Strategic Decisions in Directed Networks

Strategic Decisions in Directed Networks:

Experiments on Queueing, Route Choice, and Departure Time

By

Amnon Rapoport,
Eyran J. Gisches
and Vincent Mak

Cambridge Scholars Publishing

Strategic Decisions in Directed Networks:
Experiments on Queueing, Route Choice, and Departure Time

By Amnon Rapoport, Eyran J. Gisches and Vincent Mak

This book first published 2022

Cambridge Scholars Publishing

Lady Stephenson Library, Newcastle upon Tyne, NE6 2PA, UK

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

ISBN (10): 1-5275-7762-7
ISBN (13): 978-1-5275-7762-6

CONTENTS

ACKNOWLEDGEMENTS

The experiments that we had chosen to include in the present book have been selected from a larger set of research studies which have been conducted in the last 17 years or so with the collaboration of (in alphabetical order) Meng Cheng (Chs.11, 12), Terry E. Daniel (Chs. 14, 15), Subhasish Dugar (Ch. 7), Tamar Kugler (Ch. 7), Robin Lindsey (Ch. 14), Caiyun Liu (Ch. 10), Myounghee Moon (Chs. 11, 12), Thomas E. Nicholas (Ch. 7), James E. Parco (Chs. 2, 3), Darryl A. Seale (Chs. 2, 3, 4, 5, 11, 12), William E. Stein (Chs. 2, 3, 4, 5), Rui Yang Chs. 11, 12), Hongtao Zhang (Ch. 4), and Rami Zwick (Chs. 4, 5). A list of references is presented below. We acknowledge their valuable contributions to the formulation of the research ideas, experimental design, data collection, data analysis, or writing of the papers, and thank them for their constructive ideas and important contributions.

We also acknowledge the financial support of multiple NSF grants (listed below) that we have secured during this period, and the academic and financial support of the University of Arizona, University of Cambridge, University of California in Riverside, and Hong Kong University of Science and Technology.

A LIST OF THE ORIGINAL REFERENCES

Rapoport, A., Gisches, E.J., and Mak, V. Introductopn. (Chapter 1, This book)

Rapoport, A., Stein, W. E. Parco, J. E., and Seale D. A. (2004). Equilibrium play in single-server queues with endogenously determined arrival times. *Jornal of Economic Behavior & Organization*, 55, 67-91. (Chapter 2)

Seale, D. A., Parco, J. E., Stein, W. E. and Rapoport, A, (2005). Joining a queue or staying out: Effects of information structure and service time on arrival and staying out decisions. *Experimental Economics*, 8, 117-144 (Chapter 3)

Stein, W. E., Rapoport, A., Seale, D. A., Zhang Hongtao, and Zwick R. (2007). Batch queues with choice of arrivals: Equilibrium analysis and experimental study. *Games and Economic Behavior*, 59, 345-364. (Chapter 4)

Rapoport, A, Stein, W. E., Mak, V., Zwick, R., and Seale D. A. (2010). Endognous arrival in batch queues with constant or variable capacity. *Transportation Research, Part B*, 44, 1166-1185. (Chapter 5).

Rapoport, A., Mak, V., and Zwick, R, (2006). Navigating congested networks with variable demand: Exprimental evidence. *Journal of Economic Psychology*, 27, 648-666. (Chapter 6).

Rapoport, A., Kugler, T., Dugar, S., and Gisches, E. J. (2009). Choice of routes in congested traffic networks: Experimental tests of the Braess Paradox. *Games and Economic Behavior*, 65, 538-571. (Chapter 7)

Gisches, E. J. and Rapoport, A. (2010). Degrading network capacity may improve performance: Private versus public monitoring in the Braess Paradox. *Theory and Decision*, 73, 267-293. (Chapter 8)

Rapoport, A., Gishes, E. J., and Mak, V. (2014). Distributed decisions in networks: Laboratory study of Splittable flow. *Production and Operations Management*, 23(2), 314-331. (Chapter 9).

Liu, C., Mak, V., and Rapoport, A. (2015). Cost-sharing in directed networks: Experimental study of equilibrium choice and system dynamics. *Journal of Operations Management*, 39-40, 31-49. (Chapter 10)

Mak, V., Seale, D. A., Gisches, E. J., Rapoport, A., Cheng, M., Moon, M., and Yang, R. (2018). A network ridesharing experiment with sequential choice of Transportation mode. *Theory and Decision*, 85, 407-433. (Chapter 11)

Mak, V., Seale, D. A., Gisches, E. J., Cheng, M., Moon, M., Yang, R., and Rapoport, A. (2018). The Braess paraox and coordination failure in directed networks with mixed externalities. *Production and Operations Management*, 27(4), 717-733. (Chapter 12)

Mak, V., Gisches, E. J., and Rapoport, A. (2015). Route vs. segment: An experiment on real-time travel information in congestible networks. *Production and Operations Management*, 24(6), 947-960. (Chapter 13)

Rapoport, A., Gisches, E. J., Daniel, T., and Lindsey, R. (2014). Pre-trip information and route-choice decisions with stochastic travel conditions: Experiment. *Transportation Research Part B*, 68, 154-172. (Chapter 14).

Daniel, T., Gisches, E. J., and Rapoport, A. (2009). Departute times in Y-shaped networks and multiple bottlenecks. *American Economic Review*, 99(5) 2149-2176. (Chapter 15)

CHAPTER 1

INTRODUCTION

Networks form the infrastructure for the functioning of modern society particularly in the domains of transportation, communication, and supply chain management. Cars travel on highways, trains move on railways, phone calls and electronic messages are transmitted through telecommunication channels, and supply chains are often structured to decrease costs and increase efficiency. The design, maintenance, and structural changes of the topology of these networks is relegated to engineers, scientists, technicians, marketing experts, and administrators, where the successful functioning of such networks depends in part on the decisions of their users. What are the patterns of behavior that guide network users if, and at what time, to join a queue? What rules, norms, or customs commuters follow when they choose one route rather than another to connect their origin of travel to their destination? More generally, what are the behavioral patterns that emerge when group of agents, who act independently of one another, traverse such networks repeatedly (see, e.g., Schelling, 1978)?

This book contains a body of research that investigates the above questions over the past two decades. It consists of fourteen chapters on experimental studies of interactive decision making in directed networks. The chapters are grouped into four parts each representing a programme of research on a specific area of studies. While each chapter is self-contained and can be read independently, the chapters are also highly interrelated in terms of overarching substantive topic as well as theoretical and methodological approach. They can, therefore, be appreciated together at the level of the part or the book as a whole. Specifically, the chapters report the design, experimental procedure, and results of multiple computer-controlled laboratory experiments on the formation of queues, choice of routes in directed networks susceptible to congestion, ridesharing, and departure time decisions in traffic networks. Underlying the theoretical framework that is shared by all these studies is game theory—a loose amalgamation of mathematical models of interactive decision making in which the individual outcomes of group members, who are trying to maximize their expected utility, depend on the choices made by all the other group members. The approaches taken by game theory to model interactive decision behavior are broadly classified as either “cooperative” or “non-cooperative” depending on whether agreements by the group members are binding and enforceable. Remarking on this distinction, Aumann (1997) has noted that these two approaches should not be strictly considered as analyses of different kinds of interaction; rather, and more broadly, they should be viewed as different ways of modeling the same interactive decision-making process. As our main interest is in finding out how the participants in our experiments are *acting*, particularly in stage games played repeatedly in time, rather than merely focusing on what *outcomes* we expect to observe at the termination of the interaction, the papers collected in this book all adhere to the methodology of the non-cooperative game theoretical approach.

In addition to sharing the same theoretical framework for modeling interactive decision-making behavior, the studies that we review below also share critical features of the experimental design. In all these studies, university undergraduate students in the US and Hong Kong volunteered to take part in laboratory experiments on interactive decision making for monetary payoff contingent on their performance; no other incentives have been offered, and no questionnaires have been distributed either at the beginning or end of the experimental sessions. Specifically, no attempts have been made to measure risk-aversion, inequity-aversion, or cognitive biases, or to collect personal information. Because of the difficulties in analyzing and interpreting subjective data, which are in principle unknowledgeable by anybody else and arguably inaccessible to direct scientific measurement, and to enhance the generalizability of our findings, we have taken a reductionist approach which focuses exclusively on the collection of choice data. Whereas this approach might prove to be insufficient and too limited in modeling two-person interactions, we argue below that some of these limitations may not apply in large-scale studies of interactive behavior.

In all the experiments, the participants were seated in separate cubicles in a spacious laboratory with no possibility of face-to-face interaction. Anonymity of the participants was strictly enforced as all forms of communication of individual choices and outcomes were computer controlled. The stage games were always repeated in time; some of them comprised of two or three separate parts of the experiment (preceded by specific instructions) that were played sequentially. Written instructions were distributed to the participants at the beginning of each experiment, which they could access at any time during the session. The size of the groups (denoted by the variable *n* throughout this book unless otherwise stated), which we discuss later in some detail, varied from n=10 to n=40.

In assembling the studies into a single volume, we have had three major objectives in mind. The original articles have been published in different journals that span a wide range of scientific disciplines which are only loosely connected with one another. Among others, they include *Production and Operations Management*, *Games and Economic Behavior*, *Journal of Economic Psychology*, *Theory and Decision*, *American Economic Review*, and *Transportation Research Part B*. While interdisciplinary research is often appreciated, in practice it is now quite uncommon, and reading across disciplines is not encouraged (probably because it takes too much time and effort to cross the boundary lines between different disciplines, and the potential benefits of doing so are not immediately recognized or fully appreciated). Our first objective has been to publish these articles together and by doing so render them accessible to readers in different scientific disciplines, maximize their impact, and investigate their cumulative effects on the experimental design of large-group research. To quote Roth in the introduction of his book *Laboratory Experimentation in Economics*, "Thus, the hope for this book is that the whole will be more than the sum of its parts" (1987, p. 3).

The second objective is to stimulate new research on decision-making behavior in directed networks, which are quite common in—although not limited to—communication (e.g., internet services), transportation (e.g., traffic networks), and operations management (e.g., supply chain networks). Formally, a directed network is a graph with an *ordered* pair G = (V, E), where V is a set whose elements are called vertices, nodes, or points, depending on the discipline, and E is a set of ordered pairs of vertices called edges or links. Our research focuses on *simple directed networks* where loops are not allowed. Directed networks differ from *undirected* networks, which are defined in terms of unordered pairs of vertices. Whereas undirected networks are often used to model various forms of interaction in social groups, and have been studied extensively, both theoretically and experimentally, in social psychology, behavioral economics, operational management, and sociology (see, e.g., Easily & Kleinberg, 2010; Jackson, 2008), experimental research on decision making in directed networks has been sparse. For a recent review of experimental studies of decision making in queues, which had first been published in journals on psychology, marketing, and operations management, see Allon and Kremer (2019).

The third major objective has been to publish a collection of our joint papers on strategic behavior in *large groups* that crosses over different disciplines, as we believe that their findings have implications for the contribution of non-cooperative game theory to the *descriptive*—and not only normative—models of interactive strategic decisions. To elaborate, consider the literature on group interaction in the disciplines of social psychology, sociology, experimental and behavioral economics, transportation research, and operations management. Traditionally, they distinguish between "small" and "large" groups without committing to the value of the magical integer that separates these two classes of groups from each other. Most experimenters in these disciplines will surely agree that groups of size n=2 (e.g., the iterated Prisoner's Dilemma game, the Ultimatum game, the Trust game, two-person bargaining games) should be considered as "small" whereas groups of, say, size n>100 should be considered as "large". Obviously, "small" and "large" are fuzzy concepts that may not be separated sharply. What matters in their explication is not so much the number of the group members, but rather the differential *impact* of the individual decisions of the other n-1 members on the decisions of the pivotal member. The force of this impact, which is highly correlated with group size, depends in part on the characteristics of the game, experimental design, sophistication of the group members, and the experience they gain in playing the repeated game.

When n=2 and the design call for fixed pairs playing the same game repeatedly, anonymity of the players is not maintained and psychological factors such as altruism, punishment, reciprocity, and fairness shape to a large extent the decisions of each group member. The effects of these factors have been studied extensively in the voluminous literature on public good provision, ultimatum bargaining, the trust game, and other types of social dilemmas. We argue that despite their apparent complexity, non-cooperative games played by large groups are considerably *simpler* than small-group (particularly n = 2) games, as anonymity of the players is assured in large groups, reciprocity may no longer be applied differentially, and punishment directed at individual players for deviating from social norms or just acting in contrast to the group interest may no longer be exercised effectively. If across the iterations of the game some members of large groups wish—and a few do—to shape the future decisions of the other group members, then they respond to selected past *group* outcome statistics, not to individual choices or outcomes. Among others, these statistics may include the proportion of group members choosing the same strategy on the preceding round, mean group payoff on the preceding round, range of individual payoffs, and the number of "volunteers" who sacrifice part of their earnings on previous rounds as a signal to re-direct future group decisions to strategies that increase social welfare. As reciprocity, punishment, and altruism may no longer serve their purpose, and their effects rapidly vanish as *n* increases, players in large groups tend to perceive the group as their "environment" and act accordingly. Some will say—we realize that we are on shaky ground here—that they behave more "rationally."

We are not the first to make this point although we may be stating its implications more provocatively. Almost 50 years ago Mancur Olson pointed out in his seminal book *The Logic of Collective Action* the distinction between "small" and "large" groups and examined its implications. In discussing altruism in large groups, where agents disregard their personal welfare, he writes: "Such altruism is, however, considered exceptional, and self-interested behavior is usually thought to be the rule, at least when economic issues are not at stake" (1971, p. 2). Then he adds in the next paragraph: "None of the statements made above fully applies to the small groups, for the situation in small groups is much more complicated" (1971, p. 3). Experimental evidence supporting Olson's claim comes from the studies of the iterated two-person Prisoner's Dilemma (PD) game, which go back more than 60 years. When the PD game is played repeatedly by the same two players—the classical PD game—joint cooperative behavior tends to decline over the first 20-40 iterations, and then recovers and increases steadily (see, e.g., Rapoport & Chammah, 1965). If the game is played repeatedly by $n > 2$ players (typically $4 < n < 6$), then the rate of the dyadic cooperation declines quite early in the game, and in general play converges rapidly to the "rational" equilibrium solution of mutual defection (e.g., Goering & Kahan, 1976). In an early and well-cited review of the literature about experiments on multi-person social dilemma games, Dawes concluded: "All experimenters who have made explicit or implicit comparisons of dilemma games with varying number of players have concluded that subjects cooperate less in large groups than in small groups" (1980, p. 186). Small groups typically mean here that n=2 because "The overwhelmingly majority of experimental investigations of behavior in social dilemma games have studied subjects' responses in two-person prisoner's dilemmas that are played repeatedly by the same subjects" (1980, p. 182). The predominance of the two-person PD game research with its many variations has not changed much over the last 40 years. These sharp differences in the experimental results between small and large groups may be attributed to one or more differences between these two classes of games. In n-person ($n > 2$), but not two-person PD games, (i) the harm for defecting behavior is diffused over the n–1 players; (ii) the choices of individual strategies are made anonymously; and (iii) it is not possible for one player to attempt shaping a particular other player's behavior (e.g., such as the tit-for-tat strategy) by judicious choice of his/her own strategies.

Our findings suggest that the "rational" equilibrium solution may serve as a good prediction of aggregate behavior in large groups. We suggest that the differences between 2-person and n-person (n>2) interactive decisions are not restricted to public good and social dilemmas research, and that they increase as the value of n increases. Therefore, we question the assumption, which seems to underlie much of the existing research on small groups that its results may be rendered applicable to larger groups merely by multiplying these results by some scale factor. An implication of this suggestion is that the conclusion that the equilibrium solution has no role to play in descriptive models of interactive decision behavior (see, e.g., Selten, 1997) is premature and deserves critical examination. This conclusion is rarely qualified by noting that it is mostly based on the results of experiments on dyadic interaction. The experiments in the present volume, which we describe below, have been designed in part to address this issue.

Part 1: Decision Behavior in Queues

In a scientific sense, a queue consists of a system into which flows a stream of users who require some processing by a server with limited capacity before they leave the system. The arrival times of the users are determined either exogenously (commonly by the Poisson probability distribution) or endogenously. A common example is of batch queueing systems in transportation networks, like ferries, airport shuttles, and buses, (where potential passengers also have the option of staying out of the queue). If travelers decide to join the queue, then their waiting time depends on the arrival times of all the other passengers, the capacity of the server, queue discipline, and whether the queue is observable. Another example is a single-server queue which forms in front of a checkout counter in a supermarket, where there is a single server (cashier) that serves the customers in the order of their arrivals. Queues forming in front of retail stores on Black Friday, food vendors, voting booths on election days, sport events, public concerts, and cinemas are a common sight. The decisions whether to join a queue, and if so, at what time to arrive at the queue may be modeled as non-cooperative games and then studied experimentally in the controlled environment of the laboratory.

The first four chapters of this book describe experiments on decision making behavior in queues. Chapter 2 reports the results of a game on interactive decisions in queues with the following characteristics. Time of arrival at the queue is determined endogenously. Upon arriving at the queue, the customer decides to enter it and wait in line regardless of its length. Therefore, neither *balking* (not entering the queue upon arrival) nor *reneging* (leaving the queue after joining it) is allowed. There is a single service station, customers are served one at a time in a "first-in-first-out" (FIFO) queue discipline, and service time is fixed for all customers and commonly known. The population size is finite and commonly known. The strategy space is discrete and quite large; each customer has multiple pure strategies of time of arrival plus a single strategy of staying out of the queue. Opening and closing times of the station are fixed and commonly known (as is the case in banks, supermarkets, emission inspection stations, airlines, and most municipal offices). Finally, customers are prevented from arriving at the station and form a queue before it opens (thus, no lines in front of chain stores on the morning of Black Friday).

In the experiment, the group size is n = 20, starting time is 8:00, closing time is 18:00, and both arrival and service times are measured in minutes (hence the 601 pure strategies), The service time d, unit waiting cost c, and payoff for completing the service r are all commonly known. To fix ideas, the stage game has been structured as an emission inspection single-server station, where each driver chooses in advance the time of his/her arrival to minimize the cost of waiting time in pursuit of the service. Using a within-subject design, members of each of four groups of twenty homogenous participants were instructed to repeat the same game for 75 rounds of play. At the end of each round each participant was only informed her own payoff.

The study reports four major findings. First, the four groups do not differ significantly from one another; the patterns of arrival time and probability of staying out of the queue are consistent across the groups. Second, there is slow learning across the multiple (75) iterations of the stage game. Third, there is no support for equilibrium play on the individual level; participants in this experiment do not mix their pure strategies on each round of play. Fourth, and most importantly, the mixed-strategy equilibrium solution accounts remarkably well for the aggregate behavior; it approximates the three major statistics of aggregate behavior, namely, relative frequency distributions of arrival times, of inter-arrival arriving times, and of queue waiting times. As players do not compute mixed-strategy equilibria, these results present a challenge of constructing simple decision rules (also called "heuristics") that account for these consistent and replicable regularities on the group level.

Chapter 3 reports the results of a second experiment on decision making behavior in queues with endogenous arrival times. This experiment implements a 2 by 2 between-subject design (game type by information structure). The two versions of the game share most of the parameter values of the experiment in Chapter 2 except two. First, the design calls for two different service times, namely, 30 minutes in Game G30 and 45 minutes in Game G45. The difference between these two games has major implications for the equilibrium solutions and the participants' perception of the game. In Game G30, there are $n!$ equilibria in pure strategies where all n players enter the queue, no player waits in line, and the system idle time is zero. This equilibrium has the form of the 20 players arriving one at a time at exactly 8:00, 8:30, …, 17:30; technically, it is a simple game of coordination with a transparent solution, which is not likely to be achieved if n is large and communication between the players is forbidden. In Game G45, there are no symmetric pure-strategy equilibria, and congestion is unavoidable. Both games have unique (symmetric) mixed-strategy equilibria. In a second departure from the queueing system in Chapter 2, early arrivals in the experiment reported in Chapter 3 are allowed after 6:00 a.m., two hours before the service starts.

The two queueing systems in Chapters 2 and 3 share the same design and parameter values except for four major differences. First, the payoffs for staying out of the queue are different. Second, as mentioned earlier, the strategy spaces are different. Thirdly, the cost functions are different. Fourthly, whereas the experiment in Chapter 2 only provided private information at the end of each round, the experiment in Chapter 3 manipulated the information structure in a between-subject design (private information in four groups and public information in four other groups). The major findings are replicable patterns of arrival and staying out decisions in three of the four conditions, which are accounted quite well by the mixed-strategy equilibria. The only exception is the fourth condition in which Game G45 is played under private information. The authors do not interpret the findings of their study as evidence supporting equilibrium play. Rather, they consider them as a challenge to account simultaneously for the emergence of replicable patterns of decision behavior on the aggregate level that are accounted quite well by the mixed-strategy equilibria, at least as a first approximation, combined with individual patterns of decisions that at present defy characterization.

Unlike the two studies in Chapters 2 and 3 in which the customers are served one at a time and service time is fixed for all of them, Chapter 4 reports the results of a study of decision behavior in queues, where the customers are served in *batches*. Examples include ferries, airplanes, buses, and cable cars that operate on a strict and commonly known schedule and can carry a maximum number of passengers. Using a ferry as a motivating example, the study models the queueing system as a full information non-cooperative n-person game in strategic form, where the departure time of the ferry (T) is fixed and commonly known, service time (d) is zero, population size (n) is fixed, and passengers' arrival times are determined endogenously. Reneging (but not balking) is prohibited, and the service capacity (s) satisfies $s < n$. The latter inequality implies that not all potential passengers may embark on the ferry; some may decide to stay out and others may either join the line when the ferry is ready to depart or approach the ferry and be denied entry. Each person deciding to join the (unobservable) queue is charged a constant fee (f) and a variable cost (c) per minute waiting. If she boards the ferry, then the person gains a reward (r), whereas if she stays out of the queue, then she gains a smaller payoff (g) that satisfies $g < (r - f)$. The parameter values used in the experiment—there are quite a few of them—are $n = 20$, $f = 40$, $c = 4$, $r = 340$, $g = 60$, $s = 14$, and $T = $ 12:00 a.m.

The observability of the queue is one of the two major variables manipulated in the experiment: it may either be allowed or forbidden. The mixed-strategy equilibria for the unobservable and observable conditions are presented in the upper and lower panels of Table 1 in Chapter 4, respectively. When balking is prohibited (the queue is unobservable), each player should enter the queue at times 11:05, 11:45, 11:50, 11:55, and 12:00 (no waiting in line) with respective probabilities $p = 0.702$, 0.024, 0.006, 0.025, and 0.007, and stay out of the queue with $p = 0.236$. When balking is allowed (as the queue is observable), there are only two pure strategies under the mixed-strategy equilibrium solution, namely, enter the queue at time 11:05 with probability $p = 0.837$, and stay out, otherwise. With the cost structure presented in the experiment, observability of the queue dictates a higher probability of early arrival. Chapter 4 reports the results of an experiment using a 2 by 2 information structure (no information vs. full information) by game type (observable vs. unobservable queue) between-subject design with three groups in each condition. The stage game was iterated 60 times.

The authors of the study in Chapter 4 note that the mixed-strategy equilibrium solutions in both observability conditions are inefficient; the participants could more than *triple* their earnings by all arriving together at 12:00 (arrivals at the same time are broken randomly). Remarkably, they report no evidence of successful attempts to maximize individual and social welfare by choosing this very simple and transparent strategy (possibly potential passengers shy away from a strategy that calls for arrival in the last minute and results in an outcome that is determined randomly). Mitigating this result is the finding that with experience in repeating the stage game the aggregate decisions in all four conditions shifted away in the direction of social welfare maximization. Possibly additional experience with the game is required to shift away from equilibrium play. Subsequent chapters in the book on route choice in traffic networks susceptible to congestion provide additional evidence in support of these findings.

Ferries, buses, and airplanes may have different capacities. For example, airline companies are known to change the type of carrier in response to changes in weather or in demand, and ferries and buses may arrive at a given station only partially full. In a critical departure from the previous experiment in Chapter 4, Chapter 5 reports the results of a subsequent experiment on batch queueing systems in which the server's capacity, rather than fixed and commonly known, is a random variable with a commonly known probability distribution. Consequently, in determining whether to join the queue and, if so, at what time to do so, each player is faced simultaneously with two sources of uncertainty, namely, *strategic uncertainty* about the arrival times of the other members of the group and *environmental uncertainty* about the capacity of the server. In a second departure from the previous experiment, the experiment reported in Chapter 5 was conducted in continuous (rather than discrete) time to avoid random tie-breaking (which is often considered by customers to be unfair or biased) for passengers arriving simultaneously at the same time interval.

Chapter 5 reports the results of two between-subject experiments, namely, Experiment CON in which the server has a fixed and commonly known capacity, and Experiment VAR in which the server's capacity before its arrival is uncertain. To simplify the design of Experiment VAR, the unknown capacity of the ferry was modeled as a random variable that assumes one of two values (i.e., either "small" or "large" ferry) with equal probabilities. In Condition VAR(10,16), the capacity of the server was either small (10) or large (16) with equal probability 0.5. Condition VAR(8,18) had the same mean capacity but a larger variance. At the end of each round of play, each participant was informed of the capacity of the ferry, the number of participants who stayed out, arrival times of the participants opting to join the queue, and the payoff for the round. Early arrivals after 10:30 were allowed. With considerable experience in playing the stage game, the participants in Experiment CON slowly approached, but never reached, the equilibrium solution. Rather, they diverged from equilibrium play in both conditions in Experiment VAR by arriving at the queue significantly later than predicted (and thereby increasing their payoffs). Is this result due to an optimism bias, as reported in multiple experiments in psychology, or does it constitute evidence for departure from inefficient equilibrium play in the direction of maximizing social welfare? Additional experiments are warranted for answering this question.

Part 2: Choice of Routes in Directed Networks: Congestion

One cannot answer questions about decision-making behavior in a directed network without specifying the architecture of the network, the function that associates costs and rewards to these different choices, the information that each network user receives about the group choices and outcomes if she traverses the network repeatedly, and the number of the network users. And if one is interested in studying route choice decisions that are repeated in time, like in the morning commute, then one would be advised to supplement empirical studies, that are not well controlled, with experimental studies in the laboratory that allow manipulating specific variables, and better control for group size, group composition, incentives of the network users, and their possible biases. The four experiments reported in the chapters below constitute part of a research program designed to answer some of these questions.

Chapter 6 reports the results of an experimental study that changes systematically the demand for service in a two-route traffic network that is susceptible to the Braess Paradox (Braess, 1968). To describe the Braess Paradox (henceforth, BP), consider first the common intuition that adding one or more routes to a traffic network, and thereby increasing its capacity to serve more users, should decrease, or at best have no effect, on the cost of traveling from a common origin O to a common destination D. In a seminal paper, the mathematician Braess constructed a very simple example of a traffic network with only four vertices (two on each route) and four nodes and demonstrated that this conclusion may not necessarily hold. For this example to hold, Braess chose judiciously the costs of traversing each of the four vertices as well as the group size n. Paradoxically, he illustrated that if a third route is added to the basic network by connecting the two intermediate nodes, one in each route, and each of the n members of the group chooses independently her best route from O to D, then at the new symmetric equilibrium the travel cost for each of the n homogenous commuters *increases*.

Of course, it is typically not the case that expansion of a traffic network is detrimental to the users. Rather, the occurrence of the BP depends critically on the architecture of the network, the cost function, the group size, and its composition (homogenous vs. heterogeneous agents). The study reported in Chapter 6 was designed to answer a simple question: If the topology of the network, the cost function, and the composition of the group are held constant, how does the group size affect the BP? To answer this question experimentally, the authors constructed two networks. The *basic* network consists of two parallel routes denoted by O→A→D and O→B→D. Route O→A→D consists of the two links (segments) O→A and A→D (with a common "intermediate" node A), whereas route O→B→D consists of the links O→B and B→D (with a common "intermediate" node B). Links O→A and B→D are susceptible to congestion and the other two routes A→D and O→B are not. The *augmented* network includes a third route, denoted by O→A→B→D which is formed by connecting nodes A and B.

Equilibrium solutions in pure strategies for the basic and augmented games were constructed for n = 2, 4, 6, … 40, …, and instigated four testable hypotheses. The first three were tested and confirmed the predictions; the fourth could not be tested for administrative reasons.

● If 2 < n ≤ 14, then each agent chooses route O→A→B→D in the augmented network. The equilibrium solutions are Pareto *efficient* (social welfare is maximized).

● If 16 < n ≤ 20, then each agent chooses route O→A→B→D in the augmented network. The equilibrium solutions are Pareto *deficient* (the BP is realized).

● If 22 < n ≤ 40, then only a fraction of the players chooses route O→A→B→D in the augmented network. The size of this fraction decreases in n. The remaining players split their choices equally between routes O→A→D and O→B→D. The equilibrium solutions are Pareto *deficient* (the BP is realized).

● If 42 ≤ n, then route O→A→B→D is abandoned. Rather, the n players split their choices between routes O→A→D and O→B→D in the basic network (social welfare in maximized).

It is not feasible to test the hypotheses separately for each value of n. Rather, the authors tested them in only three conditions, namely, n=10, n=20, and n=40 and reported the following results. In Condition n=10 all the 24 groups (a total of 240 players) converged rapidly to equilibrium. In Condition n =

20, all the 12 groups converged more slowly to equilibrium. In Condition n = 40, about 90 percent of the players stayed away from route O→A→B→D as predicted.

A major finding is that, in agreement with the BP, the participants in Condition n=20 sustained heavy losses for their choices as they could have doubled their payoff by splitting equally between routes O→A→D and O→B→D and abandoning route O→A→B→D. The experiment reported in Chapter 7 has been designed in part to study the robustness of this finding in two different networks, one with the same architecture and the second with a richer architecture.

Chapter 7 reports two experiments on the BP in which college students volunteered to take part in a decision-making experiment for payoff contingent on their performance. The 108 participants in Experiment 1 were divided into 6 groups of n = 18 each. Three groups played the basic game in Chapter 6 in the first half of the session, and the augmented game in the second part (Condition ADD). To control for order effects, three additional groups played the same two games in the reverse order (Condition DELETE). Each stage game was iterated for 40 rounds. The authors note that the BP in Condition ADD is realized if route O→A→B→D is *added* to the basic game (and consequently the individual cost of travel *increases* under equilibrium play). In contrast, the BP is realized in Condition DELETE if route O→A→B→D is *deleted* from the augmented game (and consequently the individual cost of travel *decreases* under equilibrium play). The (symmetric) equilibrium solutions do not change, but the psychological effects of sustaining gains (in Condition ADD) or losses (in Condition DELETE) in equilibrium with the shift from the first to the second part of the session may be significant. In fact, the authors report no order effects; in both conditions, the route choices converged to the equilibrium solutions, although it took all 40 rounds of play to achieve this convergence.

Most of the literature on the BP only considers the minimal critical network with two parallel routes in the basic game (each with two links), three routes in the augmented game, and symmetric cost functions. Can the BP be manifested in other weighted directed networks with a richer architecture? In a seminal book, Roughgarden (2005) has answered this question positively. Experiment 2 in Chapter 7 was designed to answer this question experimentally. The basic game in Experiment 2 consisted of *three* parallel routes, namely, O→A→D, O→B→D, and O→C→E→D. The augmented game was constructed by adding *two new links* to the basic game, namely, link C→A that connects nodes C and A, and link B→E that connects nodes B and E. Consequently, each participant playing the augmented game was requested to choose one of *five* routes: O→A→D, O→B→D, O→C→E→D, O→C→A→D, and O→B→E→D of which the first three are parallel and the last two are not. Because convergence to equilibrium in Experiment 1 was slow, each stage game in Experiment 2 was iterated 80 (rather than 40) times for a total of 160 rounds of play in each session. The authors report that all the six groups playing the basic game approached the equilibrium solution. In contrast, the equilibrium solution in the augmented games was only partly supported. In (the symmetric) equilibrium, nine players choose route O→C→A→D and nine other players choose route O→B→E→D. Only two third of the participants chose these two routes by the end of the season, whereas one third of the players continued splitting their choices equally among the three other routes. We note that to maximize social welfare, *all* the *n* players should split their route choices equally among the three original routes of the basic game. It seems that one third of the participants realized this situation and therefore persisted in choosing the original routes, or that 80 rounds of play were insufficient to reach equilibrium play by all the group players.

Chapter 8 continues reporting results of our research project on the BP, which began with the earlier studies in Chapters 6 and 7 of the present volume. Its first purpose is to study experimentally the generality of the BP when the architecture of the network is further enriched beyond the previous studies. In commenting on the BP and other game-theoretical paradoxes, Cohen (1988) remarked 30 years ago that these highly abstract networks and their paradoxical implications may arise from the many aspects in which they differ from reality rather than from these aspects that they share with it. And as Valiant and Roughgarden noted almost 17 years ago, "However, remarkably little is known whether the BP is a common real-world phenomenon, or a more theoretical curiosity" (2006, p. 2). Empirical data collected following major changes in the traffic infrastructure in several major cities in the world would seem to suggest that the BP is alive and well. Our research has been designed to study the validity of this argument in the controlled environment of the laboratory.

The second and more general purpose of Chapter 8 was to study the effects of the information outcome provided to the network users when they complete their travel from their origin to their destination at the end of each round of play. The experimental design calls for repetition of the stage game. Although repetition of the stage game is critical for the experimental study of the BP, it is known to have implications for the efficiency of the equilibrium. In fact, there is a gap between the simple two-route networks investigated in Chapter 6 and in Experiment 1 of Chapter 7 and networks with a considerably richer architecture. Whereas the equilibrium solution may be deduced in simple networks even if the outcome information is private, this is no longer the case in networks with richer topology. Consequently, if the network games are repeated in time, as in all our studies, private and public outcome information at the end of each round of play may result in different dynamics. To test this hypothesis, Chapter 8 reports an experiment with two between-subject treatments of the outcome information. In one condition, (called *private monitoring*), at the end of each round of play the user only observes the traffic flow on the route segments that she traversed. In the second condition (called *public monitoring*), she is informed of the route segments and the associated costs of travel incurred by *all* the group members. Note that regret for not choosing the best route is possible under public but not under private outcome information.

The traffic network in the basic game reported in Chapter 8 had *four* parallel routes, each with two segments, connecting the origin of the travel to its destination: O→A→D, O→E→D, O→F→D, and O→B→D. To construct the augmented network, two directed links, namely A→E and B→F, were added to the basic game resulting in two additional routes, namely, O→A→E→D and O→B→F→D. Using the same experimental procedure as in Chapters 6 and 7, the authors report no significant differences between the two information conditions. Under either private or public monitoring, the players converged to the equilibrium solutions in both the basic (four routes) and the augmented (six routes) games. Similar to Experiment 1 in Chapter 7, it took 80 rounds of play to achieve convergence. The authors report no evidence that with considerable experience in playing the route choice games the patterns of choice behavior increased efficiency.

Chapter 9 extends the research program on the BP in a new direction by considering traffic networks with *asymmetric network users*. In the "unsplittable" condition, the network users are all symmetric, the group size is commonly known, route choices are decentralized, each user controls a single unit of flow, and chooses a single route from O to D. In the "splittable" condition, which is the main focus of Chapter 9, the network users are asymmetric, route choices are partly regulated, each user controls *multiple units* of flow (hereafter called "fleet"), and independently chooses one or more routes from O to D on which to distribute her fleet. Under the splittable model, network users are no longer singletons. Rather, they are grouped into *cohorts* which serve as independent unitary decision makers for coordinating the route choices of their members. Thus, whereas the choices of the cohorts are decentralized, the choices of their individual members are not; rather, they are coordinated exgoneously. Truck companies in which each manager independently coordinates the routes assigned to each of the individual trucks in her company serve as an example; route assignments to police cars, ambulances, and fire trucks in the events of fires, floods, or major crime scenes is another.

Rapoport et al. (2014) report an experiment with two between-subject conditions and groups of size n=18 players. In the unsplittable condition (game G(U18), each player controls a single unit of flow, whereas in the splittable condition $n = 4$, and the four players {#8, #5, #3, #2} control fleets of size 8, 5, 3, and 2 units of flow, respectively. The basic and augmented games were the same as in Experiment 2 in Chapter 7: a basic network game with three parallel routes and an augmented network game with five routes (some of them sharing common links). In each experimental session, the basic and augmented games were presented separately in Parts I and II, respectively. Each stage game was iterated for 50 rounds of play. In game G(U18), each user assigned a single truck to one of the routes connecting O and D, whereas in game G(S4) each of the 4 players could split her fleet over the routes in any way she wished.

The authors report several results which we briefly summarize below. In both conditions, route choice behavior approached, although never converged, to equilibrium play, and the BP was clearly manifested. Travel cost decreased if, rather than choosing routes independently, subsets of users coordinated their route choices. Importantly, if the group players are no longer assumed to be symmetric, as in all previous studies of the BP, the splittable condition allows testing equilibrium predictions about the route choices on the *individual* (rather than the group) level. The contribution of this experiment to the research program on the BP is the illustration that the BP may be manifested in directed networks of different architecture and with both homogenous and non-homogenous network users.

Part 3: Route Choice in Directed Networks: Ridesharing

The four chapters reviewed in the previous section focus on route choice in congestible networks. The three chapters reviewed in the present section focus on ridesharing or, more generally, cost-sharing. The main difference between these two sets of studies is in the type of externalities imposed by a player on her group members. The externalities resulting from the decisions of each group members are *negative* if individual benefits are a *decreasing function* of the number of other group members making the same decision. Examples include choice of routes in congestible networks in the experiments reviewed in Section 2. The externalities are *positive* as when the benefit of choosing a mode of transportation in a traffic network (e.g., private transportation, shuttle) is an *increasing function* of the number of network users who make the same choice. Examples include passengers who agree to share the cost of travel, as is the case in carpools, in shuttles, and when people communicate with one another by the telephone or the internet.

Chapter 10 is one of the first experimental studies of cost-sharing in traffic network games with positive externalities. The purpose of this study is to answer the question if, and under what conditions, network users might achieve the efficient equilibrium through repeated play (for 50 rounds) of a *cost-sharing* game which has several Pareto rankable equilibria. For this purpose, they constructed a directed network with three parallel routes connecting a fixed origin and a fixed destination: route O→D (called the "private route") with a single link, route O→SH→D with two links, and route O→CP→D also with two links. Links O→D, O→SH, and O→CP were assigned fixed costs of travel (20, 0, and 10 units of cost, respectively), whereas the cost functions for links SH→D and CP→D imposed positive externalities (150/*m*(SH→D) and 70/m(CP→D) respectively, where *m*(X→Y) ($m < n$) is the number of players traversing the corresponding link from X to Y). The authors conducted two conditions in a between-subject design with $n = 10$ members in each group, who were instructed to choose independently one of the three routes from O to D. In Condition HOM (for "homogenous" players), the cost of traversing route O→D was fixed at 20 units for all *n* players and was commonly known. In Condition HET (for "heterogenous" players), the cost of traversing the "private" route was a random variable whose value was private knowledge to the player. The equilibrium solution in pure strategies for these two games vary considerably. In Condition HOM there are three Pareto rankable equilibria in pure strategies, where all the *n* players chose either route O→D, O→CP→D, or O→SH→D (called bad", "intermediate", and "good" equilibria, respectively). In Condition HET, there are three equilibria, namely, the same "bad" and "good" equilibria as in Condition HEM, where all the *n* players choose the same route, and an "intermediate" equilibrium, where each player chooses independently between routes O→D and O→SH→D depending on the value of her (randomly drawn) cost of private travel.

The authors report that the assignment of travel cost for the choice of the private route mattered. When the price of travel on O→D was fixed and commonly known, only five of the ten groups in Condition HOM converged to the efficient (i.e., “good”) equilibrium, whereas nine out of the ten groups in Condition HET, where the price of traveling on route O→D was variable and only know privately, converged to the most inefficient (i.e., “bad”) equilibrium. We conjecture that the absence of common knowledge of the cost of private travel in Condition HET, which was drawn randomly for each group member, rendered cost-sharing less palatable.

Upon analysis of the dynamics of play, the authors reported that some of the group members sacrificed part of their payoff in an attempt to signal their willingness to choose cost-sharing through their own decisions and thereby attract other players to do so in future rounds of play, a process known as “strategic teaching” (Fudenberg & Levine, 1998). These attempts were unsystematic, uncoordinated, and generally lasted for only 2-5 consecutive rounds. In general, they were unsuccessful because they did not last long, the other players simply ignored them, did not trust them, or some combination of the above.

Additional experimental research on cost-sharing, which makes use of the same traffic network in Chapter 10. A major difference between the two studies is that in the experiment reported in Chapter 10 choice between the three modes of transportation (i.e., private car, carpool, shuttle) was conducted under the *simultaneous protocol of play*, whereas the same choice reported in Chapter 11 was conducted under the *sequential protocol of play*, where the decisions are made one at a time in an exogenously determined order. One may argue that the sequential protocol is a more common scenario, where commuters decide which mode of transportation to take after observing the choices of other group members who preceded them in the sequence (think about a group of scholars disembarking from an airplane in their travel to the conference hall). As the authors notice, real-time observability of choices among modes of transportation are often feasible, and that observability of previous choices is at present facilitated by online communication devices and social media. Nevertheless, there is often a limit to the degree of observability, and consequently it is natural and instructive to study cases in which observability of the queue is unlimited and cases in which it is limited. So again, it is a matter of information. Simultaneous protocol setups impart zero information on choice observability depending on time lags between successive observations (think about scholars disembarking from several airplanes with close arrival times landing on the same afternoon), constraints on memory, or other situational circumstances. Theoretically, when choice observability is zero, the sequential protocol is identical to the simultaneous protocol, but psychologically the two protocols may be perceived quite differently.

The authors report an experiment on choice of mode of transportation under two between-subject conditions. In Condition PART, each player *j* (*j*=1, 2, ..., *n*) was only informed of the choice of the (*single*) player who preceded her in the sequence. Consequently, player *j* in Condition PART could not ascertain how many group members who preceded her in the sequence chose one of the two modes of ridesharing. In Condition FULL, each player *j* was informed of the choices of *all* the players who preceded her in the sequence. (Under both observability conditions, player *j* = 1 had no information at all, and player *j* = 2 had complete information). In equilibrium, under both conditions, all the *n* group members choose the efficient solution, namely, route O→SH→D. Convergence to this equilibrium depends on optimal best responses of all the *n* players under backward induction.

Most of the participants in both conditions found the ridesharing game quite daunting. Under a strict definition of convergence, 11 of the 18 groups (*n*=10 in each group) converged to the efficient equilibrium within 50 rounds of play regardless of the level of choice observability. The authors also report that although the information condition had no significant effect on the number of groups that converge to equilibrium, it slowed down the rate of convergence in Condition PART in comparison to Condition FULL.

The authors of Chapter 12 report another study that introduces a new setting for the BP that involves negative (i.e., congestion) as well as positive (i.e., cost-sharing) externalities at different levels of choice observability. The idea here is that negative and positive externalities may be imposed by choices of different routes of the network because negative externalities are generated by structural changes of the network, as we have seen before, whereas positive externalities are imposed by the choice of different modes of transportation on the same network. Easley and Kleinberger summarize this situation: “with positive externalities, there exist self-fulfilling expectations and a natural set of outcomes to coordinate on; with negative externalities, any shared expectation of a fixed audience size will be self-negating, and the individual must instead sort themselves out in much more complicated ways” (2010, p. 536).

In designing an experiment to operationalize such situations, the authors of Chapter 12 had two goals in mind. The first is to construct a directed network that includes some links with positive externalities and others with negative externalities by a judicious choice of group size, cost structure that gives rise to the BP, and a network architecture that is as simple as possible. The second purpose is the comparison of route choices in basic and augmented networks under either the simultaneous or sequential protocol of play. To achieve these two goals, the authors constructed two directed networks. The basic network included three routes, namely, routes O→A→D, O→B→D, and O→ D with corresponding link costs c(O→A) = m(O→A), c(B→D) = m(B→D), c(A→D) = c(O→B)= 21, and c(O→D) = m(O→D) = [90/m(O→D) + 10], where m(X→Y) is the number of players choosing link X→Y, $0 < m < n$. To construct the augmented network, they added link A→B to the basic network with a minimal positive cost of c(A→B) = 1. Under this construction of the network, routes O→A→D, O→B→D, and O→A→B→D are susceptible to congestion, whereas the choice of route O→D imposes a positive externality. For a detailed description of the equilibrium solutions in the four cells of the 2 ×2 experimental design (basic vs. augmented network by simultaneous vs. sequential protocol of play), the reader is referred to Chapter 12.

The experiment that we present in Chapter 12 has expended our research program on route choice in traffic network in two major ways. First, the comparison of the choice of route in the basic and augmented networks by the same players under the simultaneous protocol provided experimental evidence in support of the manifestation of the BP in a new class of directed networks with both positive and negative (called “mixed”) externalities. Second, route choices provided support for the manifestation of the BP under the sequential (and not only simultaneous) protocol of play. In fact, a comparison of the two protocols resulted in a higher increase in travel cost under the sequential than the simultaneous protocol of play once the basic network was augmented by link A→B. Under both protocols of play, the authors reported no support for welfare maximization. Overall, the choice of ridesharing on route O→D significantly decreases the social cost of travel. Yet, many of the group members playing the augmented network migrated to the congestible route O→A→B→D to increase their *personal* gain. The study reported in Chapter 12 opens new avenues for research on route choice decisions in directed networks, which allow a choice between ridesharing on the one hand and travel on congestible routes on the other hand under a wider scope of protocols of play.