has a non-negative capacity γ_e and a cost $c_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which is assumed to be non-negative, continuous, and strictly increasing,

The network \mathcal{N} , the vector \mathbf{d} , the set of capacities $\{\gamma_e\}_{e \in E}$ and set of cost functions $\mathbf{c} = \{c_e\}_{e \in E}$ constitute an *instance*

$$G = (\mathcal{N}, \mathbf{d}, \mathbf{c}, \{\gamma_e\}_{e \in E}).$$

In this setup two models of routing have been studied: the model of *atomic selfish routing*, and the model of *non-atomic selfish routing*. The atomic selfish routing model is characterized by a finite number of players, and each player controls a non-negligible amount of the total traffic. In the non-atomic selfish routing model, the number of player is assumed to be very large, with anonymous players controlling a negligible fraction of the total traffic.

In the following sections we provide a detailed description of the two models, starting from the non-atomic case, and we present some of the most relevant results. In the first chapter of this work we will assume that edge capacities are infinity, since most of the original results were found under this assumption. Thus, an instance will be denoted only by the underlying network, costs and demands, $G = (\mathcal{N}, \mathbf{d}, \mathbf{c})$.

3.1 Non-Atomic Routing games

In the non-atomic selfish routing model, we consider a mass of anonymous players, each of them controlling a negligible fraction of the overall traffic. The deterministic model of non-atomic selfish routing is extensively covered in [Roughgarden \(2007\)](#), in [Roughgarden \(2016\)](#), and in [Correa and Stier-Moses \(2011\)](#), from which the theory in this section is mainly derived.

First, we formalize the way in which the traffic demand is routed through the network, which is done through the concept of *flow*. In non-atomic instances, the mass of traffic of each commodity is routed through the network according the choice of some of the available paths. The amount of traffic routed on each path defines a flow. More precisely, a flow \mathbf{y} is a vector that is indexed on the paths of the network and assigns to each path $P \in \mathcal{P}$ a non-negative number that represents the amount of traffic routed through that path P . A flow \mathbf{y} is said *feasible* if

$$\sum_{P \in \mathcal{P}_i} y_P = d_i \quad \forall i \in 1, \dots, k.$$

That is, a flow vector is said feasible if all the traffic demand is routed through the network.

We define the *load* on an edge as

$$x_e = \sum_{P: e \in P} y_P.$$

From the set of cost functions we can additively characterize the *cost of a path* with respect to a flow as

$$c_P(\mathbf{y}) = \sum_{e \in P} c_e(x_e). \quad (3.1)$$

Then, we identify the *social cost of a flow* as the sum of the costs of all paths in \mathcal{P} ,

$$c(\mathbf{y}) = \sum_{P \in \mathcal{P}} c_P(\mathbf{y}) \cdot y_P. \quad (3.2)$$

Equivalently, the cost of a flow can be expressed edge-by-edge. Indeed from (4.2) and reversing the summation we can write

$$c(\mathbf{y}) = \sum_{P \in \mathcal{P}} c_P(\mathbf{y}) \cdot y_P = \sum_{e \in E} \sum_{P: e \in P} c_e(x_e) \cdot y_P = \sum_{e \in E} c_e(x_e) \cdot x_e.$$

The concept of social cost of a flow is important because it is a natural criterion to assess to efficiency of a flow.

3.1.1 Examples: Pigou network and Braess network

Two of the most important examples of non-atomic routing instances arise in the Pigou network and in the Braess network.

First, a Pigou network is composed by two vertices - one of which is the origin, s , and other the destination, t - and by two edges - an upper edge and a lower edge. The game is played by a mass of players, identified by an amount of total traffic demand normalized to 1. Cost functions are such that the cost on the lower edge is equal to the load on that edge, and the cost on the upper edge is constant and equal to 1. Players independently choose one of the two paths from s to t in order to minimize their travelling cost. Here, for all players the lower route is always




Figure 3.1: Pigou example.

(weakly) preferred to the upper route: the lower edge is, by construction, a weakly

dominant strategy and it is at least as attractive as the upper edge even if it is fully congested: this identifies the unique equilibrium of this game.

An optimal outcome can be identified by minimizing the total cost experienced by the players, which in this case is $(1 - x) \cdot c_1(x) + x \cdot c_2(x) = 1 - x + x^2$. The optimal outcome splits the traffic in two, half up and half down. The total cost experienced by the players is then $\frac{1}{2}c_1(x) + \frac{1}{2}c_2(x) = \frac{3}{4}$. In the unique equilibrium, all the players would select the lower branch, incurring in a total cost of $1 \cdot c_2(x) = 1$. An usual measure of the efficiency of an equilibrium is the ratio between the outcome of the worst equilibrium and the optimal outcome. This ratio is called *Price of Anarchy* and in this example is $\frac{4}{3}$.

A Braess network is as follows




Figure 3.2: Braess network.

Let's consider a unit of traffic demand ($d = 1$) that has to be routed in the single source-destination network represented above. It turns out that the path $\{e_1, e_3, e_5\}$ is a dominant strategy, and all the traffic is routed through that path. The total cost generated is then 2. Also, it turns out that the best possible outcome comes from splitting the traffic in two, routing half of the traffic through $\{e_1, e_4\}$, and half through $\{e_2, e_5\}$. With this traffic allocation, the costs of each of these two paths is $\frac{3}{2}$. What emerges is that there is no profitable way of using the edge e_3 , even if it has a zero cost. Again, the ratio between the cost of the equilibrium outcome and the cost of the optimal outcome is $\frac{4}{3}$.

Considering the same examples, but without edge e_3 , leads to an important result, known as the *Braess' paradox*. Without the internal edge, the unique equilibrium flow implies to split the traffic in two: half up, through edges $\{e_1, e_4\}$, and half down, through edges $\{e_2, e_5\}$. The resulting cost of each path is $\frac{3}{2}$, and it is lower than the cost experienced by the players in the case in which e_3 is present in the network. What emerges is that adding some route possibilities deteriorates the routing performances resulting from equilibrium selfish behavior.

3.2 Congestion games and Potential games

Two general classes of games particularly important in the analysis of routing games are the classes of *congestion games* and of *potential games*. Indeed, routing games belong to the class of congestion games; moreover, any congestion game is a potential game (Rosenthal (1973)), and for any finite potential game there is a congestion game with the same potential function (Monderer and Shapley (1996)). Potential games and congestion games have many remarkable properties, which consequently also apply to the class of routing games.

In full generality, congestion games are played on a finite set of resources E . Each resource e in the set E has cost function $c_e(\cdot)$ that depends on the load on the resource e . Also, costs are assumed to be additive over resources, and the total cost incurred by a player is the sum of the costs incurred from the use of each resource. The most basic setup considers a finite set of players $I = \{1, \dots, k\}$, and each player i in I has a finite set of feasible strategies $S_i \subseteq 2^E$.

Routing games are an example of congestion games. In the non-atomic setup, the set of players coincides with the set of commodities, and the set of (pure) strategies available for commodity i is the set of feasible source-destination paths \mathcal{P}_i .

On the other side, potential games are defined by the existence of a special function called *potential function*. A potential function ψ of a finite game is a function that maps every strategy vector $S = \{S_1, \dots, S_k\}$ to a real value such that, for a generic player i , it holds that

$$\psi(S) - \psi(S') = c_i(S) - c_i(S'),$$

in which $S = (S_1, \dots, S_i, \dots, S_k)$ and $S' = (S_{-i}, S'_i)$ with $S'_i \neq S_i$.

In words, this means that if player i switches from strategy S_i to S'_i , the resulting difference in costs incurred by i is equal to the change in the potential function. Furthermore, it is a known result that a game can have at most one potential function up to a constant factor, and that every potential game has at least one pure Nash Equilibrium, that is a strategy S that minimizes locally $\psi(S)$.

Non-atomic routing games are potential games with continuous players' sets (Sandholm (2001)), and as such they always admit a Nash Equilibrium in pure strategies.

More specifically, Sandholm (2001) showed that potential games with continuous players' sets can be defined by the existence of a continuously differentiable potential function ψ whose gradient equals the vector of payoffs. This function is unique up to an additive constant and for non-atomic routing games it is such that

$$\frac{\partial \psi(\mathbf{x})}{\partial x_e} = c_e(x_e).$$

The potential function of non-atomic routing games is defined as

$$\psi(\mathbf{x}) = \sum_{e \in E} \int_0^{x_e} c_e(x) dx,$$

and its local minimizers are equilibrium loads.

Some other results, particularly on convergence and adaptive learning, hold for potential and congestion games, and therefore also for routing games. We do not analyze them in this context, but we refer to [Roughgarden \(2007\)](#), [Cominetti et al. \(2020\)](#), [Scarsini and Tomala \(2012\)](#) and [Sandholm \(2001\)](#).

3.3 Equilibria and Optimality

Informally, the idea behind a *flow in equilibrium* is that, when a positive flow is assigned to a path, this path is cost-minimizing; in other words, an equilibrium flow travels only in minimum cost paths, preventing any profitable traffic reallocation.

An *optimal flow*, on the other hand, is a flow that minimizes the total cost experienced by the players. These general ideas of optimal and equilibrium flows hold for both *non-atomic* and *atomic* instances, but the different setups induce different definitions and results among the two cases. In what follows we first start with the formalization of equilibria and optimality in *non-atomic* instances. The definitions and results for the *atomic* case are different, and will be covered later in the chapter.

In non-atomic games, an equilibrium flow is also called *Nash Flow* or *Wardrop Equilibrium*, from [Wardrop and Whitehead \(1952\)](#). Precisely, a Wardrop Equilibrium is defined as follows.

Definition 3.3.1 (Wardrop Equilibrium) *A flow \mathbf{y}^* , feasible for a non-atomic instance $(\mathcal{N}, \mathbf{d}, c)$, is a Wardrop Equilibrium if and only if, for all $i \in \{1, \dots, k\}$ and for all $P, \tilde{P} \in \mathcal{P}_i$ with $y_P > 0$, we have*

$$c_P(\mathbf{y}^*) \leq c_{\tilde{P}}(\mathbf{y}^*).$$

What follows is that the costs of all used paths are equal, and they are lower than the costs which would be experienced by any arbitrary mass of traffic on any unused path. A slightly different characterization of equilibrium flows is expressed in terms of a variational inequality. Indeed, a flow is in equilibrium if and only if the induced equilibrium load vector $\mathbf{x}^* = \{x_e^*\}_{e \in E}$ is such that

$$\sum_{e \in E} c_e(x_e^*) \cdot x_e^* \leq \sum_{e \in E} c_e(x_e^*) \cdot x_e$$

for all feasible edge loads x_e .

Notice that, since an equilibrium flow travels only on minimum cost paths, in equilibrium all paths used by each commodity have equal costs, denoted by $L_i(\mathbf{y}^*)$. The characterization of the cost of a flow at equilibrium follows.

Proposition 3.3.1 *If \mathbf{y}^* is a flow at Nash Equilibrium for a non-atomic instance $(\mathcal{N}, \mathbf{d}, c)$, then*

$$c(\mathbf{y}^*) = \sum_{i=1}^k L_i(\mathbf{y}^*) \cdot d_i.$$

An equilibrium arising from selfish behavior in general does not optimize any criterion of social efficiency. Considering the *total cost of a flow* as objective criterion, an *optimal flow* solves the following non-linear program,

$$\begin{aligned} \min_{\mathbf{x}} \quad & \sum_{e \in E} c_e(x_e) \cdot x_e \\ \text{s.t.} \quad & \sum_{P \in \mathcal{P}_i} y_P = d_i \quad \forall i \\ & x_e = \sum_{P: e \in P} y_P \quad \forall e \in E \\ & y_P \geq 0 \quad \forall P \in \mathcal{P} \end{aligned} \tag{3.3}$$

Now, we can define the total cost on an edge as $x \cdot c_e(x)$. This is also the contribution to the objective criterion by the traffic on edge e . If costs $c_e(\cdot)$ are continuously differentiable, the marginal cost on an edge is denoted by

$$c_e^*(x) = \frac{\partial c_e(x) \cdot x}{\partial x} = c_e(x) + x \cdot c_e'(x).$$

Moreover,

$$c_p^*(\mathbf{y}) = \sum_{e \in P} c_e^*(x).$$

As shown in [Roughgarden \(2007\)](#), the following proposition holds.

Proposition 3.3.2 *Let $(\mathcal{N}, \mathbf{d}, c)$ be a non-atomic instance such that for every edge the function $x \cdot c_e(x)$ is convex and continuously differentiable. A flow $\hat{\mathbf{y}}$ is optimal for (4.3) if and only if, for all $i \in \{1, \dots, k\}$, and for all $P, \tilde{P} \in \mathcal{P}_i$ with $f_P > 0$, it holds*

$$c_P^*(\hat{\mathbf{y}}) \leq c_{\tilde{P}}^*(\hat{\mathbf{y}}).$$

The similarity between the characterization of *optimal flows* and that of *equilibrium flows* is clear. As reported in [Roughgarden \(2002\)](#), the following corollary holds.

Corollary 3.3.1 *Let $(\mathcal{N}, \mathbf{d}, c)$ be a non-atomic instance with differentiable cost functions, where $x \cdot c_e(x)$ is convex for all $e \in E$ and in which marginal costs are denoted by $c_e^*(x)$. A flow $\hat{\mathbf{y}}$ feasible for $(\mathcal{N}, \mathbf{d}, c)$ is optimal if and only if it is a Wardrop Equilibrium for $(\mathcal{N}, \mathbf{d}, c^*)$.*

A Nash flow of an atomic routing game always exists. This comes from the fact that routing games are potential games with continuous players' sets. Moreover, in the non-atomic case, Nash flows are essentially unique, in the sense that multiple Nash flows are allowed, but all equilibrium flows induce the same costs. Also *optimal flows* always exist, and this comes from that fact that they are the solution of a convex optimization problem on a set can be shown to be closed and bounded.

The existence of a flow at Nash equilibrium comes from the following proposition, which is proved in [Roughgarden \(2007\)](#) through the potential function.

Theorem 3.3.1 *A non-atomic instance $(\mathcal{N}, \mathbf{d}, c)$ with continuous and non-decreasing cost functions admits at least one flow at Nash Equilibrium. Moreover, if \mathbf{y} and $\tilde{\mathbf{y}}$ are flows at Nash Equilibrium, then*

$$c_e(x_e) = c_e(\tilde{x}_e), \quad \forall e \in E.$$

The original argument for this proposition was initially provided in [Beckmann et al. \(1956\)](#), and it uses the equivalence of a *Nash Flows* and the local minima of a *potential function* which completely characterizes the costs of the game. The potential function of a non-atomic routing game is

$$\psi(\mathbf{x}) = \sum_{e \in E} \int_0^{x_e} c_e(x) dx, \quad (3.4)$$

and its local minima respect the equilibrium conditions. Since edge costs are continuous and non-decreasing, the minimization of (4.4) is a convex problem, and its local minima coincide with its global minima. Equilibrium flows can therefore be characterized as the global minimizers of the potential function associated to an instance $(\mathcal{N}, \mathbf{d}, c)$. The following proposition holds.

Proposition 3.3.3 *Let $(\mathcal{N}, \mathbf{d}, c)$ be a non-atomic instance. A flow \mathbf{y}^* feasible for $(\mathcal{N}, \mathbf{d}, c)$ is at Nash Equilibrium if and only if it is a global minimum of the corresponding potential function ψ .*

The former results concerned instances on networks endowed with infinite capacities. An analysis of non-atomic instances on capacitated networks can be found in [Correa et al. \(2004\)](#), in which they clarify that if capacities are finite, then the Wardrop principle at equilibrium doesn't hold anymore - that is, equilibrium costs of paths are not guaranteed to equal and equilibrium loads are not always unique. If capacities

are finite, a path P is called *unsaturated* if

$$x_e < \gamma_e \quad \forall e \in P.$$

The following definition of equilibrium (now called *capacitated user equilibrium*) holds.

Definition 3.3.2 *A flow \mathbf{y}^* represents a capacitated user equilibrium if no source-destination pair (s_i, t_i) has an unsaturated path with strictly smaller cost than any path used for that pair. That is, if $y_P > 0$ for $P \in \mathcal{P}_i$, then*

$$c_P(\mathbf{y}^*) \leq \min\{c_{P'}(\mathbf{y}^*) : P' \in \mathcal{P} \text{ is unsaturated}\}.$$

This definition is equivalent to the definition of *Wardrop equilibrium* considering infinite capacities on all edges. What emerges is that two paths with positive flow in the same (s_i, t_i) pair in a *capacitated user equilibrium* may have different costs. This opens the doors for multiple equilibria - which are not anymore essentially unique, in the sense that several equilibria with different total costs may emerge.

Nevertheless, if edge costs go to infinity as the edge loads approach the edge capacities, then the capacitated instance can be seen as an infinite capacity instance with a specific class of costs functions. In other words, if edge capacities are incorporated in the costs functions, then the essential uniqueness result of uncapacitated networks holds.

3.4 Price of Anarchy

The performance of a flow \mathbf{y} in a traffic routing model can be assessed by the total cost induced by \mathbf{y} , as defined in (4.2). One natural measure of the efficiency of a routing instance is given by the ratio between the total cost induced by an equilibrium flow $c(\mathbf{y}^*)$ and the total cost induced by the optimal flow $c(\hat{\mathbf{y}})$.

The *Price of Anarchy* is defined as the ratio between the worst possible equilibrium outcome and the optimal outcome. The term was first used in [Koutsoupias and Papadimitriou \(1999\)](#), and since then it has been extensively studied in several contexts. In a cost minimization game, the *Price of Anarchy* is clearly lower bounded by one, and a value close to the unity suggests that the worst possible outcome arising from selfish behavior is approximately the same of the optimal outcome. On the other side, upper bounding the Price of Anarchy is not a trivial task and constitutes a major issue in many problems and applications.

In non-atomic routing games equilibrium costs are always unique, hence the Price

of Anarchy simplifies to the ratio between the cost of any equilibrium flow and the cost of the optimal equilibrium flow,

$$PoA = \frac{c(\mathbf{y}^*)}{c(\hat{\mathbf{y}})}.$$

For non-atomic routing games it has been proven that the PoA depends on the degree of non-linearity of the cost functions, and on nothing else (Roughgarden (2002)). More precisely, the PoA in selfish routing games is high in networks with highly non-linear cost functions. First, given a general class of cost functions, let's provide the following definition.

Definition 3.4.1 (Pigou Bound) *Let \mathcal{C} be a non-empty set of cost functions. The Pigou Bound of \mathcal{C} , denoted as $\alpha(\mathcal{C})$ is*

$$\alpha(\mathcal{C}) = \sup_{c \in \mathcal{C}} \sup_{x, d \leq 0} \frac{d \cdot c(d)}{x \cdot c(x) + (d - x) \cdot c(d)}.$$

This quantity arises from the Pigou-like network as introduced in the previous section. Indeed, the equilibrium flow in a Pigou network routes all the traffic through the edge with non-constant cost, and the resulting Price of Anarchy is always of the form

$$\frac{d \cdot c(d)}{x \cdot c(x) + (d - x) \cdot c(d)}.$$

We can notice that in every instance we can explicitly solve for x , which stands for the amount of traffic that is routed through the edge with non-constant cost. Informally, the Pigou Bound searches for the worst possible specification of the function c in the set \mathcal{C} and for the worst possible traffic demand d . A famous result (Roughgarden (2002), Roughgarden (2007)) shows that for each class of cost functions the Price of Anarchy is maximized in the simplest case - that is, in Pigou-like networks. The following proposition holds.

Proposition 3.4.1 *Let \mathcal{C} be a set of cost functions and let $\alpha(\mathcal{C})$ be the Pigou Bound for \mathcal{C} . If $(\mathcal{N}, \mathbf{d}, c)$ is non-atomic with costs in \mathcal{C} , then the Price of Anarchy of $(\mathcal{N}, \mathbf{d}, c)$ is at most $\alpha(\mathcal{C})$.*

A proof can be found in Roughgarden (2016) and uses an argument in two steps, which we sketch hereafter. The first step shows that, if we consider edge costs at their equilibrium values, then the following inequality holds,

$$\sum_{e \in E} (\hat{x}_e - x_e^*) \cdot c_e(x_e^*) \geq 0, \tag{3.5}$$

in which \hat{x}_e denotes an edge load induced by an optimal flow, and x_e^* denotes an edge load induced by an equilibrium flow. The second part of the argument starts from the inequality

$$\alpha(\mathcal{C}) \geq \frac{x_e^* \cdot c_e(x_e^*)}{\hat{x}_e \cdot c_e(\hat{x}_e) + (x_e^* - \hat{x}_e) \cdot c_e(x_e^*)}.$$

This comes from that fact that $\alpha(\mathcal{C})$ is defined as the superior of the right hand side on each of its variables. Rearranging and summing over all edges $e \in E$, this entails that

$$c(\hat{\mathbf{y}}) \geq \frac{1}{\alpha(\mathcal{C})} \cdot c(\mathbf{y}^*) + \sum_{e \in E} (\hat{x}_e - x_e^*) \cdot c_e(x_e^*),$$

in which $c(\hat{\mathbf{y}}) = \sum_{e \in E} c_e(\hat{x}_e)$ is the total cost of the optimal flow, and $c(\mathbf{y}^*) = \sum_{e \in E} c_e(x_e^*)$ is the cost of a Nash flow. Since the second term is not lower than zero by (3.5), this gives the desired upper bound,

$$\frac{c(\mathbf{y}^*)}{c(\hat{\mathbf{y}})} \leq \alpha(\mathcal{C}).$$

In [Roughgarden \(2002\)](#) it is also shown that this bound is tight, in the sense that for every class of allowable cost functions, the PoA is maximized by the Pigou Bound.

A simpler result arises when the class \mathcal{C} is the set of *affine cost functions*, of the form $c(x) = a \cdot x + b$. In this case, the following proposition holds ([Roughgarden \(2002\)](#), [Correa et al. \(2008\)](#)).

Proposition 3.4.2 *If $(\mathcal{N}, \mathbf{d}, c)$ has affine cost functions, then the Price of Anarchy of $(\mathcal{N}, \mathbf{d}, c)$ is at most $\frac{4}{3}$.*

An extensive analysis of the Price of Anarchy and its asymptotic behavior in routing games can be found in [Colini-Baldeschi et al. \(2017\)](#), [Cominetti et al. \(2019\)](#), and [Cominetti et al. \(2021\)](#).

3.5 Atomic Routing games and convergence results

Again, we consider a directed graph $\mathcal{N} = (V, E)$ with k source-destination pairs and an amount d_i of demand for each pair (s_i, t_i) . Moreover, each edge is endowed with a non-negative, continuous and non-decreasing cost function and has unlimited capacity. As in the case of non-atomic instances, an atomic instance is defined as the set $(\mathcal{N}, \mathbf{d}, \mathbf{c})$. In atomic games the set of players is finite, hence atomic instances are finite simultaneous games.

Differently from the non-atomic case, each player controls a non-negligible amount

of traffic. The strategy set of a generic player i is given by the set of paths available to her, \mathcal{P}_i . Choosing a specific path $P \in \mathcal{P}_i$, a player routes d_i units of traffic on P .

A flow \mathbf{y} in an atomic instance is a non-negative vector indexed by both players and paths: we say that y_P^i is the amount of traffic that player i routes on the path P from source s_i to destination t_i . We say that an (unsplittable) flow \mathbf{y} is feasible if for all players y_P^i is equal to d_i for one s_i - t_i path, and zero for all other paths. An atomic equilibrium flow is a feasible flow such that no player can strictly decrease its cost by choosing a different path.

Definition 3.5.1 (Equilibrium Flows, atomic instances) *Let \mathbf{y} be a feasible flow of an atomic instance $(\mathcal{N}, \mathbf{d}, c)$. We say that \mathbf{y}^* is an equilibrium flow if, for all players $i \in \{1, \dots, k\}$, and for all couples $P, \tilde{P} \in \mathcal{P}_i$ of s_i - t_i paths with $y_P^i > 0$, we have*

$$c_P(\mathbf{y}^*) \leq c_{\tilde{P}}(\tilde{\mathbf{y}}).$$

In which $\tilde{\mathbf{y}}$ is identical to \mathbf{y}^ except that $\tilde{y}_P^i = 0$ and $\tilde{y}_{\tilde{P}}^i = d_i$.*

A first difference with the non-atomic case is that, on atomic instances, different equilibrium flows can have different costs. Another difference regards the existence of equilibrium flows: in atomic instances equilibrium flows may not exist. Indeed, the previous definition equates an equilibrium flow to a Nash equilibrium in pure strategies, which does not always exist. Nonetheless, it is a known result that if all the players route the same amount of traffic, or if cost functions are affine, then at least one equilibrium flow exists. Both theorems are enunciated and proved in [Roughgarden \(2007\)](#); similarly to the non-atomic case, the proof follows a potential function argument, with discrete potentials. Finally, the Price of Anarchy of atomic instances can exceed the upper bound of non-atomic instances ([Roughgarden \(2007\)](#)).

Non-atomic instances are thought as an asymptotic approximation of atomic instances with many players. A preliminary analysis of the asymptotic behavior of an equilibrium in a finite congestion games is due to [Haurie and Marcotte \(1985\)](#). They show that the asymptotic behavior of an equilibrium converges, under appropriate assumptions, to a vector flow corresponding to a Wardrop Equilibrium.

Other results on when non-atomic games can be considered the limiting case of atomic games are found in [Cominetti et al. \(2020\)](#), in which it is clarified whether and when a Wardrop equilibrium is a good approximation of a Nash equilibrium of a finite congestion game. The starting point is an atomic congestion game with unsplittable flow and many small players. The flow is unsplittable in the sense that players must route a given traffic demand on a single path, either deterministically or randomly through a mixed strategy. They study two cases of games with small players: the case in which the weight of every player is small, and the case in which

the participation probability of each player is small.

In the first case it is studied an atomic routing game in which the number of players grows to infinity and their weight decreases to zero. What emerges is that the random flows in all mixed Nash equilibria of the atomic game converge in distribution to the set of Wardrop equilibria of the corresponding non-atomic game.

In the second case, an increasingly number of players with a unity weight participates in the atomic game with decreasingly probability. Here, the Nash equilibrium flow converges to a family of Poisson random variables whose expected values constitute the Wardrop equilibrium of a suitably defined non-atomic game.

Chapter 4

Routing Games under Uncertainty

Multiple sources of uncertainty can affect routing behavior. Consequently, there are several extensions on how to introduce and model uncertainty in the deterministic routing games defined so far. For example, uncertainty may derive from the randomness in traffic inflows, or from some random noise that affects edge costs, or again from a degree of uncertainty in the participation of the players. These forms of uncertainty can be modelled as routing games under incomplete information; in this sense, recent studies are found in [Cominetti \(2015\)](#), [Wu and Amin \(2019\)](#), [Cominetti et al. \(2020\)](#), and [Macault et al. \(2022\)](#).

A different way of dealing with uncertainty comes from the study of selfish routing under incomplete information, in which users have private information that exploit in their equilibrium behavior. This form of uncertainty is modelled through a set of types, which captures all the information of a routing instance that is not common knowledge. An example arises when each commodity is aware of its own traffic demands, but doesn't know the amount of traffic deriving from the other commodities. These models of incomplete information require new equilibrium concepts, and a suitable information design may induce desirable equilibria.

In this work we mainly concentrate on a model of incomplete information with random demands and uncertain costs. The uncertainty in the costs only depends on the underlying network state, which is initially unknown - hence, differently from other models, costs are not affected by any other source of random fluctuation. For the sake of completeness, at the end of the chapter we also survey the central ideas of routing under imperfect information, and we briefly discuss the information design problem in routing games.

4.1 Introducing uncertainty

Our model under uncertainty follows the view introduced by [Macault et al. \(2022\)](#): we consider a routing instance in which the true network conditions are not known by the players and are modelled through an *unknown network state*. The unknown state influences the instance only through the edge costs, whose true specification is initially not known. Players have a belief distribution on the network state and play an equilibrium accordingly to their belief.

Formally, we introduce a *state space* as a finite probability space

$$(\mathcal{T}, 2^{\mathcal{T}}, \mu),$$

in which \mathcal{T} is the finite set of possible network states, $2^{\mathcal{T}}$ is its power set and μ is the common belief. Moreover, Θ denotes the unknown random state, and $\theta \in \mathcal{T}$ its realization.

If G is a deterministic [non-atomic routing instance](#), a non-atomic routing game with an unknown network state is denoted by

$$G_{\mu} = [G, (\mathcal{T}, 2^{\mathcal{T}}, \mu)].$$

As in the deterministic case, cost functions are assumed to be continuous and strictly increasing in the load $x \in [0, \gamma_e)$ for each edge $e \in E$ and state $\theta \in \mathcal{T}$.

In full generality, given a belief distribution $\mu \in \Delta(\mathcal{T})$, the expected edge costs are

$$c_e(x, \mu) = \int_{\mathcal{T}} c_e(x, \theta) d\mu(\theta),$$

and the expected costs of a generic path $P \in \mathcal{P}$ is

$$c_P(\mathbf{y}, \mu) = \int_{\mathcal{T}} c_P(\mathbf{y}, \theta) d\mu(\theta).$$

In the specific model considered in this thesis the set of states \mathcal{T} is finite, hence expectations are intended as sums. We can now define equilibria.

Definition 4.1.1 *A flow vector $\mathbf{y}^* \in \mathcal{Y}$ is a Wardrop Equilibrium of G_{μ} if for all paths $P, P' \in \mathcal{P}$ with $y_P^* > 0$ we have that*

$$c_P(\mathbf{y}^*, \mu) \leq c_{P'}(\mathbf{y}^*, \mu).$$

We can notice that now equilibrium flows are defined in terms of expected costs. Also, given a traffic demand d , a *load* at Wardrop Equilibrium of G_{μ} is denoted by $\mathbf{x}^*(d, \mu)$. Equivalently, a flow is at equilibrium in a capacitated instance with unknown state if no source-destination pair has an [unsaturated path](#) with strictly smaller expected cost than any path used for that pair.

4.2 Dynamic Routing games with unknown network state

Repeating a game is a fundamental ingredient to achieve any form of learning. In this section we introduce the notion of a *Dynamic Routing Game*: a routing instance repeated over time with a random traffic inflow in each stage game.

In what follows a routing game with an uncertain network state and random demands is played over time; time is discrete and γ is the total capacity of the underlying network, which is possibly infinite.

In the case of a single commodity, traffic demands are given by a sequence

$$\{D^t\}_{t \in \mathbb{N}}$$

of i.i.d. non-negative random variables with common marginal distributions bounded above by γ . Let D denote a generic element of the sequence $\{D^t\}_{t \in \mathbb{N}}$, its marginal distribution be denoted by F and its support by $\text{supp}(D)$. Assuming independence between demands and the random network state, the marginal F and the prior μ induce a unique product measure \mathbb{P} on the space

$$([0, \gamma]^\infty \times \mathcal{T}, \mathcal{B}([0, \gamma]^\infty) \otimes 2^{\mathcal{T}}).$$

The model goes as follows: in each period t a traffic demand d^t is realized and publicly observed. An *equilibrium flow* is played and the information about equilibrium loads \mathbf{x}^{*t} and costs $c(\mathbf{x}^{*t}, \theta) = [c_e(x_e^{*t}, \theta)]_{e \in E}$ is immediately broadcasted to all the players.

For every time period $t \in \mathbb{N}$, an *history* is denoted by

$$h^t := [D^1, \mathbf{x}^{*1}(D^1), c(\mathbf{x}^{*1}(D^1), \Theta), \dots, D^{t-1}, \mathbf{x}^{*(t-1)}(D^{t-1}), c(\mathbf{x}^{*(t-1)}(D^{t-1}), \Theta), D^t],$$

and the public belief at time t is μ^t .

A repeated non-atomic routing game endowed with the product measure \mathbb{P} defines a *Dynamic non-atomic routing game* (DNRG),

$$\Gamma = (G, \mathbb{P}).$$

The same model holds for multi-commodity instances, in which in every stage game we consider a random demand vector instead of a scalar,

$$\{\mathbf{D}^t\}_{t \in \mathbb{N}}$$

in which $\mathbf{D} = \{D_i\}_{i=1}^k$.

4.2.1 Social Learning

The central idea of social learning is that players, repeatedly playing a game, are collectively able to exploit the randomness of traffic demands to unveil the randomness in the network state. In each stage game the public belief is updated according to the Bayes' rule, and as soon as a state is identified by some equilibrium costs, the public belief changes.

Specifically, let μ^t denote the posterior public belief conditioned on the whole history up to period t , $\mu(\cdot|h^t)$. This distribution assigns to each state $\theta \in \mathcal{T}$ the probability

$$\mu^t(\theta) = \mathbb{P}(\Theta = \theta|h^t).$$

It follows that the prior in period zero is the distribution from which the unknown state is drawn, and it is common knowledge.

The posterior μ^t is a random variable because it depends on the random demands up to period t , and being a probability distribution it is always bounded. It follows that the sequence of posterior beliefs is a bounded martingale. Hence, by the martingale convergence theorem (see, for example, [Gut \(2005\)](#)), there exists a random variable μ^∞ such that

$$\mu^t \xrightarrow{a.s.} \mu^\infty.$$

Also, since the set of network states \mathcal{T} is finite, there exists a random time $\tau \in \mathbb{N}$ such that almost surely

$$\mu^t = \mu^\infty, \quad \forall t \geq \tau.$$

Now, let δ_θ be the degenerate Dirac measure that assigns positive probability only to state θ . Denoting by Θ the random state, δ_Θ is the random Dirac measure before the dynamic instance is played.

Building on this setup, two ideas of social learning are formalized: *Strong Learning* and *Weak Learning*. The idea of Strong Learning is that the random sequence of posteriors converges almost surely to the random Dirac measure δ_Θ . Weak Learning, on the other side, implies that true state is not necessarily discovered, but in equilibrium and asymptotically the total traffic is routed as if the true state were known.

Definition 4.2.1 *Consider a dynamic non-atomic routing game with unknown network state Γ . We say that*

1. *Strong Learning is achieved if*

$$\mu^\infty = \delta_\Theta \text{ almost surely,}$$

2. *Weak Learning is achieved when*

$$\mathbf{x}^*(\cdot, \mu^\infty) = \mathbf{x}^*(\cdot, \delta_\Theta) \text{ almost surely.}$$

Some antecedent results that prepared to the ideas of social learning date back to Kalai and Lehrer (1993) and Kalai and Lehrer (1994). The classical papers on social learning are Banerjee (1992), Bikhchandani et al. (1992), and Smith and Sørensen (2000). More recently, sequential social learning is studied in Arieli and Mueller-Frank (2021), and two application to routing games are found in Wu and Amin (2019) and Macaulf et al. (2022).

4.3 Private Information

A different source of uncertainty may derive from private information. This section contains a brief overview on how routing can also be viewed and modeled under this additional layer of uncertainty.

The routing problem is modelled as a *Bayesian Routing game* in which users may have different *types*, with private information about their own type. For example, a given type may represent the amount of traffic that a commodity can route. A *type set* is introduced, and also a *type distribution* over the set of possible *type profiles*. Two implications are that individual costs are given by the expectation over the type distribution, and that if each commodity has a unique possible type, then we are back to the model without private information.

Formally, we start with a static routing game with *incomplete information* G_μ , in which the traffic is assigned according to the common belief μ . To model private information, we assign a finite set of types T^i to each commodity $i \in I$. We introduce an *information structure* (\mathbf{T}, π) , in which

$$\mathbf{T} = T^1 \times \dots \times T^k,$$

and π is a set of probability distributions such that, given a state realization θ , the type profile $\mathbf{t} = \{t_1, \dots, t_k\} \in \mathbf{T}$ is drawn with probability

$$\pi(\mathbf{t}|\theta).$$

Here, commodity i observes $t^i \in T^i$, while does not observe \mathbf{t}^{-i} .

A Routing game with private and incomplete information can be modelled as an incomplete information game endowed with an information structure

$$[G_\mu, (\mathbf{T}, \pi)].$$

Let $y_P^i(t^i)$ denote the flow of commodity i which observes type t^i and chooses path $P \in \mathcal{P}_i$. The total flow on a path P is given by

$$y_P(\mathbf{t}) = \sum_{i \in I} y_P^i(t^i),$$

while a flow vector is

$$\mathbf{y}(\mathbf{t}) = [y_P(\mathbf{t})]_{P \in \mathcal{P}}.$$

An extension the *Wardrop Equilibrium* to Bayesian routing games is given by the *Bayes Wardrop Equilibrium* (BWE).

A BWE is a feasible flow vector $\mathbf{y}(\cdot)$ such that for all $i \in I$, $t^i \in T^i$, and $P, P' \in \mathcal{P}$, we have that if $y_P^i(t^i) > 0$, then

$$\sum_{\mathbf{t}^{-i}} \sum_{\theta} \mu(\theta) \pi(t^i, \mathbf{t}^{-i} | \theta) \cdot c_P(\mathbf{y}(\mathbf{t}), \theta) \leq \sum_{\mathbf{t}^{-i}} \sum_{\theta} \mu(\theta) \pi(t^i, \mathbf{t}^{-i} | \theta) \cdot c_{P'}(\mathbf{y}(\mathbf{t}), \theta).$$

Many extensions arise from this setup with private information. Researchers of the field have defined new equilibrium concepts and introduced correlation devices useful to induce desirable equilibrium outcomes. Indeed, models with private information accommodate for the question on how a given information structure can be designed in order to induce good equilibria.

We do not cover the details, but we point out some recent literature. A model on routing in the presence of private information is studied in [Gairing et al. \(2008\)](#). On a similar fashion, results on information incentives and on the value of private information are found in [Tavafoghi and Teneketzis \(2017\)](#), [Wu et al. \(2017\)](#), and [Acemoglu et al. \(2018\)](#). Finally, [Koessler et al. \(2021\)](#) address the information design problem in routing instances.

Chapter 5

Social Learning in non-atomic Routing Games

In this chapter we present the original part of the thesis: we examine the problem of *Social Learning* in a [DNRG](#) both when the underlying network is endowed with infinite capacities and when edge capacities are finite. [Macault et al. \(2022\)](#) showed how social learning is achieved in single source-destination networks with infinite capacities. In the same spirit, we show under which conditions social learning occurs in general networks with finite capacities; moreover, we extend the result with infinite capacities to instances with several commodities.

Before presenting the main results, we study some structural properties of capacity constrained networks. First, we notice that edge loads in capacitated networks are upper bounded: we characterize this bound, which does not always coincide with the capacity of the corresponding edge. We then reframe the definition of *feasible loads* in terms of these upper bounds. Also, from this starting point we are able to define when an unknown network state is identifiable. This preliminary analysis is key to understand when social learning occurs on capacitated networks.

Finally, we delve into the issue of *Social Learning*: we show that social learning occurs in capacitated instances if the underlying network respects a condition on its edge capacities called *weak capacity conservation*. Moreover, we show that social learning occurs in instances with infinite capacities if the underlying network respects the *series-parallel* structure in each sub-network available to its source-destination pairs.

5.1 Some structural properties

In what follows we highlight some properties of edge loads and of feasible flow vectors which result from the structure of the underlying network. We provide an expression of the loads' upper bound and we show that edge capacity conditions are respected if and only if edge loads respect their upper bounds.

This result allows us to represent the set of *feasible flows* in a multi-commodity networks as an intersection between two sets: a convex hull resulting from capacity constraints, and a set that ensures that all the traffic demand is routed. After defining the set of feasible flows, we clarify when an unknown network state is *identifiable*. Here, the upper bound on edge loads is important: indeed, to learn an unknown state, the load value which guarantees to distinguish between different states must not exceed the upper bound.

Afterwards, we study two network characterizations sufficient to achieve social learning, one in constrained networks and one in unconstrained networks.

First, for instances with infinite capacities, we recall a specific network structure that arises in single source-destination networks. This network structure is called *series-parallel*; it can be generalized to multi-commodity networks such that the sub-networks in which each commodity can route its traffic are themselves series-parallel.

Regarding instances with finite capacities, we start providing some examples of learning failure; then we introduce a new network characterization that concerns edge capacities, which is formalized in the ideas of *capacity conservation* and of *weak capacity conservation*.

5.1.1 Load upper bound and capacity conditions

We first introduce some notation. Since the network \mathcal{N} is directed, each edge has a *tail node* and an *head node*. The set of edges that exit from edge e is denoted by H_e and the set of edges that enter in edge e by T_e . More precisely, let

$$H_e = \{\text{set of exiting edges from the head node of edge } e\},$$

$$T_e = \{\text{set of entering edges into the tail node of edge } e\}.$$

Also, we recall that a flow vector $\mathbf{y} = \{y_P\}_{P \in \mathcal{P}}$ is said *feasible* if all the traffic demand is routed and if the implied edge loads are such that

$$x_e \leq \gamma_e \quad \forall e \in E,$$

in which γ_e is the edge capacity.

However, if capacities are finite, the maximum load reachable on an edge may be smaller than the edge capacity. Nevertheless, this maximum load is completely determined by the set of edge capacities. We remark that the edge upper bound does not depend on equilibrium behavior, but results from the structure of the underlying network. The following proposition characterizes the *load upper bound* of a feasible flow vector in a capacitated instance.

Proposition 5.1.1 *Given a routing instance $(\mathcal{N}, \mathbf{d}, \mathbf{c}, \{\gamma_e\}_{e \in E})$, feasible edge loads are upper bounded as follows,*

$$\forall e \in E \quad x_e \leq \hat{x}_e = \min \left\{ \min_{e \in E_e} \gamma_e, \sum_{P: e \in P} \min_{e \in P} \gamma_e \right\}.$$

In which for a given edge $e \in E$, $E_e \subseteq E$ is defined recursively in two steps,

1. $e \in E_e$
2. $e' \in E_e$ iff for some $e \in E_e$, $H_e = \{e'\}$ or $T_e = \{e'\}$.

The [proof](#) can be found the appendix.

In words, we can refer to the set E_e as the largest collection of edges that contains edge e and stops as soon as more than one edge exits or enters in a edge contained in E_e .

The upper bound \hat{x}_e is important in defining when a unknown network state is *identifiable*, and thus can be learned. To this extent, the following proposition is useful.

Proposition 5.1.2 *Given a routing instance $(\mathcal{N}, \mathbf{d}, \mathbf{c}, \{\gamma_e\}_{e \in E})$, it holds that*

$$\forall e \in E \quad x_e \leq \gamma_e \iff \forall e \in E \quad x_e \leq \hat{x}_e$$

Where \hat{x}_e is defined as above.

The [proof](#) can be found in the appendix.

This proposition clarify that a load is feasible if and only if it respects the upper bound contemporaneously on all its edges.

5.1.2 Feasibility condition and identifiability

At this point, we can characterize the *feasibility* of a flow by considering two conditions: one condition that assures that edge capacities are respected, and a second condition which assures that all the traffic is routed.

The condition on edge capacities is that

$$\forall e \in E, \quad \sum_{P: e \in P} y_P \leq \hat{x}_e,$$

while the complete routing condition is

$$\forall i \in I, \quad \sum_{P: P \in P_i} y_P = d_i.$$

The set of possible flow vectors in which the *capacity conditions* are respected by

$$\mathbf{C} = \left\{ y \in \mathbb{R}_+^{|\mathcal{P}|} : \forall e \in E, \sum_{P: e \in P} y_P \leq \hat{x}_e \right\}$$

and the set in which the *demand conditions* are respected by

$$\mathbf{D} = \left\{ y \in \mathbb{R}_+^{|\mathcal{P}|} : \forall i \in I, \sum_{P: P \in P_i} y_P = d_i \right\}.$$

We can define the set of *feasible flows* as the set in which both conditions are satisfied,

$$\mathcal{Y} = \left\{ y \in \mathbb{R}_+^{|\mathcal{P}|} : \mathbf{C} \cap \mathbf{D} \right\}.$$

From these we can derive the classical feasibility conditions, indeed, summing over all i ,

$$\sum_{i \in I} \sum_{P: P \in P_i} y_P = \sum_{P \in \mathcal{P}} y_P = \sum_i d_i = d_{tot}$$

and

$$\sum_{P: e \in P} y_P = x_e \leq \hat{x}_e \leq \gamma_e.$$

We can now characterize when an identifiable network state.

Definition 5.1.1 (Identifiability) *A random state Θ is identifiable if, for a feasible flow vector $\mathbf{y} \in \mathcal{Y}$,*

$$\forall \theta, \hat{\theta} \in \mathcal{T}, \theta \neq \hat{\theta}, \quad \exists e \in E \text{ s.t. } c_e(x, \theta) \neq c_e(x, \hat{\theta}).$$

In words, a random state is identifiable if for all pairs of possible state realizations there exists an edge such that, for at least one value $\bar{x}_e \leq \hat{x}_e$, we have

$$c_e(\bar{x}_e, \theta) \neq c_e(\bar{x}_e, \hat{\theta}).$$


5.2 Learning failures

In capacity constrained networks, equilibrium loads are not guaranteed to reach the edge capacity for all the edges. Eventually, equilibrium loads cannot even reach the load upper bound on all edges; this may prevent *strong learning*. In particular, if the load value that allows to identify states cannot be reached, then strong learning cannot occur.

We provide some examples in which learning does not occur. In general, learning failure may derive from structural properties of the network or from behavior in equilibrium.

Example 1. Learning failure with some finite capacities


We consider a case in which learning may fail if a network has an explicit capacity only on some of its edges. Consider the following network,



with costs

$$c_1(\cdot) = x \quad c_2(\cdot, \theta) = \begin{cases} x, & \text{if } \theta_1 \\ x + 10, & \text{if } \theta_2 \end{cases} \quad c_3(\cdot) = x$$

in which states can occur with equal probability, $\mu(\theta_1) = \mu(\theta_2) = \frac{1}{2}$. To each edge of the network we associate a triple $(\gamma_e, c_e(\cdot, \mu), \hat{x}_e)$, in which $c_e(\cdot, \mu)$ denotes the expected cost and \hat{x}_e denotes the load upper bound.



This instance has two paths,

$$P_1 = \{e_1, e_3\}, \quad P_2 = \{e_2, e_3\}.$$

The network total capacity is $\gamma = 1$, and for a given value of traffic demand $d \in [0, 1)$ we always have that


$$c_{P_1}(\cdot) < c_{P_2}(\cdot, \mu).$$

Hence, the edge e_2 which allows to indentify states is never reached in equilibrium.

Example 2. Learning failure with all finite capacities

Now, we consider an example of network with limited capacities on all edges, and such that states can be identified for a value $x_e < \hat{x}_e$.

Let's consider the network



with costs

$$c_1(\cdot) = \frac{1}{3-x} \quad c_3(\cdot, \theta) = \begin{cases} \frac{5}{6-x}, & \text{if } \theta_1 \\ \frac{5}{6-x} + 2\frac{(x-1) \cdot I_{\{x>1\}}}{6-x}, & \text{if } \theta_2 \end{cases} \quad c_2(\cdot) = \frac{1}{4-x}$$

and in which an equal prior probability is assigned to both states, $\mu(\theta_1) = \mu(\theta_2) = \frac{1}{2}$. Here, network capacity is $\gamma = 3$ and the expected cost of edge e_3 is

$$c_3(\cdot, \mu) = \frac{5}{6-x} + (x-1) \cdot \frac{I_{\{x>1\}}}{6-x}.$$


To learn the true state we need that a value strictly greater than one is routed through path $P_2 = \{e_1, e_3\}$. Nevertheless, we can notice that even if a flow slightly higher than one crosses edge e_3 , its costs would be even higher than the cost of e_2 loaded with maximal demand, which in this case is three. Indeed,

$$c_3(1 + \epsilon, \mu) = \frac{5}{5-\epsilon} + \frac{\epsilon}{5-\epsilon} > 1 \quad \text{and} \quad c_2(d_{max}) = c_2(3) = 1.$$

Hence, an equilibrium flow greater than the unity will never be routed through path P_2 , and the true state will never be learned.

Example 3. Learning failure with non-identifiable states

Let's consider the following network,



with costs


$$c_1(\cdot) = x \quad c_3(\cdot, \theta) = \begin{cases} \frac{1}{4-x}, & \text{if } \theta_1 \\ \frac{1}{4-x} + (x-3) \cdot I_{\{x>3\}}, & \text{if } \theta_2 \end{cases} \quad c_2(\cdot) = \frac{1}{3-x}$$

In this case, the true state is not learnable because the states are identifiable only for edge loads that exceed the load upper bound. We recall that this scenario is excluded if the state is [identifiable](#).

Example 4. Learning failure in a multi-commodity network with bounded marginals

The aim of this example is to show that even the total demand can cover all the possible values between zero and the network capacity, learning may fail.

Let the network \mathcal{N} be



and suppose that the joint random traffic demand follows

$$p(D_1, D_2) = \begin{cases} \frac{1}{2}, & D_1 \stackrel{\text{a.s.}}{=} 0, \quad D_2 \sim U[0, 6], \\ \frac{1}{2}, & D_1 \sim U[6, 8], \quad D_2 \sim U[0, 6], \\ 0, & \text{otherwise} \end{cases}$$

Here, the total traffic demand can potentially reach any value between zero and the network capacity, since

$$\text{supp}(D_1 + D_2) = [0, 14].$$

Nevertheless, there are some edge loads that will never be reached, and this may prevent learning. In this case, a load on edge e_1 such that $x_1 < 1$ cannot be reached. It follows that if a state can be identified only for a load value $\bar{x}_1 < 1$, then learning fails.

This example highlights that, other than some conditions on structural network capacities, it is needed a condition on the joint random demand.

Example 5. Learning failure in a multi-commodity network with infinite capacities


Now we consider an example of learning failure with two commodities, infinite capacities, and random marginal demands with support that goes from zero to infinity.

Suppose that the two commodities share the same source and destination node, but have access to different sub-networks. Also, suppose that

$$\text{supp}(\mathbf{D}) = \{(D_1, D_2) \in \mathbb{R}_+^2 : D_1 = D_2\},$$

which implies that the marginals are such that $\text{supp}(D_i) = [0, \infty)$.


The underlying network is as follows,



and edge costs are

$$c_1(\cdot) = c_5(\cdot) = x \quad c_2(\cdot) = c_4(\cdot) = 1 + \epsilon x \quad c_3(\cdot, \theta) = \begin{cases} \epsilon x, & \text{if } \theta_1 \\ \epsilon x + 10, & \text{if } \theta_2. \end{cases}$$

Each of the two commodities has access to two of the three available paths: the first commodity can route its traffic through path $P_1 = \{e_1, e_4\}$ and path $P_3 = \{e_1, e_3, e_5\}$, while the second commodity can use paths $P_2 = \{e_2, e_5\}$ and $P_3 = \{e_1, e_3, e_5\}$. This implies that the sub-networks available to the two commodities are both series-parallel,



We can notice that, in equilibrium, the edge e_3 is never used, hence the state is never learned. Indeed, denoting by $d = d_1 = d_2$ a general traffic demand realization, we can notice that in the unique equilibrium the first commodity uses only path P_1 and the second commodity uses only path P_2 , with costs

$$c_{P_1} = c_{P_2} = d + 1 + \epsilon d.$$

Under this flow assignment, path P_3 has cost

$$c_{P_3} = 2d + 5,$$

which is higher than the other costs and hence P_3 is never used.

We can notice that this issue would not emerge if $\text{supp}(\mathbf{D}) = \mathbb{R}_+^2$. In that case, the support comprehends some demand realizations that guarantee the use of path P_3 .

5.2.1 Capacity conservation

We begin this section with a brief intermezzo on network capacities. In what follows γ is the total capacity of a network. It corresponds to the sum of the edge capacities of the smallest *cut* of the network. Specifically, a cut $\mathcal{C} \subseteq E$ in a single source-destination directed multigraph is a set of edges such that it is not possible to route from the source to the destination without passing through \mathcal{C} .

The capacity of a cut $\gamma_{\mathcal{C}}$ is defined additively as the sum of the capacities of the edges that compose it,

$$\gamma_{\mathcal{C}} = \sum_{e \in \mathcal{C}} \gamma_e.$$

The *smallest cut* is the cut with smallest capacity. The capacity of the smallest cut is the network capacity, and denotes the maximum amount of traffic that can be routed through the network.

In the same spirit, a multi commodity network \mathcal{N} is composed by the union of the sub-networks \mathcal{N}_i available to each commodity. Denoting by $I = \{1, \dots, k\}$ the set of commodities and by γ_i the capacity of each sub-network \mathcal{N}_i , it holds that

$$\gamma \leq \sum_{i \in I} \gamma_i,$$

and this comes from the fact that some of the edges may be shared between two sub-networks.

In a multi-commodity network we can refer to a cut as a set of edges such that it is not possible to go from any source to any destination without crossing \mathcal{C} . Again, the capacity of the network corresponds to the capacity of the smallest cut.

We now introduce a network condition such that social learning occurs when a DNRG is played on a capacitated network. From now on, for every internal node $v \in V$ we define the sets of entering edges into v and exiting edges from v respectively as T_v and H_v ,

$$T_v = \{ \text{set of entering edges into node } v \},$$

$$H_v = \{ \text{set of exiting edges from node } v \}.$$

We say that a capacitated network \mathcal{N} respects the *capacity conservation* property if on all internal nodes $v \in V$ it is satisfied

$$\sum_{e \in H_v} \gamma_e = \sum_{e \in T_v} \gamma_e.$$

In words, this means that for each node $v \in V$ that is not the source or the destination, the sum of the capacities of the edges that enter in v is equal to the sum of the capacities of the edges that exit from v .




Figure 5.2: Two examples of capacity conservation in single s-t Networks

Notice that the property of capacity conservation can be respected only if the underlying graph is a directed acyclic graph.

Let us first consider the case in which a routing instance has a single commodity whose traffic demand is d . By construction, this specific network topology is

such that, if the network is fully congested, the load on each edge coincides with the edge capacity. In other words, denoting by γ the total network capacity, if a network \mathcal{N} respects capacity conservation we have that

$$d = \gamma \implies \forall e \in E, \quad x_e = \gamma_e.$$

Another desirable property that comes with this particular construction is that the load upper bound on an edge coincides with the capacity of that edge, i.e.,

$$\forall e \in E, \quad \hat{x}_e = \gamma_e.$$

This can be understood since in general it holds that

$$x_e \leq \hat{x}_e \leq \gamma_e.$$

Indeed,

$$\begin{aligned} d = \gamma &\implies \forall e \in E, \quad x_e = \gamma_e \\ &\implies \forall e \in E, \quad x_e = \hat{x}_e. \end{aligned}$$

That is, if the traffic demand is equal to the network capacity, all edge loads reach their load upper bound.

This is a key desirable property, from which we introduce a weaker form of capacity conservation, called *weak capacity conservation*.

Definition 5.2.1 (Weak capacity conservation) *Let $(\mathcal{N}, d, \mathbf{c}, \{\gamma_e\}_{e \in E})$ be a routing instance and \hat{x}_e be the load upper bound of edge $e \in E$. The network \mathcal{N} respects weak capacity conservation if*

$$d = \gamma \implies \forall e \in E, \quad x_e = \hat{x}_e.$$

In words, an oriented network respects weak capacity conservation if it is endowed with a set of edge capacities such that, when the traffic demand d reaches the total network capacity γ , then the load on each edge x_e reaches the load upper bound \hat{x}_e . Again, we notice that this definition of weak capacity conservation is not compatible with the existence of cycles in the graph.

In our dynamic model with a random demand D , a realization d exactly equal to the network capacity γ happens with probability zero. Nevertheless, since edge costs are strictly increasing, equilibrium loads are continuous in d (see, for example, [Cominetti et al. \(2021\)](#)).

Hence, we can say that

$$\text{supp}(D) = [0, \gamma) \implies \forall e \in E, \quad \text{supp}(x_e^*) = [0, \hat{x}_e),$$

and this will prove to be sufficient to achieve *strong learning*. Here, x_e^* denotes an equilibrium load on edge e , which is a deterministic non-negative function of the demand realization d (again, for this result we refer to [Cominetti et al. \(2021\)](#)).

An example of weak capacity conservation is given by the following network.




Figure 5.3: Weak capacity conservation in a single-commodity network.

A similar argument holds for multi-commodity instances. However, edges that are available to different commodities deserve a careful consideration.

In particular, we consider the case in which commodities whose sub-networks share a common edge have the right of a proportion of the capacity on the edge shared. The predetermined allocation of the common edges allows to determine the capacity γ_i of the sub-network \mathcal{N}_i available to each commodity $i \in I$. Recall that each random demand D_i is bounded above by γ_i .

Under this construction, *capacity conservation* implies that

$$\gamma = \sum_{i \in I} \gamma_i,$$

in which γ is the total network capacity and γ_i is the capacity of the sub-network available to commodity i .

Thus, in a multi commodity network we can generalize the property of weak conservation of capacities as

$$d_i = \gamma_i \quad \forall i \in I \quad \implies \quad x_e = \hat{x}_e \quad \forall e \in E.$$

Again by continuity of each edge load in the traffic demands, we have

$$\text{supp}(\mathbf{D}) = \times_{i \in I} [0, \gamma_i] \quad \implies \quad \text{supp}(x_e^*) = [0, \hat{x}_e] \quad \forall e \in E.$$

An example of weak capacity conservation in the case of multiple commodities follows.

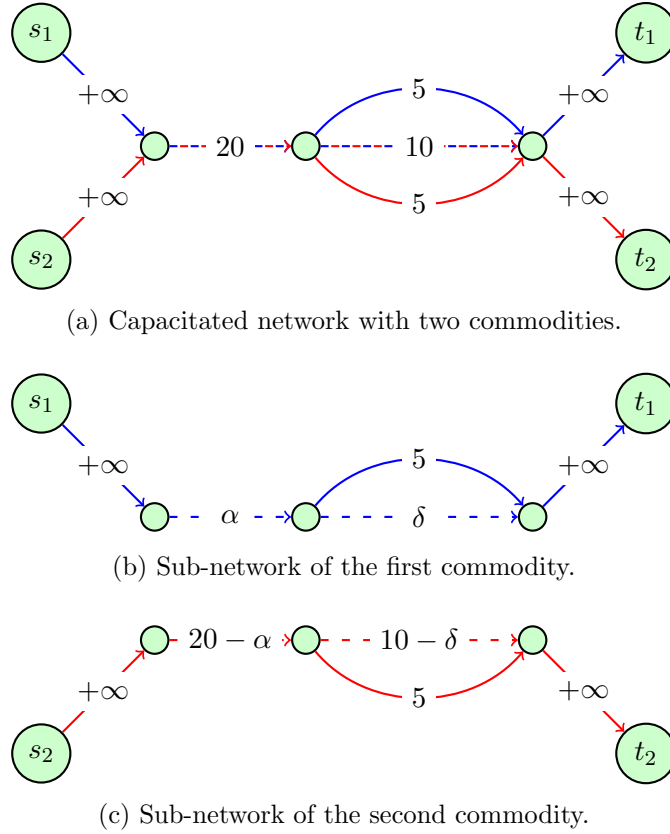


Figure 5.4: Multi-commodity network (a) that respects weak capacity conservation and in which edges are labeled with their capacities. (b) and (c) are the sub-networks available to the two commodities.

5.2.2 SP networks and generalization

Now we focus on networks with possibly infinite edge capacities. We start with the definition of the class of single-commodity *series-parallel* networks. This particular network topology respects desirable properties (Duffin (1965), Milchtaich (2006), Chen et al. (2016)) which result useful to exclude some occurrences of learning failure (for some examples we refer to Macaul et al. (2022)). We define the class of series-parallel network recursively as follow.

Definition 5.2.2 (Series-parallel network) *A single origin-destination network \mathcal{N} is called series-parallel (SP) if it can be defined sequentially as follows:*

1. *Either \mathcal{N} has a single edge,*
2. *or \mathcal{N} consists of two SP networks connected in series, by merging the destination of the first with the origin of the second,*

3. or \mathcal{N} consists of two SP networks connected in parallel, by merging the origin of the first with the origin of the second and the destination of the first with the destination of the second.

The key aspect is that if a network is not series-parallel, then it contains some sub-networks such that some edges are not used in equilibrium even when the demand is high.




Figure 5.5: Network (A) is SP. Network (B) is not SP due to the red edge.

Recall that \mathcal{P}_i stands for the set of paths available to commodity i in a instance with multiple commodities. Now, we can characterize a multi-commodity instance that respects similar desirable properties by considering the case in which each set \mathcal{P}_i form itself a series-parallel network. We call the resulting network U-SP, since the set of all paths is defined as the union of the paths available to each commodity, which are SP.

In what follows each SP network has an index, i , which represents one commodity in the non-atomic routing game played on the U-SP. Let $\mathcal{S}(SP)$ denote the set of all possible SP networks.

Definition 5.2.3 A multigraph \mathcal{N} with a set $I = \{1, \dots, i, \dots, k\}$ of origin-destination pairs is U-SP if

$$\forall i \in I, \quad \mathcal{P}_i \in \mathcal{S}(SP)$$

in which \mathcal{P}_i is the set of paths available to commodity i .

This means that each commodity routes its traffic on a series-parallel network.

5.3 Learning in capacity constrained networks

In this and in the following paragraph we devote our attention to the conditions under which social learning occurs.

Recall that the players start in the first period with a prior distribution on the network state, which is common knowledge. In each stage a random demand is

realized and the equilibrium flow is played; all the new information about the experienced costs is immediately broadcast to all players, who update their common belief according to the Bayes' rule. Strong learning occurs when the posterior public belief converges almost surely to the degenerate distribution which assigns positive probability only to the true state.

Before stating the main theorem, we highlight the importance of costs assumptions. In our model, edge costs are continuous and strictly increasing: this will guarantee that each edge load is an unbounded function of the traffic demands. Moreover, we will assume that edge costs go to infinity as the load reaches the edge capacity. That is, for all states $\theta \in \mathcal{T}$, we assume that

$$\lim_{x_e \rightarrow \gamma_e} c_e(x_e, \theta) = +\infty. \quad (5.1)$$

If this is not the case, equilibrium costs may not be unique in a capacitated instance (Correa et al. (2004)). This assumption guarantees uniqueness since any such capacitated instance is essentially the same of a non-atomic instance with infinite capacities and a specific cost structure. In other words, under this assumption edge capacities are encompassed in the cost functions.

Assumptions 1-3:

1. edge costs are strictly increasing in the load for all possible states;
2. edge costs are continuous in the load for all possible states;
3. the following limit condition holds,

$$\lim_{x_e \rightarrow \gamma_e} c_e(x_e, \theta) = +\infty \quad \forall \theta \in \mathcal{T}.$$

Now we can state the main results of social learning in capacity constrained networks. In what follows, Γ is a *dynamic non-atomic routing game* (DNRG) as presented here, an unknown network state is identifiable according to Definition 5.1.1 and weak capacity conservation follows Definition 5.2.1.

Theorem 5.3.1 (Learning in capacity constrained networks) *Let Γ be a single commodity DNRG with identifiable unknown network state such that the network \mathcal{N} respects weak capacity conservation and Assumptions 1-3 hold.*

If

$$\text{supp}(D) = [0, \gamma),$$

then Strong Learning occurs.

The proof can be found in the appendix.

The key idea of the proof is that, under such conditions, for any possible edge load x_e between zero and the load upper bound there exists a realization of the demand that generates it. By continuity of the costs in the loads and of the loads in the demands, it is enough to reach a value that belongs to a small neighborhood of the required demand that guarantees a change in the belief. Thanks to the condition on the support of D , any such neighborhood has positive probability.

A general result holds for the multi-commodity case. A major remark is that it is not sufficient that the total traffic demand $\sum_{i \in I} D_i$ has as support the set $[0, \gamma]$. A case of learning failure is shown in [Example 4](#).

Theorem 5.3.2 *Let Γ be a multi-commodity DNRG with identifiable unknown network state such that the network \mathcal{N} respects weak capacity conservation and [assumptions 1-3](#) hold.*

If

$$\text{supp}(\mathbf{D}) = \times_i [0, \gamma_i),$$

then Strong Learning occurs.

The [proof](#) can be found in the appendix.

5.4 Learning in networks with infinite capacities

The work of [Macault et al. \(2022\)](#) covers the single origin-destination case with infinite capacities. The authors proved the following result.

Theorem 5.4.1 *Let Γ be a single commodity DNRG with identifiable unknown network state such that the network \mathcal{N} is SP, edges have infinite capacities and [Assumptions 1-3](#) hold.*

If

$$\text{supp}(D) = [0, +\infty),$$

then Strong Learning occurs.

A similar result holds for instances with more than one commodity.

Theorem 5.4.2 *Let Γ be a multi-commodity DNRG with identifiable unknown network state such that the network \mathcal{N} is U-SP, edges have infinite capacities and [Assumptions 1-3](#) hold.*

If

$$\text{supp}(\mathbf{D}) = \mathbb{R}_+^k,$$

then Strong Learning occurs.

The [proof](#) can be found in the appendix.

The learning procedure is the same as in the single origin case, with equilibrium edge loads are that are continuous in the demands and unbounded.

A noteworthy and non-trivial aspect that emerges only in multi-commodity instances is that edge loads are not monotone in traffic demands (see, for example, [Fisk \(1979\)](#)). Nevertheless, thanks to the *series-parallel* structure of the sub-network available to each commodity, edge loads are still unbounded.

Moreover, it is important that the support of the demand vector \mathbf{D} is \mathbb{R}_+^k , since the condition $\text{supp}(D_i) = [0, \gamma_i)$ for all $i \in I$ may not be sufficient (as shown in [Example 5](#)).

Chapter 6

Conclusion

In this work we study some of the most recent results concerning the traffic assignment problems, which can be modelled as a non-cooperative game played on a network. In a selfish routing game an amount of traffic inflow has to be routed from some sources to some destinations of the underlying directed graph; the traffic assignment on the paths of the network is called *flow*.

In the first part of the thesis we study routing games under complete information and introduce the ideas of equilibrium flows, optimal flows, and efficiency of routing instances. Then, we study some models that arise under uncertainty: several sources of randomness can emerge in real traffic situations and have been theoretically studied before, such as in the level of traffic demands, or in the costs experienced by the players. We consider routing instances with an unknown network state that influences edge costs, in which the traffic is routed by minimizing the expected cost of each path. Learning the true underlying state may be possible when this kind of incomplete information games are repeatedly played with a random level of demand in each stage game. Specifically, there are two possible kinds of social learning in non-atomic routing games: *strong learning*, when the true network state is identified, and *weak learning*, when players learn how to play as if the true state were known.

In the second part of the thesis we define when a network state is *identifiable* and show under which conditions strong learning occurs. With assumptions of continuity, monotonicity, and unboundeness of cost functions, we show that the learning conditions differ between capacitated instances and instances with infinite edge capacities. In capacitated instances a sufficient condition to achieve learning is that, under full congestion, edge loads reach their upper bound. On the other hand, in infinite capacity instances it is sufficient that the sub-network available to each commodity is *series-parallel*.

6.1 Future Research

A crucial hidden assumption of the learning procedure considered in this thesis is that in each stage game the traffic is routed according to a Wardrop Equilibrium. However, this solution concept may not match real traffic outcomes, and may be unnatural in some applications: for example, paths may be too long to determine reasonably the traffic assignment at the beginning of the game. On the other hand, it may be more realistic that players decide sequentially on edges, and not on paths. Also, users may experience some forms of discounting on the costs experienced in each edge. To this end, the broad literature on Dynamic Discrete Choice may provide several insights, since the choice on edges would be discrete and the traffic assignment may be modelled as a finite horizon discrete choice problem. An associated research question may concern the conditions under which learning occurs if alternative solution concepts are employed.

A second research direction may derive from the way in which the network uncertainty is modelled and how, eventually, some inference on this uncertainty can be drawn. In this model we consider an unknown network state which is drawn before the beginning of the history and remains fixed. We could consider a different model in which a network state is drawn independently in each stage game. If this happens, it may be possible to understand something about the distribution from which the network state is drawn.

Bibliography

- Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. Informational braess' paradox: The effect of information on traffic congestion. *Operations Research*, 66(4):893–917, 2018.
- Itai Arieli and Manuel Mueller-Frank. A general analysis of sequential social learning. *Mathematics of Operations Research*, 46(4):1235–1249, 2021.
- Robert J Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.
- Robert J Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, pages 1–18, 1987.
- Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated games with incomplete information*. MIT press, 1995.
- Abhijit V Banerjee. A simple model of herd behavior. *The quarterly journal of economics*, 107(3):797–817, 1992.
- Martin Beckmann, Charles B McGuire, and Christopher B Winsten. Studies in the economics of transportation. *Cowles Commission for Research in Economics by Yale University Press*, 1956.
- Petra Berenbrink and Oliver Schulte. Evolutionary equilibrium in bayesian routing games: Specialization and niche formation. *Theoretical computer science*, 411(7-9):1054–1074, 2010.
- Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016a.
- Dirk Bergemann and Stephen Morris. Information design, bayesian persuasion, and bayes correlated equilibrium. *American Economic Review*, 106(5):586–91, 2016b.
- Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992.
- John Adrian Bondy and Uppaluri Siva Ramachandra Murty. *Graph theory*. Springer Publishing Company, Incorporated, 2008.
- Benjamin Brooks, Alexander P Frankel, and Emir Kamenica. Information hierarchies. *Available at SSRN 3448870*, 2019.
- Xujin Chen, Zhuo Diao, and Xiaodong Hu. Network characterizations for excluding braess’s paradox. *Theory of Computing Systems*, 59(4):747–780, 2016.
- George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 67–73, 2005.
- Riccardo Colini-Baldeschi, Roberto Cominetti, Panayotis Mertikopoulos, and Marco Scarsini. The asymptotic behavior of the price of anarchy. In *International Conference on Web and Internet Economics*, pages 133–145. Springer, 2017.
- Roberto Cominetti. Equilibrium routing under uncertainty. *Mathematical Programming*, 151(1):117–151, 2015.
- Roberto Cominetti and Alfredo Torrico. Additive consistency of risk measures and its application to risk-averse routing in networks. *Mathematics of Operations Research*, 41(4):1510–1521, 2016.
- Roberto Cominetti, Marco Scarsini, Marc Schröder, and Nicolas E Stier-Moses. Price of anarchy in stochastic atomic congestion games with affine costs. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 579–580, 2019.
- Roberto Cominetti, Marco Scarsini, Marc Schröder, and Nicolás Stier-Moses. Convergence of large atomic congestion games. *arXiv preprint arXiv:2001.02797*, 2020.
- Roberto Cominetti, Valerio Dose, and Marco Scarsini. The price of anarchy in routing games as a function of the demand. *Mathematical Programming*, pages 1–28, 2021.
- José R Correa and Nicolás E Stier-Moses. Wardrop equilibria. *Encyclopedia of Operations Research and Management Science*. Wiley, 2011.
- José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.
- José R Correa, Andreas S Schulz, and Nicolás E Stier-Moses. A geometric approach

- to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64(2):457–469, 2008.
- Richard J Duffin. Topology of series-parallel networks. *Journal of Mathematical Analysis and Applications*, 10(2):303–318, 1965.
- Michael J. Evans and Jeffrey S. Rosenthal. Probability and statistics: The science of uncertainty. *The American Statistician*, 59:276–276, 2005.
- Caroline Fisk. More paradoxes in the equilibrium assignment problem. *Transportation Research Part B: Methodological*, 13(4):305–309, 1979.
- Lester Randolph Ford and Delbert Ray Fulkerson. Flows in networks. In *Flows in Networks*. Princeton University Press, 1962.
- Françoise Forges. Games with incomplete information: from repetition to cheap talk and persuasion. *Annals of Economics and Statistics*, pages 3–30, 2020.
- Drew Fudenberg, Fudenberg Drew, David K Levine, and David K Levine. *The theory of learning in games*, volume 2. MIT press, 1998.
- Martin Gairing, Burkhard Monien, and Karsten Tiemann. Selfish routing with incomplete information. *Theory of Computing Systems*, 42(1):91–130, 2008.
- Jacob K Goeree, Thomas R Palfrey, and Brian W Rogers. Social learning with private and common values. *Economic Theory*, 28(2):245–264, 2006.
- Allan Gut. *Probability: a graduate course*, volume 200-5. Springer, 2005.
- Saeed Hadikhaneloo, Rida Laraki, Panayotis Mertikopoulos, and Sylvain Sorin. Learning in nonatomic games, part i: Finite action spaces and population games. *arXiv preprint arXiv:2107.01595*, 2021.
- John M Harris, Jeffrey L Hirst, and Michael J Mossinghoff. *Combinatorics and graph theory*. Springer Publishing Company, Incorporated, 2008.
- Alain Haurie and Patrice Marcotte. On the relationship between nash—cournot and wardrop equilibria. *Networks*, 15(3):295–308, 1985.
- Ron Holzman and Nissan Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, 21(1-2):85–101, 1997.
- Ehud Kalai and Ehud Lehrer. Rational learning leads to nash equilibrium. *Econometrica*, 61(5):1019–1045, 1993. ISSN 00129682, 14680262.
- Ehud Kalai and Ehud Lehrer. Weak and strong merging of opinions. *Journal of Mathematical Economics*, 23(1):73–86, 1994.
- Emir Kamenica. Bayesian persuasion and information design. *Annual Review of Economics*, 11:249–272, 2019.

- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Anna R Karlin and Yuval Peres. *Game theory, alive*, volume 101. American Mathematical Soc., 2017.
- Jon Kleinberg and Eva Tardos. *Algorithm design*. Pearson Education India, 2006.
- Max Klimm and Philipp Warode. Computing all wardrop equilibria parametrized by the flow demand. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 917–934. SIAM, 2019.
- Frederic Koessler, Marco Scarsini, and Tristan Tomala. Information design in large games. *arXiv preprint arXiv:2107.06312*, 2021.
- Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Annual symposium on theoretical aspects of computer science*, pages 404–413. Springer, 1999.
- Emilien Macault, Marco Scarsini, and Tristan Tomala. Social learning in nonatomic routing games. *Games and Economic Behavior*, 2022.
- Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Number 9780195102680 in OUP Catalogue. Oxford University Press, 1995.
- Laurent Mathevet, Jacopo Perego, and Ina Taneva. On information design in games. *Journal of Political Economy*, 128(4):1370–1404, 2020.
- Nimrod Megiddo. Optimal flows in networks with multiple sources and sinks. *Mathematical Programming*, 7(1):97–107, 1974.
- Panayotis Mertikopoulos. *Online optimization and learning in games: Theory and applications*. PhD thesis, Grenoble 1 UGA-Université Grenoble Alpes, 2019.
- Igal Milchtaich. Network topology and the efficiency of equilibrium. *Games and Economic Behavior*, 57(2):321–346, 2006.
- Dov Monderer and Lloyd S Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- Fernando Ordóñez and Nicolás E Stier-Moses. Wardrop equilibria with risk-averse users. *Transportation Science*, 44(1):63–86, 2010.
- Christos Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 749–753, 2001.
- Robert W Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

- Tim Roughgarden. *Selfish routing*. Cornell University, 2002.
- Tim Roughgarden. Routing games. *Algorithmic game theory*, 18:459–484, 2007.
- Tim Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, 2016. doi: 10.1017/CBO9781316779309.
- Tim Roughgarden and Eva Tardos. Introduction to the inefficiency of equilibria. *Algorithmic game theory*, 17:443–459, 2007.
- William H Sandholm. Potential games with continuous player sets. *Journal of Economic Theory*, 97(1):81–108, 2001.
- Marco Scarsini and Tristan Tomala. Repeated congestion games with bounded rationality. *International Journal of Game Theory*, 41(3):651–669, 2012.
- Carl P. Simon and Lawrence Edward Blume. *Mathematics for Economists*. W. W. Norton and Company, 1994.
- Lones Smith and Peter Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.
- Hamidreza Tavafoghi and Demosthenis Teneketzis. Informational incentives for congestion games. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1285–1292. IEEE, 2017.
- Hamidreza Tavafoghi and Demosthenis Teneketzis. Strategic information provision in routing games, 2019.
- John Glen Wardrop. Road paper. some theoretical aspects of road traffic research. *Proceedings of the institution of civil engineers*, 1(3):325–362, 1952.
- John Glen Wardrop and James Ivor Whitehead. Correspondence. some theoretical aspects of road traffic research. *Proceedings of the institution of civil engineers*, 1(5):767–768, 1952.
- David P Williamson. *Network flow algorithms*. Cambridge University Press, 2019.
- Manxi Wu and Saurabh Amin. Learning an unknown network state in routing games. *IFAC-PapersOnLine*, 52(20):345–350, 2019.
- Manxi Wu, Jeffrey Liu, and Saurabh Amin. Informational aspects in a class of bayesian congestion games. In *2017 American Control Conference (ACC)*, pages 3650–3657. IEEE, 2017.
- Manxi Wu, Saurabh Amin, and Asuman E Ozdaglar. Value of information in bayesian routing games. *Operations Research*, 69(1):148–163, 2021.
- Yixian Zhu and Ketan Savla. A computational approach for information design in non-atomic routing games, 2020.

Appendix

Proofs

Proposition 5.1.1.: *Given a routing instance $(\mathcal{N}, \mathbf{d}, \mathbf{c}, \{\gamma_e\}_{e \in E})$, feasible edge loads are upper bounded as follows,*

$$\forall e \in E \quad x_e \leq \hat{x}_e = \min \left\{ \min_{e \in E_e} \gamma_e, \sum_{P: e \in P} \min_{e \in P} \gamma_e \right\}$$

In which for a given edge $e \in E$, $E_e \subseteq E$ is defined recursively in two steps,

1. $e \in E_e$
2. $e' \in E_e$ iff for some $e \in E_e$, $H_e = \{e'\}$ or $T_e = \{e'\}$.

Proof. We show that any load profile such that $x_e > \hat{x}_e$ for at least for one edge violates the feasibility constraint. By contradiction, suppose that for at least one edge it holds that

$$x_e > \min \left\{ \min_{e \in E_e} \gamma_e, \sum_{P: e \in P} \min_{e \in P} \gamma_e \right\}. \quad (1)$$

In order, we consider two cases: when the first element of the minimum is not smaller than the second, and then when it is strictly smaller.

- First, let

$$\min_{e \in E_e} \gamma_e \leq \sum_{P: e \in P} \min_{e \in P} \gamma_e.$$

This means that (1) can be written as

$$x_e > \min_{e \in E_e} \gamma_e,$$

by the conservation of flow and given the structure of E_e it means there there exists an edge $e' \in E_e$ such that

$$x_{e'} > \gamma_{e'},$$

which violates feasibility.

• On the other hand, let

$$\min_{e \in E_e} \gamma_e > \sum_{P: e \in P} \min_{e \in P} \gamma_e.$$

So, our contradictory hypothesis (1) is

$$x_e > \sum_{P: e \in P} \min_{e \in P} \gamma_e.$$

By the definition of load on an edge, the previous inequality can be written as

$$\begin{aligned} \sum_{P: e \in P} y_P &> \sum_{P: e \in P} \min_{e \in P} \gamma_e, \\ \implies \exists P \in \mathcal{P} \quad s.t. \quad y_P &> \min_{e \in P} \gamma_e \\ \implies \exists e \in P \quad s.t. \quad x_e = \sum_{P: e \in P} y_P &> \gamma_e, \end{aligned}$$

which is against feasibility. \square

Proposition 5.1.2 *Given a routing instance $(\mathcal{N}, \mathbf{d}, \mathbf{c}, \{\gamma_e\}_{e \in E})$, it holds that*

$$\forall e \in E \quad x_e \leq \gamma_e \iff \forall e \in E \quad x_e \leq \hat{x}_e$$

Where \hat{x}_e is defined as above.

Proof. First, we can notice that

$$\forall e \in E, \quad x_e \leq \hat{x}_e \implies \forall e \in E, \quad x_e \leq \gamma_e.$$

Indeed, it is sufficient to notice that

$$\min_{e' \in E_e} \gamma_{e'} \leq \gamma_e.$$

Then,

$$x_e \leq \hat{x}_e \leq \gamma_e.$$

Secondly, we can show that

$$\forall e \in E, \quad x_e \leq \hat{x}_e \iff \forall e \in E, \quad x_e \leq \gamma_e.$$

By contradiction, suppose that for at least one edge we have

$$x_e > \hat{x}_e = \min \left\{ \min_{e' \in E_e} \gamma_{e'}, \sum_{P: e \in P} \min_{e \in P} \gamma_e \right\}.$$

By *Proposition 5.1.1* this implies that on at least one edge $e \in E$ it holds

$$x_e > \gamma_e,$$

which concludes the proof. \square

Remark *Costs assumptions:*

1. edge costs are strictly increasing in the load for all possible states;
2. edge costs are continuous in the load for all possible states;
3. the following limit condition holds,

$$\lim_{x_e \rightarrow \gamma_e} c_e(x_e, \theta) = +\infty \quad \forall \theta \in \mathcal{T}.$$

Theorem 5.3.1 *Let Γ be a single commodity DNRG with identifiable unknown network state such that the network \mathcal{N} respects weak capacity conservation and [assumptions 1-3](#) hold.*

Then, if

$$\text{supp}(D) = [0, \gamma),$$

Strong Learning occurs.

Proof. To prove this theorem, we go through the steps of the original theorem of [Macault et al. \(2022\)](#). First, given a public belief μ^t , we define the set $L(\mu^t)$ as

$$L(\mu^t) = \{d \in \mathbb{R} : \mu^t \neq \mu^{t+1} \text{ if } D^t = d\}.$$

This means that $L(\mu^t)$ is the set of demand realizations such that, under the resulting equilibrium load profile \mathbf{x}^{*t} , the posterior μ^{t+1} differs from the posterior μ^t . Moreover, if $\mu^t \neq \mu^\infty$, then this set is non-empty. In other words, this means that whenever

$$\mu^t \neq \delta_\Theta,$$

there are at least two values θ_1 and θ_2 such that

$$\mu^t(\theta_1) \in (0, 1) \quad \text{and} \quad \mu^t(\theta_2) \in (0, 1).$$

This implies, by identifiability of states, that there exists an edge $e \in E$ such that for those θ_1 and θ_2 we have that

$$c_e(\cdot, \theta_1) \neq c_e(\cdot, \theta_2)$$

for some load $x_e \leq \hat{x}_e$. Here, \hat{x}_e represents the maximum possible load achievable on edge e . Let \bar{x}_e be a load value that guarantees identifiability, i.e. such that

$$c_e(\bar{x}_e, \theta_1) \neq c_e(\bar{x}_e, \theta_2).$$

It is known from [Cominetti et al. \(2021\)](#) that if the costs functions are strictly increasing¹, then equilibrium loads are continuous in the demand.

¹It is important to notice the original result of [Cominetti et al. \(2021\)](#) applies to networks with infinite capacities and strictly increasing cost functions. Here, since the limit condition on edge costs hold, the capacitated instance is substantially equivalent to a instance with infinite capacities and suitable costs functions.

This implies that there exists a demand interval \mathcal{D} such that

$$\forall d^t \in \mathcal{D}, \quad x_e^* \in I(\bar{x}_e).$$

Moreover, since the network respects weak capacity conservation, we have that

$$\text{supp}(D) = [0, \gamma) \implies \forall e \in E, \forall x_e \in [0, \hat{x}_e), \Pr(x_e^* \in I(x_e)) > 0.$$

Now, by continuity of $c_e(\cdot, \theta) \forall \theta \in \mathcal{T}$, and by the identifiability condition we have that

$$c_e(x_e, \theta_1) \neq c_e(x_e, \theta_2), \quad \forall x_e \in I(\bar{x}_e).$$

A load value in a neighborhood of \bar{x}_e guarantees identifiability of states, and when such a value is reached the common belief changes. Indeed, one of the two costs is realized and

$$\mu^t(\theta_1) = 0 \text{ or } \mu^t(\theta_2) = 0.$$

Since $\text{supp}(D) = [0, \gamma)$, then the demand interval \mathcal{D} has positive probability, and

$$P(D^t \in \mathcal{D} \text{ for some } t \in \mathbb{N}) = 1.$$

This argument is repeated until all states are excluded except for the true one. Thus, the posterior probability does not change anymore and coincides with the degenerate distribution which assign all the probability mass to the true state,

$$\mu^t = \delta_\theta.$$

This means that, for all possible realizations, after the realization of a state the social belief converges to the degenerate Dirac measure. This also means that, before the realization of the state θ , there exists a random time τ such that $\mu_t = \delta_\Theta$ for all $t \geq \tau$, and thus strong learning occurs. \square

Theorem 5.3.2 *Let Γ be a multi-commodity DNRG with indentifiable unknown network state such that the network \mathcal{N} respects weak capacity conservation and [assumptions 1-3](#) hold.*

Then, if

$$\text{supp}(D) = \times_i [0, \gamma_i),$$

Strong Learning occurs.

Proof. Again, let μ^t be the posterior public belief, and let $L(\mu^t)$ be the set of demand realizations $\mathbf{d} = \{d_i\}_{i \in I}$ such that under an equilibrium load profile \mathbf{x}^{*t} it holds that,

$$\mu^t \neq \mu^{t+1} \quad \text{if} \quad \begin{bmatrix} D_1^t \\ \vdots \\ D_k^t \end{bmatrix} = \mathbf{D}^t = \mathbf{d} = \begin{bmatrix} d_1 \\ \vdots \\ d_k \end{bmatrix}.$$

If $\mu^t \neq \mu^\infty$, then $L(\mu^t) \neq \emptyset$ and there exist two values $\theta_1, \theta_2 \in \mathcal{T}$ that have positive probability according to the prior belief, i.e.,

$$\begin{cases} \mu^t(\theta_1) \in (0, 1) \\ \mu^t(\theta_2) \in (0, 1) \end{cases}.$$

Moreover, by identifiability of states, there exists an edge $e \in E$ and a load value $x_e < \hat{x}_e$ such that

$$c_e(x_e, \theta_1) \neq c_e(x_e, \theta_2).$$

This means there there exists at least one load value x_e such that the cost of that edge is different in the two states; we call this value \bar{x}_e , thus

$$c_e(\bar{x}_e, \theta_1) \neq c_e(\bar{x}_e, \theta_2).$$

We need to show there exists a value $\mathbf{d} = \{d_i\}_{i \in I}$ such that the equilibrium load x_e^* on an edge e induced by \mathbf{d} belongs to a neighborhood of \bar{x}_e .

Here, equilibrium costs functions are continuous in the demand, and also equilibrium loads are continuous in the demand. Moreover, under weak capacity conservation we have that

$$\text{supp}(\mathbf{D}) = \times_i [0, \gamma_i) \quad \implies \quad \forall e \in E, \forall x_e \in [0, \hat{x}_e), \Pr(x_e^* \in I(x_e)) > 0.$$

In words, if

$$\text{supp}(\mathbf{D}) = \times_i [0, \gamma_i),$$

then, for all edges, all values of edge loads between zero and the upper bound are contained in a neighborhood which happens with positive probability.

By the previous argument, there exist an neighborhood of the demand vector, say $\tilde{\mathbf{D}}$, which induces an equilibrium load x_e^* that belongs to a neighborhood of \bar{x}_e . That is, $\tilde{\mathbf{D}}$ is such that, for all $\mathbf{d}^t \in \tilde{\mathbf{D}}$, the expected cost on edge e is

$$c_e(x_e^*, \mu^t), \quad x_e^* \in I(\bar{x}_e).$$

We recall that only the true state determines the cost that is experienced by the players - and does not usually coincide with the expected cost that determines the

equilibrium behavior. By continuity of $c_e(\cdot, \theta)$ for all $\theta \in \mathcal{T}$ and by the identifiability condition we have that

$$c_e(x_e, \theta_1) \neq c_e(x_e, \theta_2), \quad \forall x_e \in I(\bar{x}_e).$$

Then, whenever a realized value of the demands is $\mathbf{d}^t \in \tilde{\mathbf{D}}$, only one of the two costs is experienced by the players, and the state which corresponds to the not realized cost is removed from the support of the posterior belief. This implies that the common belief is updated such that

$$\mu^t(\theta_1) = 0 \quad \text{or} \quad \mu^t(\theta_2) = 0.$$

Then, since the demand vector has full support,

$$P\left(\mathbf{d}^t \in \tilde{\mathbf{D}}, \text{ for some } t \in N\right) = 1,$$

and this concludes the proof because with probability one the beliefs are updated until the true state is the only state with positive probability. \square

Theorem 5.4.2 *Let Γ be a multi-commodity DNRG with identifiable unknown network state such that the network \mathcal{N} is U-SP, edges have infinite capacities and assumptions 1-3 hold.*

Then, if

$$\text{supp}(\mathbf{D}) = \mathbb{R}_+^k,$$

Strong Learning occurs.

Proof. Let μ^t be the posterior public belief at time t and $L(\mu^t)$ be the set of demand realizations such that under the equilibrium load profile \mathbf{x}^{*t} the posterior μ^{t+1} differs from μ^t .

If $\mu^t \neq \mu^\infty$, then $L(\mu^t) \neq \emptyset$ and there exist two values $\theta_1, \theta_2 \in \mathcal{T}$ that have positive probability according to the prior belief. Moreover, by identifiability of states, there exists an edge $e \in E$ and a load value $x_e \leq \hat{x}_e$ such that

$$c_e(x_e, \theta_1) \neq c_e(x_e, \theta_2).$$

This means that there exists at least one value x_e such that the cost of that edge is different in the two states; let \bar{x}_e be this value, hence

$$c_e(\bar{x}_e, \theta_1) \neq c_e(\bar{x}_e, \theta_2).$$

Now, we show there exists a value $\tilde{\mathbf{d}} = \{\tilde{d}_i\}_{i \in I}$ such that the equilibrium load x_e^* on an edge e induced by $\tilde{\mathbf{d}}$ belongs to a neighborhood of \bar{x}_e . Then, since costs functions are continuous, we notice that to attain learning it is sufficient that the traffic

demand reaches a value in a neighborhood of $\tilde{\mathbf{d}}$; finally, since the random demand vector \mathbf{D} has full support, this neighborhood is reached with positive probability as the game is repeated.

Thanks to the *series-parallel* structure of the sub-network accessible to each commodity for every edge $e \in E$ the equilibrium load x_e^* is unbounded. This is a known result of single commodity instances (see [Macault et al. \(2022\)](#)), and it holds in also with several commodity - even though in this case equilibrium loads are not generally monotone in traffic demands (see, for example, [Fisk \(1979\)](#)). It generalizes because the single commodity case is encompassed in the support of the demand vector of the multi-commodity case: we can consider, without loss of generality, a variable demand $d_1 \in \mathbb{R}$ and notice that $(d_1, 0, \dots, 0) \in \text{supp}(\mathbf{D}) = \mathbb{R}_+^k$.

Hence, there exist an neighborhood, say $\tilde{\mathbf{D}}$, which induces an equilibrium load x_e^* that belongs to a neighborhood of \bar{x}_e . That is, $\tilde{\mathbf{D}}$ is such that, for all $\mathbf{d}^t \in \tilde{\mathbf{D}}$, continuity of costs and identifiability of states guarantee that in equilibrium

$$c_e(x_e^*, \theta_1) \neq c_e(x_e^*, \theta_2), \quad \forall x_e^* \in I(\bar{x}_e).$$

Then, whenever a realized value of the demands is $\mathbf{d}^t \in \tilde{\mathbf{D}}$, only one of the two costs is experienced by the players, and the common belief is updated such that

$$\mu^t(\theta_1) = 0 \quad \text{or} \quad \mu^t(\theta_2) = 0.$$

Then, since the demand vector has full support,

$$P\left(\mathbf{d}^t \in \tilde{\mathbf{D}}, \text{ for some } t \in N\right) = 1.$$

Since the game is infinitely repeated, a value in this neighborhood will be reached almost surely, and the belief distribution will be updated until only the true network state will be assigned positive probability. This implies strong learning, because the posterior belief will coincide with the degenerate distribution which assigns positive probability only to the true state. \square

List of Symbols

- $I = \{1, \dots, k\}$: set of commodities
- \mathcal{P}_i : paths available to commodity i
- $\mathcal{P} = \cup_i \mathcal{P}_i$: set of all available paths
- \mathcal{N} : oriented graph with a set of vertices V and edges E
- $\mathbf{d} = \{d_i\}_{i \in I}$: traffic demand vector
- $\mathbf{c} = \{c_e\}_{e \in E}$: vector of edge cost functions
- $G = (\mathcal{N}, \mathbf{d}, \mathbf{c}, \{\gamma_e\}_{e \in E})$: deterministic Routing Instance with finite edge capacities
- $G = (\mathcal{N}, \mathbf{d}, \mathbf{c})$: deterministic Routing Instance with infinite edge capacities
- $\mathbf{y} = \{y_P\}_{P \in \mathcal{P}}$: flow vector
- y_P : flow on path P
- x_e : load on edge e
- \mathbf{y}^* : flow vector at Wardrop Equilibrium
- ψ : Potential Function (in Potential Games)
- γ_e : capacity of edge e
- SP_i : series-parallel network indexed on commodity i
- $U\text{-}SP_I$: U-SP network indexed on commodities $I = \{1, \dots, k\}$
- $\mathcal{S}(SP)$: set of series-parallel networks
- \bar{x}_e : load value that guarantees a change in beliefs in a Dynamic Non-atomic Routing game
- \hat{x}_e : load upper bound on edge e
- \mathbf{x}^* : equilibrium load profile

- \mathcal{T} : finite set of states of nature
- $G_\mu = [G, (\mathcal{T}, 2^{\mathcal{T}}, \mu)]$: non-atomic routing game with unknown network state
- \mathbb{P} : product measure generated by the marginal distribution F of the traffic demand and the prior μ
- $\Gamma = (G, \mathbb{P})$: dynamic non-atomic routing game with unknown network state and random demands
- μ^t : posterior belief a time t
- \mathcal{C} : cut in an oriented multigraph
- (\mathbf{T}, π) : information structure

Acknowledgment

I would like to express my deepest gratitude and appreciation to Professor Marco Scarsini for being an ideal teacher, mentor, and thesis supervisor. Thank you for the precious insight and continuous guidance.

I am also thankful to Professor Fabio Antonelli for being my co-supervisor during the draft of the thesis.

I am grateful to my Professors in LUISS, in Tilburg University, and in Paris Dauphine, without whom I would never have reached this point. Also, I am thankful to the QTEM staff of LUISS for the bright opportunity and the administrative help.

I would like to extend my sincere thanks to Professor Carlo Gotta for the priceless advice during these months.

I must recognize the key role of my housemates in Verbs, especially Humayd and Alessandro, whose help was crucial during the draft of this work. Also, my friends and colleagues in Paris, Tilburg, Roma, and Rionero made this year special and gave me enthusiasm even in the most difficult moments.

I am deeply indebted to Beatrice for her constant presence and for being the first to influence my interest in quantitative research.

Finally, a special thanks goes to my family. For the unconditional, unequivocal, and persistent support. I cannot thank you enough.

Summary: Routing Games, Information and Learning

LUISS Guido Carli

Department of Economics and Finance

by Francesco Giordano

June 10, 2022

1 Introduction and contribution

In this thesis we present some modern aspects of *Selfish Routing* models. *Routing games* are the tools to describe and model realistic instances of routing, such as telecommunication and road traffic.

We first present the classic setup, which consists of a static game without uncertainty in which a traffic inflow has to be routed from a set of sources to a set of destinations: we mention the main results concerning equilibria, optimality, and efficiency.

Afterwards, we consider some recent contributions where some forms of uncertainty are introduced – as in the level of traffic demands, or in the costs experienced by the players. Several sources of randomness emerge in real traffic situations and have been theoretically studied before. In the analysis of how uncertainty can affect routing behavior we focus on a form of incomplete information that can eventually be unveiled through the assumption of Bayesian rationality.

We consider routing instances with an unknown state that influences edge costs, in which the traffic is routed by minimizing the expected cost of each path. Learning the true underlying state may be possible when this kind of incomplete information games are repeatedly played, with a random level of demand in each stage game. Specifically, there are two possible kinds of social learning in non-atomic routing games: *strong learning*, when the true network state is identified, and *weak learning*, when players learn how to play as if the true state were known.

The original theoretical contribution is presented in the second part of the the-

sis. A recent result of Macault et al. (2022) clarified when *social learning* occurs in single commodity non-atomic routing games with infinite edge capacities: the authors showed that strong learning occurs almost surely if the network respects a topological condition called *series-parallel*.

From this starting point, in the thesis we precisely define when a network state is *identifiable* and extend the network conditions required to achieve almost sure learning also in multi-commodity and capacitated instances. This thesis extends the existing results in two different directions:

1. it considers multi-commodity models;
2. it allows edges to have traffic capacity.

Considering strictly increasing, continuous and unbounded costs, what emerges is that the learning conditions differ between capacitated instances and instances with infinite edge capacities.

Specifically, in capacitated instances a sufficient condition to achieve learning is that the underlying network respects a condition such that, under full congestion, the load on each edge reaches the edge load upper bound. This condition is viewed as a weaker form of *capacity conservation*; under capacity conservation, the total capacity entering in each node equals to the total capacity exiting from that node. It is noteworthy to notice that this condition does not require any restriction on the network topology.

On the other hand, in infinite capacity instances a sufficient condition to achieve learning is that the sub-network available to each commodity is *series-parallel*.

2 Deterministic Routing Games

A traffic instance can be modelled as a non-cooperative game played on a network. In a selfish routing game, some masses of traffic inflow - each called *commodity* - have to be routed from source vertices to destinations vertices of the underlying directed graph; the traffic assignment on the paths of the network is called *flow*.

In the non-atomic setting the set of players is defined in terms of infinitely divisible masses of traffic, and each player has a negligible influence on edge costs.

A *non-atomic instance* is defined by

$$G = (\mathcal{N}, \mathbf{d}, \mathbf{c}, \{\gamma_e\}_{e \in E}),$$

in which \mathcal{N} is the underlying network, $\mathbf{d} = \{d_i\}_{i \in I}$ is the vector of traffic demands, $\{\gamma_e\}_{e \in E}$ is the set of edge capacities, and $\mathbf{c} = \{c_e\}_{e \in E}$ is the set of cost functions. If

edge capacities are infinite, an instance is denoted only by the underlying network, costs and demands, $G = (\mathcal{N}, \mathbf{d}, \mathbf{c})$.

The symbol \mathcal{P} denotes the the set of paths available to the players. The amount of traffic routed on each path defines a *flow*. More precisely, a flow $\mathbf{y} = \{y_P\}_{P \in \mathcal{P}}$ is a vector that assigns to each path $P \in \mathcal{P}$ a non-negative number representing the amount of traffic routed through that path P . A flow vector is said feasible if all the traffic demand is routed through the network and edge capacities are respected.

We define the *load* on an edge as

$$x_e = \sum_{P:e \in P} y_P.$$

A *flow* is in equilibrium when it uses only minimum cost paths, preventing any profitable traffic reallocation. In non-atomic games, an equilibrium flow is also called *Nash Flow* or *Wardrop Equilibrium*. When edges have infinite capacity the costs of all used paths are equal, and they are lower than the costs which would be experienced by any arbitrary mass of traffic on any unused path.

A *Nash flow* of an atomic routing game always exists. This follows from the fact that routing games are potential games with continuous players' sets. This can be shown thanks to the equivalence of a *Nash Flows* and the local minima of a *potential function* which completely characterizes the costs of the game. Moreover, in the non-atomic case Nash flows are essentially unique, in the sense that multiple Nash flows are allowed, but all equilibrium flows induce the same costs.

An *optimal flow*, on the other hand, is a flow that minimizes the total cost experienced by the players. Also optimal flows always exist, and this stems from that fact that they are the solution of a convex optimization problem on a closed and bounded set.

In capacitated instances the Wardrop principle at equilibrium does not hold anymore - that is, equilibrium costs of paths are not guaranteed to equal and equilibrium loads are not always unique. The equilibrium concepts for capacitated instances analogous to the Wardrop Equilibrium is called *capacitated user equilibrium*.

3 Routing Games under Uncertainty

Multiple sources of uncertainty can affect routing behavior. Consequently, there are several extensions on how to introduce and model uncertainty in deterministic routing games. In the work we focus on a model of incomplete information with random demands and uncertain costs. The uncertainty in the costs only depends

on the underlying network state, which is initially unknown. We consider a routing instance in which the true network conditions are not known by the players and are modelled through an *unknown network state*. Players have a belief distribution on the network state and play an equilibrium accordingly to their belief.

We consider a finite set of possible network states denoted by \mathcal{T} . Moreover, Θ denotes the unknown random state, and $\theta \in \mathcal{T}$ its realization. Also, players are assumed to have a the common belief, denoted by μ , about the unknown state.

A feasible flow vector is in equilibrium if it minimizes the expected costs of the paths given the belief μ . In other words, a flow is at equilibrium in an instance with unknown state if no commodity has an path with strictly smaller expected cost than any other path used.

A *Dynamic Non-atomic Routing Game* (DNRG) is a routing instance repeated over time with a random traffic inflow in each stage game. Time is discrete and γ is the total capacity of the underlying network, which is possibly infinite.

Traffic demands are modelled as a sequence of independent identically distributed (i.i.d.) non-negative random vectors with common marginal distributions bounded above by the network capacities. The model goes as follows: in each period a traffic demand is realized and publicly observed. An *equilibrium flow* is played and the information about equilibrium loads and costs is immediately broadcasted to all the players.

Repeatedly playing a game, players may be collectively able to exploit the randomness of traffic demands to unveil the randomness in the network state. In each stage game the public belief is updated according to the Bayes' rule, and as soon as a state is identified by some equilibrium costs, the public belief changes.

The posterior public belief conditioned on the whole history up to period t is denoted by μ^t . The posterior μ^t is a random variable because it depends on the random demands up to period t , and being a probability distribution it is always bounded. It follows that the sequence of posterior beliefs is a bounded martingale. Hence, by the martingale convergence theorem, it converges.

Building on this setup, two ideas of social learning are formalized: *Strong Learning* and *Weak Learning*. The idea of Strong Learning is that the random sequence of posteriors converges almost surely to a distribution which assigns positive probability only to the true state. Weak Learning, on the other side, implies that true state is not necessarily discovered, but in equilibrium and asymptotically the total traffic is routed as if the true state were known.

4 Social Learning in non-atomic routing games

In this section we examine the problem of *Social Learning* in a Dynamic Non-atomic Routing Game both in the case in which the underlying network is endowed with infinite capacities, and in the case in which edge capacities are finite.

Before presenting the main results, we study some structural properties of capacity constrained networks. First, we notice that edge loads in capacitated networks are upper bounded: we characterize this bound, which does not always coincide with the capacity of the corresponding edge. We then reframe the definition of *feasible loads* in terms of these upper bounds. Also, from this starting point we are able to define when an unknown network state is identifiable. This preliminary analysis is key to understand when social learning occurs on capacitated networks.

Finally we show that social learning occurs in capacitated instances if the underlying network respects a condition on its edge capacities called *weak capacity conservation*. Moreover, we show that social learning occurs in instances with infinite capacities if the underlying network respects the *series-parallel* structure in each sub-network available to its source-destination pairs.

Let us first clarify some notation. Since the network \mathcal{N} is directed, each edge has a *tail node* and an *head node*. We denote the set of edges which exit from edge e by H_e and the set of edges that enter in edge e by T_e . More precisely, let

$$H_e = \{\text{set of exiting edges from the head node of edge } e\},$$

$$T_e = \{\text{set of entering edges into the tail node of edge } e\}.$$

Also, we recall that a flow vector $\mathbf{y} = \{y_P\}_{P \in \mathcal{P}}$ is said *feasible* if all the traffic demand is routed and if the implied edge loads are such that

$$x_e \leq \gamma_e \quad \forall e \in E,$$

in which γ_e is the edge capacity.

However, if capacities are finite, the maximum load reachable on an edge may be smaller than the edge capacity. Nevertheless, this maximum load is completely determined by the set of edge capacities. We characterize the *load upper bound* of a feasible flow vector in a capacitated instance and denote it by \hat{x}_e .

The upper bound \hat{x}_e is important in defining when a unknown network state is *identifiable*, and thus can be learned. We clarify that a load is feasible if and only if it respects the upper bound contemporaneously on all its edges.

At this point, we can characterize the *feasibility* of a flow by considering two conditions: one condition that assures that edge capacities are respected, and a second condition which assures that all the traffic is routed. We can define the set of *feasible flows* as the set in which both conditions are satisfied.

Having clear when a flow vector is feasible, we can characterize when an unknown network state can be identified. A random state is *identifiable* if for all pairs of possible state realizations there exists an edge such that, for at least one value $\bar{x}_e \leq \hat{x}_e$, we have

$$c_e(\bar{x}_e, \theta) \neq c_e(\bar{x}_e, \hat{\theta}).$$

In capacity constrained networks, equilibrium loads are not guaranteed to reach the edge capacity for all the edges. Eventually, equilibrium loads cannot even reach the load upper bound on all edges; this may prevent *strong learning*. In particular, if the load value that allows to identify states cannot be reached, then strong learning cannot occur. In the thesis we provide some examples in which learning does not occur.

4.1 Capacity conservation

In what follows we denote by γ the total capacity of a network. It corresponds to the sum of the edge capacities of the smallest *cut* of the network.

Specifically, a cut $\mathcal{C} \subseteq E$ in a single source-destination directed multigraph is a set of edges such that it is not possible to route from the source to the destination without passing through \mathcal{C} . The capacity of a cut $\gamma_{\mathcal{C}}$ is defined additively as the sum of the capacities of the edges that compose it. The *smallest cut* is the cut with smallest capacity. The capacity of the smallest cut is the network capacity, and denotes the maximum amount of traffic that can be routed through the network.

In the same spirit, a multi commodity network \mathcal{N} is composed by the union of the sub-networks \mathcal{N}_i available to each commodity. In a multi-commodity network we can refer to a cut as a set of edges such that it is not possible to go from any source to any destination without crossing \mathcal{C} .

Again, the capacity of the network corresponds to the capacity of the smallest cut.

We introduce a network condition such that social learning occurs when a DNRG is played on a capacitated network. From now on, for every internal node $v \in V$ we define the sets of entering edges into v and exiting edges from v respectively as T_v and H_v ,

$$T_v = \{ \text{set of entering edges into node } v \},$$

$$H_v = \{ \text{set of exiting edges from node } v \}.$$

We say that a capacitated network \mathcal{N} respects the *capacity conservation* property if on all internal nodes $v \in V$ it is satisfied

$$\sum_{e \in H_v} \gamma_e = \sum_{e \in T_v} \gamma_e.$$

In words, this means that for each node $v \in V$ that is not the source or the destination, the sum of the capacities of the edges that enter in v is equal to the sum of the capacities of the edges that exit from v .

A fully congested network that respects *capacity conservation* is such that

$$d = \gamma \implies \forall e \in E, \quad x_e = \hat{x}_e.$$

That is, if the traffic demand is equal to the network capacity, all edge loads reach their load upper bound.

This a key desirable property, from which we introduce a weaker form of capacity conservation, called *weak capacity conservation*.

Definition 4.1 (Weak capacity conservation) *Let $(\mathcal{N}, d, \mathbf{c}, \{\gamma_e\}_{e \in E})$ be a routing instance and \hat{x}_e be the load upper bound of edge $e \in E$. The network \mathcal{N} respects weak capacity conservation if*

$$d = \gamma \implies \forall e \in E, \quad x_e = \hat{x}_e.$$

In our dynamic model with a random demand D , a realization d exactly equal to the network capacity γ happens with probability zero. Nevertheless, since edge costs are strictly increasing, equilibrium loads are continuous in d .

Hence, we can say that

$$\text{supp}(D) = [0, \gamma) \implies \forall e \in E, \quad \text{supp}(x_e^*) = [0, \hat{x}_e),$$

and this will reveal sufficient to achieve *strong learning*. Here, x_e^* denotes an equilibrium load on edge e , which is a deterministic non-negative function of the demand realization d .

A similar argument holds for multi-commodity instances. However, edges that are available to different commodities deserve a careful consideration.

In particular, we consider the case in which commodities whose sub-networks share a common edge have the right of a proportion of the capacity on the edge shared. The predetermined allocation of the common edges allows to determine the capacity

γ_i of the sub-network \mathcal{N}_i available to each commodity $i \in I$.

Under this construction, *capacity conservation* implies that

$$\gamma = \sum_{i \in I} \gamma_i,$$

in which γ is the total network capacity and γ_i is the capacity of the sub-network available to commodity i .

Thus, in a multi commodity network we can generalize the property of weak conservation of capacities as

$$d_i = \gamma_i \quad \forall i \in I \quad \implies \quad x_e = \hat{x}_e \quad \forall e \in E.$$

Again by continuity of each edge load in the traffic demands, we have

$$\text{supp}(\mathbf{D}) = \times_{i \in I} [0, \gamma_i) \quad \implies \quad \text{supp}(x_e^*) = [0, \hat{x}_e) \quad \forall e \in E.$$

4.2 SP networks and generalization

Then we focus on networks with possibly infinite edge capacities. We start with the definition of the class of single-commodity *series-parallel* networks. This particular network topology respects desirable properties that exclude some occurrences of learning failure.

The key aspect is that if a network is not series-parallel, then it contains sub-networks such that some edges are not used in equilibrium even when the demand is high.

Recall that \mathcal{P}_i stands for the set of paths available to commodity i in a instance with multiple commodities. We can characterize a multi-commodity instance that respects similar desirable properties by considering the case in which each set \mathcal{P}_i form itself a series-parallel network. We call the resulting network U-SP. This means that each commodity routes its traffic on a series-parallel network.

4.3 Learning in capacity constrained networks

Recall that the players start in the first period with a prior distribution on the network state, which is common knowledge. In each stage a random demand is realized and the equilibrium flow is played; all the new information about the experienced costs is immediately broadcasted to all the players, which update their common belief according to the Bayes' rule. Strong learning occurs when the posterior public

belief converges almost surely to the degenerate distribution which assigns positive probability only to the true state.

Now we report state the main results of social learning in capacity constrained networks. Some fundamental cost assumptions are the following.

Assumptions 1-3:

1. edge costs are strictly increasing in the load for all possible states;
2. edge costs are continuous in the load for all possible states;
3. the following limit condition holds,

$$\lim_{x_e \rightarrow \gamma_e} c_e(x_e, \theta) = +\infty \quad \forall \theta \in \mathcal{T}.$$

Theorem 4.1 (Learning in capacity constrained networks) *Let Γ be a single commodity DNRG with identifiable unknown network state such that the network \mathcal{N} respects weak capacity conservation and [assumptions 1-3](#) hold.*

Then, if

$$\text{supp}(D) = [0, \gamma),$$

Strong Learning occurs.

The key idea of the proof is that, under such conditions, for any possible edge load x_e between zero and the load upper bound there exists a realization of the demand that generates it. By continuity of the costs in the loads and of the loads in the demands, it is enough to reach a value that belongs to a small neighborhood of the required demand that guarantee a change in the belief. Thanks to the condition on the support of D , any such neighborhood has positive probability.

A general result holds for the multi-commodity case. A major remark is that it is not sufficient that each demand D_i has as support the set $[0, \gamma_i)$. It is needed that the joint distribution of the demand vector \mathbf{D} has a support as described in the statement of the theorem.

Theorem 4.2 *Let Γ be a multi-commodity DNRG with identifiable unknown network state such that the network \mathcal{N} respects weak capacity conservation and [assumptions 1-3](#) hold.*

Then, if

$$\text{supp}(\mathbf{D}) = \times_i [0, \gamma_i),$$

Strong Learning occurs.

4.4 Learning in networks with infinite capacities

The precursory work of Macault et al. (2022) covers the single origin-destination case with infinite capacities. The authors proved the following result.

Theorem 4.3 *Let Γ be a single commodity DNRG with identifiable unknown network state such that the network \mathcal{N} is SP, edges have infinite capacities and [assumptions 1-3](#) hold.*

Then, if

$$\text{supp}(D) = [0, +\infty),$$

Strong Learning occurs.

A similar result holds for instances with more than one commodity.

Theorem 4.4 *Let Γ be a multi-commodity DNRG with identifiable unknown network state such that the network \mathcal{N} is U-SP, edges have infinite capacities and [assumptions 1-3](#) hold.*

Then, if

$$\text{supp}(D) = \mathbb{R}_+^k,$$

Strong Learning occurs.

The learning procedure is the same as in the single origin case, with equilibrium edge loads are that are continuous in the demands and unbounded.

5 Conclusion

In the thesis we study some of the most recent proceedings of the traffic assignment problem. In a selfish routing game, an amount of traffic inflow has to be routed from some sources to some destinations of the underlying directed graph; the traffic assignment on the paths of the network is called *flow*.

In the first part of the thesis we study routing games under complete information and introduce the ideas of equilibrium flows, optimal flows, and efficiency of routing instances. Then, we study some models that arise under uncertainty. We consider routing instances with an unknown network state which influences edge costs, in which the traffic is routed by minimizing the expected cost of each path. Learning the true underlying state may be possible when this kind of incomplete information games are repeatedly played with a random level of demand in each stage game. Specifically, there are two possible kinds of social learning in non-atomic routing games: *strong learning*, when the true network state is identified, and *weak learning*, when players learn how to play as if the true state were known.

In the second part of the thesis, we show under which conditions social learning occurs. We define when a network state is *identifiable*, and show the network conditions required to achieve almost sure learning. Considering strictly increasing, continuous and unbounded costs, what emerges is that the learning conditions differ between capacitated instances and instances with infinite edge capacities. In capacitated instances a sufficient condition to achieve learning is that, under full congestion, the load on each edge reaches the edge load upper bound. This condition is viewed as a weaker form of *capacity conservation*; under capacity conservation, the total capacity entering in each node equals to the total capacity exiting from that node. On the other hand, in infinite capacity instances a sufficient condition to achieve learning is that the sub-network available to each commodity is *series-parallel*.