

Degrading network capacity may improve performance: private versus public monitoring in the Braess Paradox

Eyran J. Gisches · Amnon Rapoport

Published online: 14 December 2010

© The Author(s) 2010. This article is published with open access at Springerlink.com

Abstract The Braess Paradox (BP) is a counterintuitive finding that degrading a network that is susceptible to congestion may *decrease* the equilibrium travel cost for each of its users. We illustrate this paradox with two networks: a basic network with four alternative routes from a single origin to a single destination, and an augmented network with six alternative routes. We construct the equilibrium solutions to these two networks, which jointly give rise to the paradox, and subject them to experimental testing. Our purpose is to test the generality of the BP when the network is enriched as well as the effects of the information provided to the network users when they conclude their travel. To this end, we compare experimentally two information conditions when each of the two networks is iterated in time. Under *public monitoring* each user is accurately informed of the route choices and payoffs of all the users, whereas under *private monitoring* she is only informed of her own payoff. Under both information conditions, over iterations of the basic and augmented games, aggregate route choices converge to equilibrium.

Keywords Braess Paradox · Congested traffic networks · Route choice · Experiment

Electronic supplementary material The online version of this article (doi:[10.1007/s11238-010-9237-0](https://doi.org/10.1007/s11238-010-9237-0)) contains supplementary material, which is available to authorized users.

E. J. Gisches (✉)

Department of Management Information Systems, Eller College of Management,
University of Arizona, 1130 E. Helen St., Tucson, AZ 85721, USA
e-mail: eyran@email.arizona.edu

A. Rapoport

Department of Management and Marketing, Anderson Graduate School of Management, University
of California Riverside, Riverside, CA 92521, USA
e-mail: amnonr@ucr.edu

1 Introduction

1.1 The Braess Paradox

The Braess Paradox (BP) (1968) is a major finding in the equilibrium analysis of route choice in directed networks that are susceptible to congestion. It demonstrates that degrading a network that is subject to congestion by removing one or more links may under certain circumstances *decrease* the cost of travel for all its users. The BP may be illustrated in networks modeled as non-atomic games (Roughgarden 2007) where the number of commuters is very large and, as a consequence, each commuter only controls a negligible fraction of the overall traffic. Alternatively, as in this study, it may be illustrated in networks modeled as atomic routing games, where each commuter has a non-negligible effect on the travel costs of all the other commuters. Our study focuses on directed traffic networks with a common origin and common destination in which each commuter imposes the same amount of traffic.

The two networks in Fig. 1 illustrate the BP. There are n commuters ($n = 18$ in this study) who independently have to choose a route $O \rightarrow D$ from vertex (node) O (for origin) to vertex D (for destination), while minimizing total cost of travel. Each link is assigned a cost function $c_{ij}(f_{ij})$ (also called “latency function”), which describes the cost of travel incurred by a commuter traversing link ij as a function of the number of commuters, f_{ij} , who travel on the same link. For example, the cost of travel on the route segment $O \rightarrow B$ in Fig. 1b is fixed at $c_{OB}(f_{OB}) = 210$; it is not affected by congestion. In contrast, the cost of traveling on link $B \rightarrow D$ is $c_{BD}(f_{BD}) = 10 \times f_{BD}$ which does depend on congestion. The four cost functions in Fig. 1b are instances of a *linear cost function* of the form $c_{ij} = a \times f_{ij} + b$, where $a, b \geq 0$.

Consider next the three-route network in Fig. 1a (left panel), referred to as the *augmented network*, and assume as before that $n = 18$. This network has five links and three alternative routes. Travel on the link $A \rightarrow B$ is costless: $c_{AB}(f_{AB}) = 0$. There is a unique equilibrium in pure strategies where all the $n = 18$ commuters independently choose the route $O \rightarrow A \rightarrow B \rightarrow D$ (Rapoport et al. 2009). The associated cost of travel in this equilibrium is $c_{OA}(18) + c_{AB}(18) + c_{BD}(18) = 360$. Next, remove the link $A \rightarrow B$ to construct the two-route network in Fig. 1b (right panel), referred to as the *basic network*. In the ensuing pure-strategy equilibrium for the basic network, $n/2 = 9$ commuters choose route $O \rightarrow A \rightarrow D$ and $n/2 = 9$ other commuters choose route $O \rightarrow B \rightarrow D$. The equilibrium cost of travel in this structurally degraded network decreases from 360 to 300. Thus, paradoxically, degrading the augmented net-




Fig. 1 Basic (b) and augmented (a) networks with two alternative routes

work by removing one of its links may improve its performance under equilibrium flow of selfish agents.

The BP may be compared to two other well-known paradoxes in non-cooperative game theory, namely, the Prisoner's Dilemma and Centipede (e.g., [Aumann 1992, 1995, 1998](#)) games. Considered jointly, these paradoxes deepen our understanding of the possible counterintuitive implications of the Nash equilibrium concept and the assumptions needed to sustain it. There is a major difference between the BP and the two other paradoxes. When the Prisoner's Dilemma game is played only once the analysis simply shows that the equilibrium solution is inefficient. There is nothing paradoxical about this result; Nash equilibria may not necessarily maximize social welfare. The paradoxical result emerges when this game is played repeatedly a finite and commonly known number of times and mutual defection on each round is derived by backward induction (e.g., [Luce and Raiffa 1957](#)). The paradoxical result of exiting on step 1 of the Centipede game is also derived by backward induction. In contrast to these two games, the counterintuitive outcome of the BP results from a *structural* change in the *architecture* of the network rather than the controversial process of backward induction.

1.2 Previous research

1.2.1 Theoretical

The discovery of the BP has generated a large body of research, almost all of it theoretical, in communication, transportation, and computer science. Researchers have proposed classification of networks in which the addition of one or more links could degrade performance ([Frank 1981](#); [Steinberg and Zangwill 1983](#)), discovered new paradoxes ([Arnott et al. 1993](#); [Cohen and Kelly 1990](#); [Dafermos and Nagurney 1984](#); [Pas and Principio 1997](#); [Smith 1978](#); [Steinberg and Stone 1988](#)), proved that detecting the BP in networks is algorithmically hard ([Roughgarden 2006](#)), and quantified the degree of degradation in network performance from unregulated traffic ([Koutsoupias and Papadimitriou 1999](#); [Roughgarden and Tardos 2002](#)). Almost all the existing theoretical work on the problem of detecting and illustrating the BP has considered networks similar to the one presented in Fig. 1a and b, other special networks, or focused on the special case where only a single link is added to the basic network. A notable exception is a class of Braess graphs studied by [Roughgarden \(2005\)](#).

[Steinberg and Zangwill \(1983\)](#) and [Taguchi \(1982\)](#) have proposed analytical conditions that partially characterize whether removing a given route would improve the equilibrium flow in a network. Those results were later generalized by [Dafermos and Nagurney \(1984\)](#). On the basis of their investigation, Steinberg and Zangwill concluded that the BP is a common rather than rare event. In a more recent article, [Valiant and Roughgarden \(2006\)](#) considered a class of large random networks with links that have random linear cost functions, a single origin, and a single destination. They proved that as the number of vertices in the random network becomes very large, the probability of detecting routes whose removal decrease travel time approaches one.

But is the BP a mere theoretical curiosity? As early as 20 years ago arguments were raised against the relevance of the BP to real-life situations. The argument goes that these are highly abstract networks and their seemingly paradoxical implications arise from the many aspects in which they differ from reality rather than from these aspects that they share with it (Cohen 1988). Valiant and Roughgarden have expressed a similar sentiment in writing “However, remarkably little is known about whether the BP is a common real-world phenomenon, or a mere theoretical curiosity” (2006, p. 2). If the BP is a rare event in selfish routing networks, restricted to judiciously chosen combinations of parameter values and very simple networks, then interest in it ought to be limited. After all, there are numerous theories in the social and physical sciences that yield curious results when applied to “pathological” cases. But if a substantial fraction of networks in communication and transportation are susceptible to the BP, then the problem of removing links from the augmented network or, alternatively, adding links to the basic network gains practical significance and should, therefore, be approached with considerable caution.

1.2.2 Empirical

On the empirical (i.e., field data) side, there has been mostly anecdotal evidence suggesting that the BP might actually have occurred in certain road networks. In a postscript to his exposition of the BP about 40 years ago, Murchland (1970) remarked that Knödel had noted that, in agreement with the BP, major road improvements in the center of the city of Stuttgart had failed to yield the benefits expected. In an article published with the provocative title “What if they closed 42nd street and nobody noticed?” (Kolata 1990), The New York Times also hinted at the counterintuitive consequences of road closures. Fisk and Pallottino (1981) provided examples that occurred on a modeled network of the city of Winnipeg and supported the BP. Additional empirical evidence has been reported in a recent article by Youn et al. (2008), who analyzed a large portion of the Boston road network with 246 directed links, and concluded that the “Braess’s paradox is more than academic curiosity” (2008, p. 4).

1.2.3 Experimental

A complementary approach is to simulate traffic networks that are susceptible to the BP (e.g., Fig. 1a and b) in the laboratory, have subjects independently choose routes on these networks before and after one or more links are removed from the augmented network or added to the basic network, and find out whether systematic and replicable patterns of behavior emerge and if they do, whether they support equilibrium play. If replicable patterns of behavior do emerge, then they may contribute to answering the question whether the BP is a significant phenomenon of real value or, alternatively, that it is devoid of practical interest. A major advantage of this approach is that the assumptions underlying the BP about the constant population size, parameters of the cost functions, and information structure can be implemented with precision. This is not the case in field studies (Morgan et al. 2009). Another advantage is that it allows the study of choice patterns and learning from experience by iterating the congestion game

with the same population of users while keeping the cost functions and information structure fixed over iterations.

Experiments designed to study the occurrence of the BP include [Aoki et al. \(2007\)](#); [Morgan et al. \(2009\)](#), and [Rapoport et al. \(2006, 2008, 2009\)](#). [Rapoport et al. \(2009, Experiment 1\)](#) studied the basic and augmented networks in Fig. 1b and 1a, respectively. In one of their studies, subjects first chose routes in 40 iterations of the basic network and then in 40 additional iterations of the augmented network. In a second study, the order of presentation of the basic and augmented networks was reversed. [Rapoport et al. \(2009\)](#) reported several major findings. First, the order of presentation of the two networks had no effect on the observed distribution of route choices. Second, an important determinant of performance in the experiment was persistent variability in traffic flow that decreased very slowly with experience. Thirdly, and most importantly, mean route choices in the augmented network slowly converged with experience to the pure-strategy equilibrium solution where all subjects choose route $O \rightarrow A \rightarrow B \rightarrow D$. Keeping the cost functions fixed but systematically varying the population size ($n = 10, 20, 40$ in three different conditions), [Rapoport et al. \(2006\)](#) reported additional evidence in strong support of the BP. Statistically significant, but weaker support for the BP was reported by [Aoki et al. \(2007\)](#), [Morgan et al. \(2009\)](#) and [Rapoport et al. \(2008, 2009, Exp. 2\)](#).

1.3 Designing the present study

Two main features of the design of the previous studies critically limit their generalizability. First, with the exception of [Rapoport et al. \(2009, Exp. 2\)](#), all of these experiments have focused on the simple case of a basic network with only two alternative routes. There is a large gap between the very simple networks studied in the laboratory (see, e.g., Fig. 1a and b) and the considerably richer networks examined theoretically by, for example, [Youn et al.](#) Whereas equilibrium play may become transparent when the network has only two routes, particularly if they are symmetric, this is not the case in networks with more routes. The second feature has to do with the type of monitoring at the end of each round when the network game is played repeatedly. Real-world traffic networks are often characterized by *private monitoring*, where each user only observes the traffic flow on the route segments that she actually traverses. But in the experiments reported above (with the exception of [Morgan et al.](#)), at the end of each round each subject was informed of the route segments chosen and associated costs incurred by *all* the users (*public monitoring*). A major shortcoming of these experiments resulting from the very simple structure of the networks that they implemented is that the effects of private and public monitoring could not be distinguished: even under private monitoring the information about the route choices and payoffs of all the users could easily and accurately be deduced.

When network games are repeated in time, public and private monitoring may result in very different dynamics. In fact, if the game is repeated in time, then the players may find a way to deviate from the one-shot equilibrium play and collectively choose routes that increase their payoff. We know from the folk theorem for finitely repeated games with public monitoring that an outcome, which is better for each user than the

Nash equilibrium of the one-shot game, can be implemented as an equilibrium of the repeated game. The folk theorem maintains that if the Nash equilibrium exceeds the minimax payoff,¹ as is the case in our design, then players can adopt many payoff profiles in the first stage of the repeated game and maintain them by the threat of holding deviating players to their minimax payoff. No such general result exists regarding the possibility of playing off Nash equilibrium strategies when monitoring is private (see, e.g., Mailath and Samuelson 2006 for a survey of repeated games and variations of the folk theorem under alternative ways of monitoring). Therefore, we will use the one shot equilibrium as the benchmark for evaluating the choices of our subjects. The same approach was used by Rapoport et al. and Morgan et al.

This study aims to overcome the two design shortcomings mentioned above. We examine patterns of route choice and experimentally test equilibrium solutions in a basic network with *four* (rather than two) alternative routes from a common origin to a common destination and then add to it *two* (rather than one) directed route segments resulting in an augmented network with *six* (rather than four) alternative routes. We use a within-group design with the same subjects independently choosing routes in both networks, and iterate the network games multiple times. Outcome information at the end of each round is manipulated in a between-subject design, with two conditions. In one condition of public monitoring each subject is fully informed at the end of each round of the route choices and payoffs of all group members. In the second condition of private monitoring, at the end of each round each subject is only informed of the number of users on her own route. Information about route choices and costs of travel of all the users may no longer be deduced if monitoring is private. Comparison of these two conditions should reveal how the type of monitoring affects route choice. If shown to manifest itself in a more complex network and under more stringent information conditions, then the BP would gain vital support of its practical importance to network design.

The rest of this article is organized as follows. Section 2 presents the model used in our study and the equilibrium solutions. Section 3 describes the experimental method. Section 4 reports the results, and Sect. 5 discusses them. Online Appendix A collects the proofs, and Online Appendix B presents the subject instructions.

2 Model

Figure 2 displays the two traffic networks that were presented to our subjects in the experiment. Figure 2b (right panel) exhibits the basic network that includes four routes with no cross roads connecting them, namely, routes $O \rightarrow A \rightarrow D$, $O \rightarrow E \rightarrow D$, $O \rightarrow F \rightarrow D$, and $O \rightarrow B \rightarrow D$. The cost functions associated with the segments of these routes are of three kinds: the constant cost function $c_{ij}(f_{ij}) = 130$, which is not susceptible to congestion; the cost function $c_{ij}(f_{ij}) = 10 \times f_{ij}$, which is moderately affected by congestion; and the cost function $c_{ij}(f_{ij}) = 20 \times f_{ij}$, which

¹ In Games 2a and 2b the minimax strategy would be $O \rightarrow F \rightarrow D$ or $O \rightarrow B \rightarrow D$ with payoff (in our design the payoff equals a reward of 290 minus travel cost) of -20 which is lower than the Nash equilibrium payoff of 50 (Game 2a) or 100 (Game 2b).




Fig. 2 Basic (b) and augmented (a) networks with four alternative routes

is more heavily affected by congestion. Routes $O \rightarrow A \rightarrow D$ and $O \rightarrow E \rightarrow D$ have the same total cost functions on road segments that are arranged in opposite way. The same holds for routes $O \rightarrow F \rightarrow D$ and $O \rightarrow B \rightarrow D$.

The basic network game has multiple pure-strategy equilibria in which routes $O \rightarrow A \rightarrow D$, $O \rightarrow E \rightarrow D$, $O \rightarrow F \rightarrow D$, and $O \rightarrow B \rightarrow D$ are independently chosen by 3, 3, 6, and 6 commuters, respectively. The multiple equilibria differ in the identity of the players choosing each route but have the same distribution of users over the four routes. The equilibria are Pareto unrankable. For each route, the individual cost of travel is $C = 190$. There is also a symmetric mixed-strategy equilibrium in which these four routes are chosen with probability 0.157, 0.157, 0.343, and 0.343, respectively. Under mixed-strategy equilibrium play, the expected cost of travel increases from 190 to 203.33. The upper panel of Table 1 summarizes the equilibrium solutions for the basic network. Proofs are presented in Online Appendix A.

To construct the augmented network, two directed links were jointly added to the basic network in Fig. 2b, namely, links $A \rightarrow E$ and $B \rightarrow F$, both of them costless (see Fig. 2a). Link $A \rightarrow E$ connects the heavily affected by congestion road segments of routes $O \rightarrow A \rightarrow D$ and $O \rightarrow E \rightarrow D$, and link $B \rightarrow F$ connects the two moderately affected by congestion road segments of routes $O \rightarrow B \rightarrow D$ and $O \rightarrow F \rightarrow D$. There are multiple pure-strategy equilibria in which the original four routes of the basic network are no longer chosen, and the two new routes $O \rightarrow A \rightarrow E \rightarrow D$ and $O \rightarrow B \rightarrow F \rightarrow D$ are chosen by 6 and 12 commuters, respectively. The individual equilibrium cost of travel on each of these three-segment routes increases from $C = 190$ to $C = 240$, an increase of 26.3%. There is also a symmetric mixed-strategy equilibrium in which the two routes $O \rightarrow A \rightarrow E \rightarrow D$ and $O \rightarrow B \rightarrow F \rightarrow D$ are chosen with respective probability of 0.314 and 0.686 (lower panel of Table 1). The expected cost of travel under symmetric mixed-strategy equilibrium play is 253.33.

3 Method

3.1 Subjects

One hundred and eighty undergraduate and graduate students at the University of Arizona volunteered to participate in a computer-controlled decision-making experiment for payoff contingent on their performance. Male and female students participated in almost equal proportions. The subjects were divided into 10 groups (sessions) of

Table 1 Summary of the equilibrium predictions for the basic and augmented networks in Fig. 2

Routes	$O \rightarrow A \rightarrow D$	$O \rightarrow E \rightarrow D$	$O \rightarrow F \rightarrow D$	$O \rightarrow B \rightarrow D$	$O \rightarrow E \rightarrow A \rightarrow D$	$O \rightarrow B \rightarrow F \rightarrow D$	Equilib. cost
<i>Basic network (Game 2b)</i>							
Pure st. equilibrium	3	3	6	6	NA	NA	190
Sym. mixed st. equilib.	0.157	0.157	0.343	0.343	NA	NA	203.33
<i>Augmented network (Game 2a)</i>							
Pure st. equilibrium	0	0	0	0	6	12	240
Sym mixed st. equilib.	0	0	0	0	0.314	0.686	253.33

In both networks, each session includes $n = 18$ subjects

$n = 18$ members each. Five groups participated in the public monitoring (PUBLIC) condition in which participants were fully informed at the end of each round of the number of travelers on all the links and their costs of travel. Five other groups participated in the private monitoring (PRIVATE) condition in which each participant was only informed of the number of network users traveling on the same road segments that she chose and the cost of travel on this route. Each session lasted about 100 min. Excluding a \$5 show-up bonus, the mean payoff across all the sessions was \$15.57.

3.2 Procedure

The experimental sessions were conducted in a computerized laboratory with multiple terminals located in separate cubicles. After entering the laboratory, each subject drew a poker chip that determined her seating. At any point during the experiment questions about the procedure were answered individually by the experimenter.

Each session was divided into Parts I and II. Written instructions were handed at the beginning of each part and the subjects could read them at their own pace (see Online Appendix B for the instructions of Condition PRIVATE). The instructions for Part I presented and explained the augmented network game in Fig. 2a and the procedure for choosing one of the six routes. The instructions for Part II did the same for the basic network game in Fig. 2b. The instructions for Conditions PUBLIC and PRIVATE only differed from each other in the description of the type of monitoring at the end of each round. In Part I, the subjects in all the 10 sessions chose routes in the augmented network that was repeated for 60 identical rounds. After completing Part I, they were handed instructions for Part II and asked to choose routes in the basic network for 60 additional rounds. Each set of instructions exhibited the network, explained the cost functions, illustrated the computation of the travel cost, and described the procedure for choosing a route. To choose a route, the subject had to click on all the links comprising it. After all the 18 members of each group chose their routes, a new screen was displayed with full information (Condition PUBLIC) presented in tabular form about the number of participants choosing each of the routes and their payoff for the round. In addition, the number of travelers on each link was exhibited on the network diagram. In Condition PRIVATE, only information about the number of users choosing the same route as the subject and her payoff for the round was presented.

On each round in both Parts I and II, each participant was provided with a reward (endowment) of $R = 290$ payoff units. Individual payoff for the round was calculated by *subtracting* the travel cost from the reward. The value of R was chosen deliberately so that, in equilibrium, a player wins $290 - 190 = 100$ payoff units in Part II and only $290 - 240 = 50$ units in Part I. Consequently, individual payoff in the basic network was set to be twice as large as the payoff in the augmented network, thereby considerably enhancing the effect of the BP. With the value of the reward set at $R = 290$, on any particular round participants could either gain or lose money. For example, if the group were split with nine subjects choosing route $O \rightarrow A \rightarrow E \rightarrow D$ and nine others choosing route $O \rightarrow B \rightarrow F \rightarrow D$ in the augmented network (Fig. 2a), then nine participants would have lost 70 units each and nine others would have gained 110 units.

The participants were paid their earnings at the end Part II. Payoffs were accumulated across 12 randomly chosen rounds (six in each information condition) and converted to money at the rate of 60 units = \$1.00. The 12 payoff rounds were randomly drawn in front of all the participants only after the experiment was completed to eliminate any wealth effects. Overall, each participant in any session ended up with a positive gain, so that the problem of overall negative earnings was not encountered.

Four features of the design warrant brief discussion. First, to eliminate reputation effects, communication between the players was prohibited. Second, we opted for a within-subject design in playing the basic and augmented games. If the effect of the BP is realized under a within-subject design with the same subjects participating in both network games, then the effect should prove much more convincing. Third, the same reward of $R = 290$ was assigned to each subject on each round in both parts of the session. In general, when only costs are involved and the cost functions are all linear, then adding a link may increase travel cost by at most 1/3 (Roughgarden and Tardos 2002). In our design, where costs are subtracted from a constant reward, the effect of the BP is tripled.² Fourth, we had the subjects experience the augmented network first.³ Had they played the basic network first, then avoiding routes $O \rightarrow A \rightarrow E \rightarrow D$ or $O \rightarrow B \rightarrow F \rightarrow D$ in the augmented network might have been attributed to previous experience. Playing the augmented network first with no prior experience negates this possibility.⁴

4 Results

The results are organized in four separate sections. Section 4.1 reports the results of tests of hypotheses about route choice behavior in Games 2a and 2b. Section 4.2 presents analyses of the mean payoff by condition across all the 120 rounds of play. Section 4.3 focuses on individual differences. The analyses reported in Sects. 4.1, 4.2, and 4.3 jointly provide strong evidence of learning: mean frequencies of choice in the augmented network and mean payoffs slowly converge with experience to equilibrium. To account for these patterns of behavior and the effects of experience on the distributions of route choice under private and public monitoring, we propose a simple, reinforcement-based learning model in Sect. 4.4 and test its implications.

When group members repeatedly interact with one another over multiple iterations of the stage game, the statistical unit of analysis is the group, not the individual player. Consequently, we conduct non-parametric tests on group statistics to compare the

² When expressed in terms of cost of travel, the move from the augmented network in Part I to the basic network in Part II reduces the cost from 240 to 190 or by 21%. However, by subtracting the cost of travel from a fixed reward the same change increases the profit from 50 to 100, or by 100%. This resulting effect is three times the maximal theoretical cost increase as shown by Roughgarden and Tardos.

³ In an alternative within-subject design, participants first play the iterated basic network game in Part I and then the iterated augmented game in Part II. This alternative design examines the effect of *adding* links to the basic network rather than *deleting* the same links from the augmented network. Rapoport et al. (2009) compared the two designs to each other in Games 1a and 1b and reported no significant differences between them.

⁴ In Part I players could still refrain from choosing route $O \rightarrow A \rightarrow E \rightarrow D$ or route $O \rightarrow B \rightarrow F \rightarrow D$ thereby avoiding negative consequences of the BP.

two conditions to each other. As the number of groups in each condition is relatively small, these tests are not powerful. Therefore, in a few cases we relax the assumption of subject independence and conduct more powerful tests assuming that subjects within group are mutually independent. In justification of this assumption, we note that because the groups are relatively large and reputation building is not possible, the effect of each individual player on the decisions of the other group members is sufficiently small to safely be neglected.

4.1 Route choice

Table 2 presents the mean frequency of route choices. The upper panel presents the results for the augmented network, and the lower panel for the basic network. In each panel, the mean frequencies of choice are shown separately for Conditions PUBLIC and PRIVATE, and within each condition they are displayed separately by session. For each session and each route, the mean frequencies of route choice are computed across all the 60 iterations of the stage game. Overall means across sessions are presented in boldface. The bottom three rows of each panel display the pure-strategy equilibrium solution, the W statistic of the Wilcoxon two-sided rank sum test for comparing the two experimental conditions, and the associated significance level.

Inspection of the overall mean frequency of route choice across the five sessions in the upper panel of Table 2 suggests no systematic differences between the two monitoring conditions in the augmented network. Of the six statistical tests only one resulted in a significant difference, with subjects in Condition PRIVATE choosing, on average, route $O \rightarrow B \rightarrow D$ more often (1.347) than the subjects in Condition PUBLIC (0.900). Analysis of the overall mean frequency of route choice across the five sessions in the lower panel of Table 2 yielded the same conclusion for the basic network. For each of the four routes, the session means for the two conditions are almost identical. All four tests for Game 2b resulted in non-significant differences between the two monitoring conditions. Considered jointly, the results of the 10 tests cannot reject the null hypothesis of equality of mean route choice in Conditions PUBLIC and PRIVATE.

We turn next to an investigation of the dynamics of route choice and tests of the equilibrium solutions. Figure 3 exhibits the mean frequency of route choice for each route separately and each round of play in Condition PUBLIC. Mean frequencies computed across the five sessions of Game 2a are displayed in the upper panel of Fig. 3, and mean frequencies for Game 2b in the lower panel. Recall that in the equilibrium solution for Game 2b (lower panel) routes $O \rightarrow A \rightarrow D$ and $O \rightarrow E \rightarrow D$ are each chosen by 3 players and routes $O \rightarrow F \rightarrow D$ and $O \rightarrow B \rightarrow D$ are each chosen by six players. The results in Table 2 (lower panel) support this prediction (mean frequency equals 5.727 and 6.107 for routes $O \rightarrow F \rightarrow D$ and $O \rightarrow B \rightarrow D$ and 3.087 and 3.080 for routes $O \rightarrow A \rightarrow D$ and $O \rightarrow E \rightarrow D$). The lower panel of Fig. 3 exhibits no evidence for learning in Part II. This is supported by the non-significant correlations between mean frequency of route choice and round number. They assume the values $r = 0.050$ ($p = 0.704$), $r = -0.003$ ($p = 0.983$), $r = 0.160$ ($p = 0.222$), and $r = -0.169$ ($p = 0.198$), for routes $O \rightarrow A \rightarrow D$, $O \rightarrow E \rightarrow D$, $O \rightarrow F \rightarrow D$, and $O \rightarrow B \rightarrow D$, respectively.

Table 2 Mean frequency of route choice by condition (PUBLIC or PRIVATE), session (1–5), network type (Games 2a and 2b), and route of travel across rounds 1–60

Condition	Session	Route $O \rightarrow A \rightarrow D$	Route $O \rightarrow E \rightarrow D$	Route $O \rightarrow F \rightarrow D$	Route $O \rightarrow B \rightarrow D$	Route $O \rightarrow A \rightarrow E \rightarrow D$	Route $O \rightarrow B \rightarrow F \rightarrow D$
<i>Augmented network (Game 2a)</i>							
PUBLIC	1	0.883	0.733	1.300	0.817	5.083	9.183
	2	1.083	1.350	1.433	1.033	4.033	9.067
	3	0.633	0.900	0.833	0.733	4.733	10.167
	4	0.883	1.267	1.133	0.850	4.317	9.550
	5	0.683	0.817	1.367	1.067	4.800	9.267
Overall	0.833	1.013	1.213	0.900	4.593	9.447	9.447
PRIVATE	1	0.683	1.000	1.550	1.233	4.217	9.317
	2	0.867	1.267	1.700	1.350	3.900	8.917
	3	1.250	1.233	1.400	1.217	3.617	9.283
	4	1.033	0.950	1.350	1.350	4.167	9.150
	5	0.833	1.017	1.633	1.583	4.300	8.633
Overall	0.933	1.093	1.527	1.347	4.040	9.060	9.060
Equilib.	0	0	0	0	6	12	12
W	25.5	24.5	18.0	15.0	37.0	33.0	33.0
p	NS	NS	NS	0.008	NS	NS	NS
<i>Basic network (Game 2b)</i>							
PUBLIC	1	3.183	2.967	5.833	6.017	–	–
	2	3.217	3.183	5.900	5.700	–	–
	3	3.133	2.967	5.567	6.333	–	–
	4	2.850	3.150	5.467	6.533	–	–
	5	3.050	3.133	5.867	5.950	–	–
Overall	3.087	3.080	5.727	6.107	–	–	–

Table 2 continued

Condition	Session	Route					
		$O \rightarrow A \rightarrow D$	$O \rightarrow E \rightarrow D$	$O \rightarrow F \rightarrow D$	$O \rightarrow B \rightarrow D$	$O \rightarrow A \rightarrow E \rightarrow D$	$O \rightarrow B \rightarrow F \rightarrow D$
PRIVATE	1	2.700	3.150	6.000	6.150	-	-
	2	3.133	3.183	5.983	5.700	-	-
	3	3.133	2.967	5.883	6.017	-	-
	4	3.117	3.050	5.800	6.033	-	-
	5	2.983	3.083	5.350	6.583	-	-
	Overall	3.013	3.087	5.803	6.097	-	-
Equilib.		3	3	6	6	-	-
W		32.0	27.0	24.0	26.0	-	-
p		NS	NS	NS	NS	-	-

Wilcoxon two-sided rank sum test

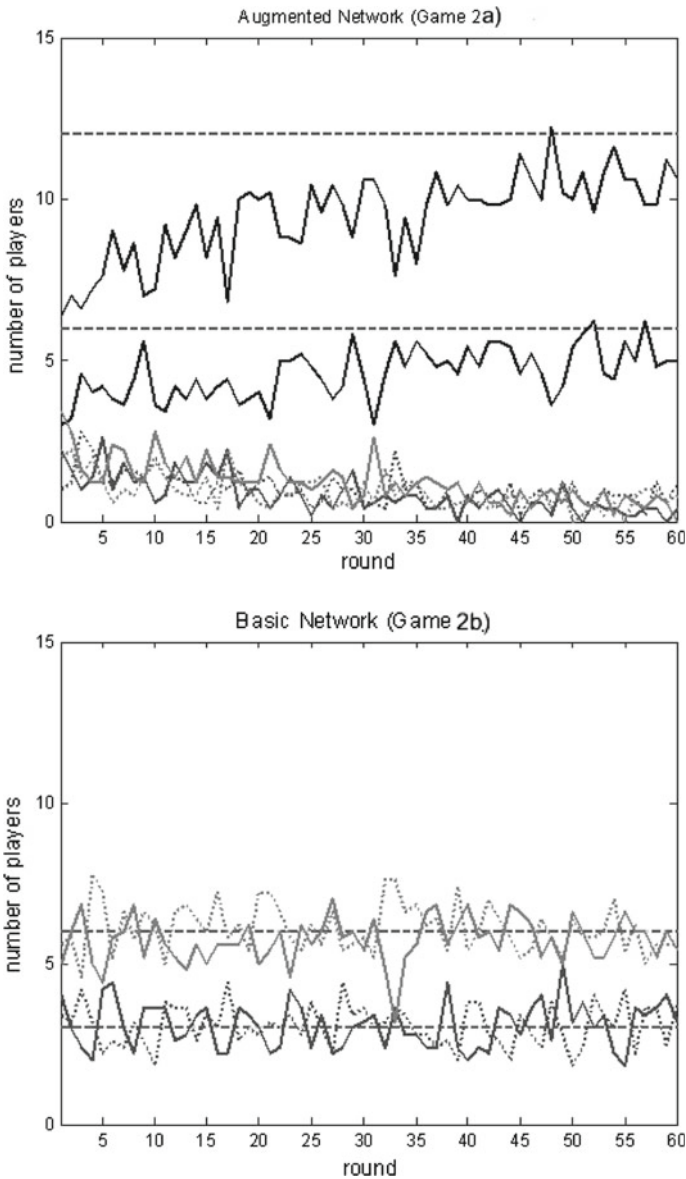


Fig. 3 Mean observed route choice across sessions by network type (Games 2a and 2b) and round of play: Condition PUBLIC

In contrast, the mean frequencies of route choice in Game 2a (upper panel) exhibit strong evidence of learning. The mean frequencies for the “original” four routes of the basic network ($O \rightarrow A \rightarrow D$, $O \rightarrow E \rightarrow D$, $O \rightarrow F \rightarrow D$, and $O \rightarrow B \rightarrow D$) decrease by a factor of 3.6 from 2.15 on round 1 to 0.60 on round 60. The mean frequency of route $O \rightarrow A \rightarrow E \rightarrow D$ increases from 6.4 on round 1 to 10.6 on

round 60 and slowly converges from below to the equilibrium solution. Similarly, the mean frequency of route $O \rightarrow B \rightarrow F \rightarrow D$ increases from 3 on round 1 to 5 on round 60.

Additional evidence for equilibrium play in Condition PUBLIC comes from an analysis of the individual, rather than aggregate, route choice on the last 10 rounds. For each subject separately, we counted the number of times she chose route $O \rightarrow A \rightarrow E \rightarrow D$ (f_{OAED}), route $O \rightarrow B \rightarrow F \rightarrow D$ (f_{OBDF}), and any of the four original routes of the basic network (f_{other}), where $f_{OAED} + f_{OBDF} + f_{\text{other}} = 10$. Our results show that $f_{\text{other}} = 0$ for 43 of the 90 subjects and $f_{\text{other}} \leq 2$ for 17 additional subjects. Jointly taken, these results imply that the non-zero frequencies of choice of the four routes in the last block of Game 2a are mostly due to a minority of the subjects.⁵ In support of equilibrium play, the majority of the subjects abandoned the four “original” routes and divided their choices between the two “new” routes $O \rightarrow A \rightarrow E \rightarrow D$ and $O \rightarrow B \rightarrow F \rightarrow D$ in approximately the 1:2 ratio predicted by the equilibrium solution.

Arguably, the most important finding of this study is exhibited in Fig. 4 that displays the same information as Fig. 3 for Condition PRIVATE. A comparison of Figs. 3 and 4 suggests no major differences due to the type of monitoring. The only difference is that convergence to equilibrium play is slower in Condition PRIVATE than in Condition PUBLIC. The two conditions do not significantly differ from each other in the last block of 10 rounds. Table 3 presents the mean frequency of route choice for rounds 51–60. The effects of learning may be observed by comparing Tables 2 and 3. With only two exceptions (routes $O \rightarrow F \rightarrow D$ and $O \rightarrow B \rightarrow D$), all the non-parametric tests that compare the two conditions in the last block of rounds (for both Games 2a and 2b) do not reject the null hypothesis of equality of means.

4.2 Payoffs

Public monitoring in the traffic network games is clearly detrimental; on average, the subjects in Condition PRIVATE earned 11.3% more than the subjects in Condition PUBLIC in Game 2a and 2.8% more in Game 2b. This occurred not only in Game 2a where, as shown in the previous section, subjects in Condition PRIVATE exhibited a slower rate of learning than subjects in Condition PUBLIC, but also in Game 2b where equilibrium was already reached on the first few rounds of play in both conditions. Table 4 presents the evidence. It displays the mean payoff across the five sessions for Games 2a (upper panel) and 2b (lower panel). In each panel, mean payoffs are presented for blocks of 10 rounds (i.e., rounds 1–10, 11–20, ..., 51–60). The bottom two rows of each panel show the Wilcoxon test statistic W and the associated level of significance, p . In each panel and each block in Table 4, the mean payoff in Condition PRIVATE is higher than the mean payoff in Condition PUBLIC. In each panel, three of the six tests yielded significant differences between the two means. Analysis of the overall individual payoffs supports this conclusion. For each game type separately, we

⁵ These subjects conceivably persisted in choosing these routes in order to signal to the other group members to mimic their choices and thereby maximize social welfare.

Table 3 Mean frequency of route choice by condition (PUBLIC or PRIVATE), session (1–5), network type (Games 2a and 2b), and route of travel across rounds 51–60

Condition	Session	Route $O \rightarrow A \rightarrow D$	Route $O \rightarrow E \rightarrow D$	Route $O \rightarrow F \rightarrow D$	Route $O \rightarrow B \rightarrow D$	Route $O \rightarrow A \rightarrow E \rightarrow D$	Route $O \rightarrow B \rightarrow F \rightarrow D$
<i>Augmented network (Game 2a)</i>							
PUBLIC	1	0.20	0.50	0.60	0.70	5.80	10.20
	2	0.50	1.00	0.60	0.70	4.60	10.60
	3	0.20	0.70	0.50	0.50	5.20	10.90
	4	0.20	1.10	0.60	0.60	5.10	10.40
	5	0.40	0.60	0.40	0.40	5.60	10.60
Overall		0.30	0.78	0.54	0.58	5.26	10.54
PRIVATE	1	0.10	0.70	1.30	0.90	5.20	9.80
	2	0.20	0.80	0.80	0.80	4.70	10.70
	3	2.00	0.20	1.00	0.60	4.10	10.10
	4	0.40	0.70	0.90	0.80	5.20	10.00
	5	0.20	0.60	1.20	1.40	5.80	8.80
Overall		0.58	0.6	1.04	0.90	5.00	9.88
Equilib.		0	0	0	0	6	12
W		28.50	30.50	15	17.50	29.50	36
p		NS	NS	0.008	0.048	NS	NS
<i>Basic network (Game 2b)</i>							
PUBLIC	1	3.10	3.30	5.40	6.20	–	–
	2	3.70	3.60	5.70	5.00	–	–
	3	3.30	2.80	5.90	6.00	–	–
	4	2.80	3.00	5.30	6.90	–	–
	5	3.10	3.60	6.40	4.90	–	–
Overall		3.20	3.26	5.74	5.80	–	–

Table 3 continued

Condition	Session	Route $O \rightarrow A \rightarrow D$	Route $O \rightarrow E \rightarrow D$	Route $O \rightarrow F \rightarrow D$	Route $O \rightarrow B \rightarrow D$	Route $O \rightarrow A \rightarrow E \rightarrow D$	Route $O \rightarrow B \rightarrow F \rightarrow D$
PRIVATE	1	2.60	3.50	5.80	6.10	-	-
	2	3.60	2.80	5.60	6.00	-	-
	3	3.10	3.10	5.90	5.90	-	-
	4	3.40	3.00	5.40	6.20	-	-
	5	3.60	2.70	5.40	6.30	-	-
Overall	3.26	3.02	5.62	6.10	-	-	-
Equilib.	3	3	6	6	-	-	-
W	26	33	28.5	25	-	-	-
p	NS	NS	NS	NS	NS	-	-

Wilcoxon two-sided rank sum test




Fig. 4 Mean observed route choice across sessions by network type (Games 2a and 2b) and round of play: Condition PRIVATE

compared the overall individual payoffs (90 in each condition) between the two conditions by the Wilcoxon–Mann–Whitney test. The null hypothesis that the two groups of subjects are drawn from the same population was rejected in Game 2a ($z = -4.553$, $p < 0.001$) and Game 2b ($z = -2.970$, $p < 0.01$).

Table 4 Mean payoff across sessions by network type (Games 2a and 2b), condition (PUBLIC and PRIVATE), and blocks of 10 rounds

Condition	Rounds 1–10	Rounds 11–20	Rounds 21–30	Rounds 31–40	Rounds 41–50	Rounds 51–60	Overall	Equil.
<i>Augmented network (Game 2a)</i>								
PUBLIC	62.456	59.678	57.200	54.811	50.833	54.356	56.556	50
PRIVATE	66.956	70.256	64.900	63.767	56.333	55.389	62.933	50
<i>W</i>	91.0	68.0	77.0	69.5	90.5	103.0	2886.5	–
<i>p</i>	NS	0.006	0.038	0.008	NS	NS	<0.001	–
<i>Basic network (Game 2b)</i>								
PUBLIC	89.067	90.533	92.611	92.100	90.900	90.611	90.970	100
PRIVATE	90.878	92.589	93.644	94.300	94.344	95.178	93.489	100
<i>W</i>	87.0	80.0	94.5	78.0	72.5	61.5	2715.0	–
<i>p</i>	NS	NS	NS	0.045	0.015	0.001	<0.001	–

Wilcoxon two-sided rank sum test

Figure 5 exhibits the mean payoff across the five sessions by round of play. The top panel displays the results for Condition PUBLIC and the bottom panel for Condition PRIVATE. The equilibrium payoffs (100 and 50 for Games 2b and 2a, respectively) are portrayed as straight horizontal lines. Both panels show that the mean payoffs in Game 2b approach 100 from below (they cannot exceed 100) and the mean payoffs in Game 2a approach it from above.⁶ For each condition and game type separately, we computed the Spearman correlation between the mean payoff across the five sessions and the round number. The correlations for Game 2a were negative and significant: $r = -0.354$ ($p < 0.01$) and $r = -0.629$ ($p < 0.001$) for Conditions PUBLIC and PRIVATE, respectively. The correlations for Game 2b were positive and significant for Condition PRIVATE ($r = 0.441$, $p < 0.001$) but not significant for Condition PUBLIC. We repeated the same correlation analysis for each session separately and obtained essentially the same pattern in each of the 10 sessions.⁷ We conclude that with experience in repeatedly playing the basic game

⁶ Individual payoff in the basic network (Game 2b) ranges from -200 (290–490 when all the 18 group members choose either route $O \rightarrow A \rightarrow D$ or $O \rightarrow E \rightarrow D$) to 150 (290–140 when only a single subject chooses either route $O \rightarrow F \rightarrow D$ or $O \rightarrow B \rightarrow D$). However, the mean payoff in Game 2b cannot exceed the equilibrium payoff of 100. Individual payoff in the augmented network (Game 2a) ranges between -200 (when all the 18 group members choose either route $O \rightarrow A \rightarrow D$ or $O \rightarrow E \rightarrow D$) to 270 (290–20) when only a single subject chooses the route segments $O \rightarrow B$, $B \rightarrow F$, and $F \rightarrow D$. The mean payoff may be higher or smaller than the equilibrium payoff of 50.

⁷ The five correlations for Sessions 1, 2, 3, 4, and 5 in Game 2a of Condition PUBLIC were all negative and assumed the values -0.461 ($p < 0.001$), -0.305 ($p < 0.05$), -0.394 ($p < 0.01$), -0.071 (NS), and -0.302 ($p < 0.05$). The five correlations in Game 2b assumed the values -0.010 (NS), -0.039 (NS), 0.213 (NS), 0.129 (NS), and -0.068 (NS). The corresponding correlations for Condition PRIVATE were -0.423 ($p < 0.001$), -0.494 ($p < 0.001$), -0.304 ($p < 0.05$), -0.320 ($p < 0.05$), and -0.428 ($p < 0.001$) for Game 2a (all negative), and 0.336 ($p < 0.01$), 0.247 (NS), 0.231 (NS), 0.238 (NS), and 0.009 (NS) for Game 2b (all positive).

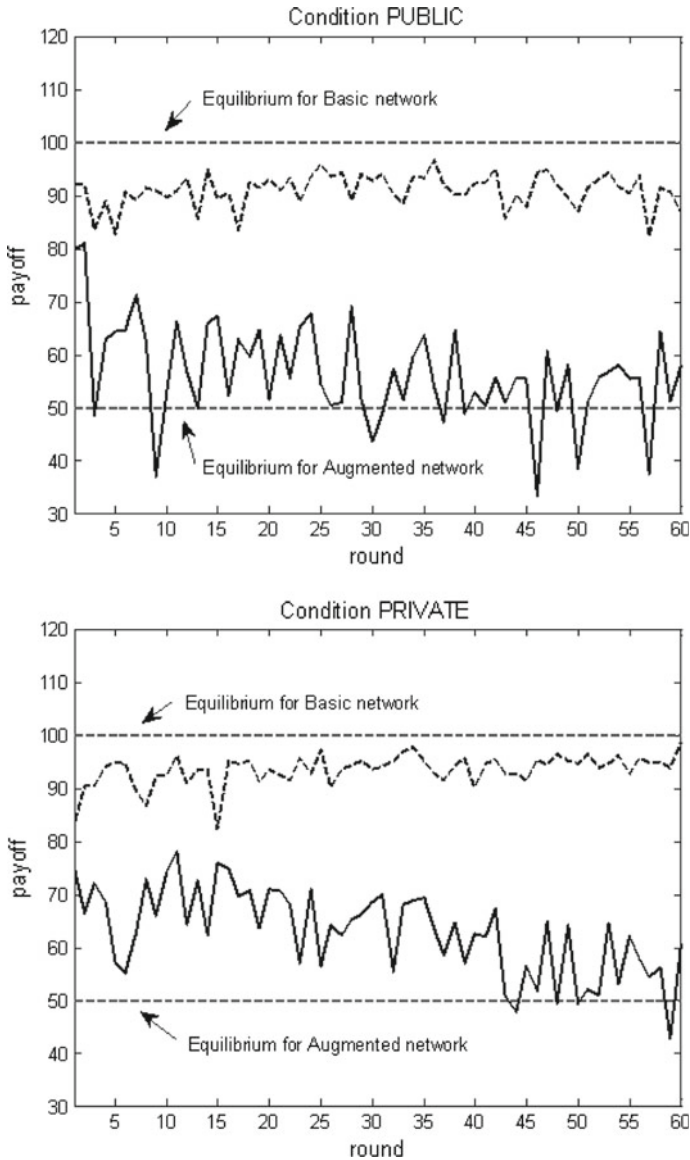


Fig. 5 Mean payoff by network type (Games 2a and 2b) condition (*top panel* Condition PUBLIC, *bottom panel* Condition PRIVATE)

subjects learn to better distribute themselves over the four routes and thereby increase their payoffs. In contrast, repeated selfish play in the augmented network resulted in subjects falling in the BP trap and thereby reducing their payoffs by almost a factor of 2.

4.3 Heterogeneous individual profiles

Figure 6 exhibits the individual profiles of the 18 players who participated in Session 4 of Condition PUBLIC. This session is typical; it is representative of the individual profiles of route choice in the other sessions in both conditions. The horizontal axis in each profile displays the round number and the vertical axis the six actual routes. Each sub-graph exhibits the 60 route choices in the basic and augmented networks. Inspection of the individual profiles reveals considerable individual differences. For example, Player 18 never switched her route across all the 60 rounds in the basic game (always choosing route $O \rightarrow B \rightarrow D$) and all the 60 rounds in the augmented game (always choosing route $O \rightarrow B \rightarrow F \rightarrow D$). In contrast, Players 1, 4, 9, and 10 switched their routes on almost every round in both networks. Players 6, 7, and 15 switched their routes in the augmented network much more often than in the basic network. Individual profiles of route choice defy simple classification. It is then even more surprising that this combination of heterogeneous profiles, or similar combinations in the other sessions, converged to equilibrium play in both the basic and augmented networks.

4.4 Learning

A major aim of this study was to test the predictive power of the pure-strategy equilibrium solution. Consequently, our major emphasis has been on aggregate rather than individual choice behavior. However, individual patterns of play are important for understanding the group dynamics because learning takes place on the individual rather than aggregate level.

Why do some network users switch their routes very frequently and others do not? One explanation for networks with only two alternative routes is in terms of different types of players. Selten et al. (2004) identified two types of players called *myopic* and *contrarians*. Both of them base their decision on the route choices on round t , not on the outcomes of earlier rounds. *Myopic* subjects switch from route j to route j' if the payoff associated with traveling on route j' on round t exceeds the payoff associated with traveling on route j on the same round. *Contrarian* subjects do not switch because they believe that the number of myopic subjects is sufficiently large that staying on the same route j on round $t+1$ is their best response. More generally, one may account for switching/not switching of routes in terms of a cognitive hierarchy model (e.g., Camerer et al. 2004) where level 0 players are myopic, and players using $K > 1$ steps of reasoning anticipate the decisions of lower-step thinkers and best-respond to the mixture of their decisions.

An alternative explanation would be that subjects develop preference for some routes based on their travel experience. A common way to describe this adaptive process is by using a reinforcement learning model (e.g., Camerer and Ho 1999). Our results suggest this type of model would better describe the observed behavior for several reasons. First, and most importantly, in Condition PRIVATE anticipatory behavior is not possible thus ruling out the possibility of applying most kinds of hierarchical level thinking or at least sophisticated versions of them. The comparable behavior

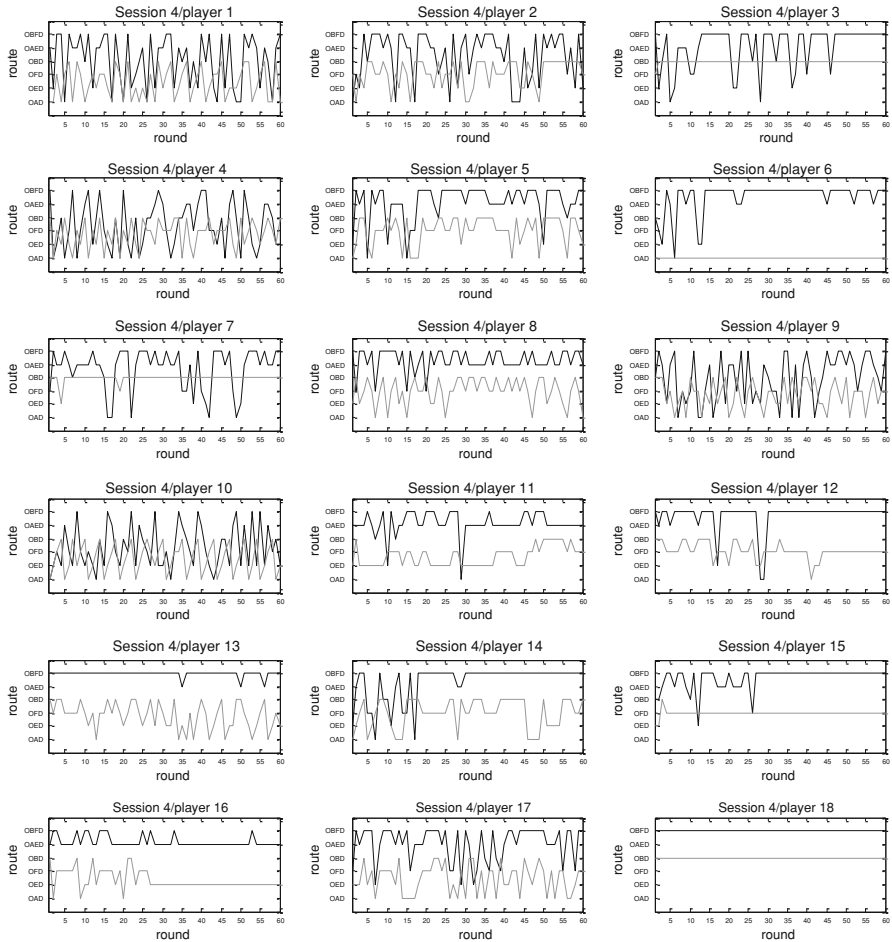


Fig. 6 Individual profiles of route choice of the subjects in Session 4 of Condition PUBLIC

patterns in both information conditions suggest that a similar adaptive model might have been used in both. Second, in our games, a (simple) hierarchical model would result in only four possible route distributions in the basic game and six distributions in the augmented network (resulting from the choice of the “naïve” population segment and the consequent best responses by the higher levels). Our data present much richer choice patterns which may be reproduced with a simple reinforcement learning type model. Third, we observe similar correlations between individual payoff and frequency of switches suggesting payoff affects (negatively) players’ tendency to switch routes. Hereafter, we attempt to account for the observed *aggregate* behavior using *individual* level reinforcement learning model.

To explain the dynamics of play exhibited in Figs. 3 and 4, we have searched for a parsimonious reinforcement learning model that *simultaneously* accounts for the aggregate route choice in both conditions and both networks. We chose the

Experience-Weighted Attractions (EWA) model of Camerer and Ho (1999) because of its generality, simplicity, and predictive success. In the EWA model, the attraction of each strategy is updated on every round of play by the payoff gained from choosing it. These attractions are then translated into probabilities of choosing each strategy.

For player i , let $\pi_i^j(s_i(t), s_{-i}(t))$ denote the payoff earned by player i on trial t for choosing strategy j , where $s_i(t)$ is the strategy chosen by player i , and $s_{-i}(t)$ are the strategies chosen by all other players. The reinforcement that player i attaches to strategy j is denoted by

$$R_i^j(t) = [\delta + (1 - \delta) \cdot I(s_i^j, s_i(t))] \cdot [\pi_i^j(s_i(t), s_{-i}(t)) - \pi_0],$$

where δ is a weight given to the strategies not chosen,⁸ $I(x, y)$ is an indicator (or bias) function which equals 1 if $x = y$ (i.e., strategy actually chosen by the player) and 0 otherwise, and π_0 is an aspiration level.⁹ This reinforcement is then used to update strategy j 's attraction by

$$A_i^j(t) = A_i^j(t - 1) + R_i^j(t).$$

On round $t+1$, player i stochastically determines her route choice using a logistic response function $p_i^j(t + 1) = e^{\lambda \cdot A_i^j(t)} / \sum_l e^{\lambda \cdot A_i^l(t)}$, where λ is an attraction sensitivity estimator.

We systematically searched for a parameter combination that best describes the observed route choice in the two conditions and two networks.¹⁰ Figures 7 and 8 exhibit the aggregate route choice for both types of game and both information conditions. The theoretical mean route choice functions in Figs. 7 and 8 closely resemble the observed mean route choice functions in Figs. 3 and 4. They illustrate that a simple and parsimonious reinforcement-based learning model is sufficient to account for the major findings of this study regarding route choice. These include:

1. Converge to equilibrium behavior in the two basic networks that is reached within a very few trials.
2. Mean route choices across the 60 rounds that approach, although never reach, equilibrium behavior in the two augmented networks.
3. No discernible differences in mean route choice between the two information conditions except for quicker convergence to equilibrium play in the augmented network in Condition PUBLIC.

⁸ In the PRIVATE condition where players were informed only about their own payoff, $\delta = 0$ for all the routes not chosen.

⁹ The original EWA model does not incorporate an aspiration level. However, Borgers and Sarin (2000) among others have suggested including it in learning models. We set this aspiration level π_0 to be the equilibrium payoff of 100 in the basic network and 50 in the augmented network.

¹⁰ We set the initial attractions at $A_i^j(t = 1) = (3/18) \times 290 \times N_1$, where 290 is the size of the reward (endowment) and N_1 is the mean number of subjects who chose route j on round 1. We also set $\delta = 0.5$ and $\lambda = 0.005$ for all the subjects in each session.

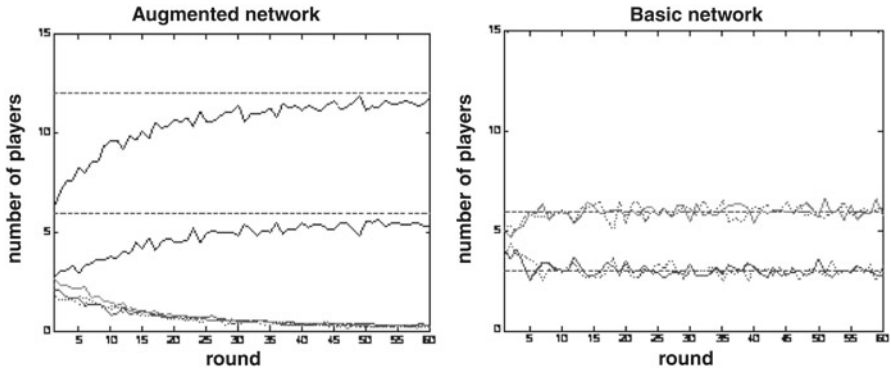


Fig. 7 Predicted mean route choice in Condition PUBLIC by the EWA learning model




Fig. 8 Predicted mean route choice in Condition PRIVATE by the EWA learning model

The learning functions in Figs. 7 and 8 are considerably smoother than the ones in Figs. 3 and 4 and, therefore, do not fully capture the trial-to-trial variability in the observed choices. An error parameter that would have increased variability could be introduced, but we opted not to do so in order to keep the number of parameters at a minimum.

5 Discussion

Our study is concerned with selfish routing in directed congestible networks where each player controls a non-negligible amount of traffic. It makes two contributions of potential importance to network design. First, we provide experimental evidence that, in agreement with BP, degrading a fairly complex network by simultaneously removing two of its links may decrease the equilibrium cost of traversing this network for all of its users. Support for equilibrium behavior on the aggregate, but not individual, level is obtained in a within-subject experimental design where the same players are asked to choose routes in the original and degraded networks. Moreover, it is obtained under a cost structure in which the equilibrium payoff doubles in value when the network

is so degraded. This evidence is to be contrasted with experimental studies of two other major paradoxes in non-cooperative game theory, namely, the two-person iterated Prisoner's Dilemma game and the two-player Centipede game (e.g., [McKelvey and Palfrey 1992](#)), which provide no support for equilibrium play. We attribute this difference in results to the fact that altruism, reciprocity, and punishment that play a major role in dyadic interaction have no effect on route choice with a relatively large number of players in each group as in our study.

The second contribution concerns the effects of type of monitoring. We report almost identical patterns of route choice in the original and degraded networks under public and private monitoring. These patterns appear in each of the separate sessions. Our analysis in Sect. 4.4 suggests that learning to play the one-shot equilibrium strategy may take place even under private monitoring. Taken together, these two contributions further demonstrate the robustness of the BP and its importance to network design.

Repetition of a game is known to have critical implications for the efficiency of the equilibria. The folk theorem states that the efficient outcome of a one-shot game can be supported as an equilibrium of the infinitely repeated game, and therefore, selfish players may learn to choose strategies that maximize social welfare. In classical game theory, the folk theorem requires public monitoring (e.g., [Friedman 1971](#)), but recent study ([Scarsini and Romania 2010](#)) shows that players may learn to play in a socially optimal way even under private monitoring. We do not observe patterns of route choice that increase efficiency. This may be attributed to the fact that the number of players in our experiment is too large, the number of rounds is too small, or both to allow for spread of information in the population of players. Future experiments with smaller number of network users or more rounds of play could examine whether efficient route choices could be observed with experience in repeatedly playing traffic network games susceptible to the BP.

Acknowledgements This research has been supported by grant 1008393 awarded to the University of Arizona by the National Science Foundation. We thank Filippo Rossi and Maya Rosenblatt for their help in data collection.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Aoki, K., Ohtsubo, Y., Rapoport, A., & Saijo, T. (2007). Effects of prior investment and personal responsibility in a simple network game. *Current Research in Social Psychology*, 13(2), 10–21.
- Arnott, R., de Palma, A., & Lindsey, R. C. (1993). A structural model of peak-period congestion—A traffic bottleneck with elastic demand. *American Economic Review*, 83(1), 161–179.
- Aumann, R. J. (1992). Irrationality in game theory. In P. Dasgupta, D. Gale, O. Hart, & E. Maskin (Eds.), *Economic analysis of markets and games* (pp. 214–227). Cambridge, MA: MIT Press.
- Aumann, R. J. (1995). Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 8(1), 6–19.
- Aumann, R. J. (1998). On the centipede game. *Games and Economic Behavior*, 23(1), 97–105.
- Borges, T., & Sarin, R. (2000). Naive reinforcement learning with endogenous aspirations. *International Economic Review*, 41(4), 921–950.
- Braess, D. (1968). Über ein Paradoxon der Verkehrsplanung. *Unternehmensforschung*, 12, 258–268.

- Camerer, C. F., & Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4), 827–874.
- Camerer, C. F., Ho, T., & Chong, J. (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119(3), 861–898.
- Cohen, J. E. (1988). The counterintuitive in conflict and cooperation. *American Scientist*, 76(6), 577–584.
- Cohen, J. E., & Kelly, F. P. (1990). A paradox of congestion in a queuing network. *Journal of Applied Probability*, 27(3), 730–734.
- Dafermos, S., & Nagurney, A. (1984). On some traffic equilibrium-theory paradoxes. *Transportation Research Part B—Methodological*, 18(2), 101–110.
- Fisk, C., & Pallottino, S. (1981). Empirical-evidence for equilibrium paradoxes with implications for optimal planning strategies. *Transportation Research Part A—Policy and Practice*, 15(3), 245–248.
- Frank, M. (1981). The Braess Paradox. *Mathematical Programming*, 20(3), 283–302.
- Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *Review of Economic Studies*, 38(113), 1–12.
- Kolata, G. (1990, December 25). What if they closed 42bd street and nobody Noticed? *The New York Times*.
- Koutsoupias, E., & Papadimitriou, C. (1999). Worst-case equilibria. In *STACS'99—16th annual symposium on Theoretical aspects of computer science* (Vol. 1563, pp. 404–413). Springer, Berlin.
- Luce, R. D., & Raiffa, H. (1957). *Games and decisions: Introduction and critical survey*. New York: Wiley.
- Mailath, G. J., & Samuelson, L. (2006). *Repeated games and reputations: Long-run relationships*. Oxford: Oxford University Press.
- McKelvey, R. D., & Palfrey, T. R. (1992). An Experimental-Study of the Centipede Game. *Econometrica*, 60(4), 803–836.
- Morgan, J., Orzen, H., & Sefton, M. (2009). Network architecture and traffic flows: Experiments on the Pigou–Knight–Downs and Braess Paradoxes. *Games and Economic Behavior*, 66(1), 348–372.
- Murchland, J. D. (1970). Braess's Paradox of traffic flow. *Transportation Research*, 4(4), 391–394.
- Pas, E. I., & Principio, S. L. (1997). Braess' Paradox: some new insights. *Transportation Research Part B: Methodological*, 31(3), 265–276.
- Rapoport, A., Kugler, T., Dugar, S., & Gisches, E. J. (2008). Braess Paradox in the laboratory: Experimental study of route choice in traffic networks with asymmetric costs. In T. Kugler, J. C. Smith, T. Connolly, & Y. J. Son (Eds.), *Decision modeling and behavior in complex and uncertain environments* (pp. 309–337). New York: Springer.
- Rapoport, A., Kugler, T., Dugar, S., & Gisches, E. J. (2009). Choice of routes in congested traffic networks: Experimental tests of the Braess Paradox. *Games and Economic Behavior*, 65(2), 538–571.
- Rapoport, A., Mak, V., & Zwick, R. (2006). Navigating congested networks with variable demand: experimental evidence. *Journal of Economic Psychology*, 27(5), 648–666.
- Roughgarden, T. (2005). *Selfish routing and the price of anarchy*. Cambridge, MA: MIT Press.
- Roughgarden, T. (2006). On the severity of Braess's Paradox: Designing networks for selfish users is hard. *Journal of Computer and System Science*, 72(5), 922–953.
- Roughgarden, T. (2007). Routing games. In N. Nisan, T. Roughgarden, E. Tardos, & V. V. Vazirani (Eds.), *Algorithmic game theory* (pp. 461–486). Cambridge, MA: Cambridge University Press.
- Roughgarden, T., & Tardos, E. (2002). How bad is selfish routing? *Journal of the ACM*, 49(2), 236–259.
- Scarsini, M., & Romania, V. (2010). Repeated congestion games with bounded rationality. Working Paper.
- Selten, R., Schreckenberg, M., Chmura, T., Pitz, T., Kube, S., Hafstein, S. F., Chrobok, R., et al. (2004). Experimental investigation of day-to-day route-choice behaviour and network simulations of autobahn traffic in North Rhine-Westphalia. In *Human behaviour and traffic networks* (pp. 1–21). Springer, Berlin.
- Smith, M. J. (1978). In a road network, increasing delay locally can reduce delay globally. *Transportation Research*, 12(6), 419–422.
- Steinberg, R., & Stone, R. E. (1988). The prevalence of paradoxes in transportation equilibrium problems. *Transportation Science*, 22(4), 231–241.
- Steinberg, R., & Zangwill, W. I. (1983). The prevalence of Braess' Paradox. *Transportation Science*, 17(3), 301–318.
- Taguchi, A. (1982). Braess' Paradox in a two-terminal transportation network. *Journal of the Operations Research Society of Japan*, 25(4), 376–388.

- Valiant, G., & Roughgarden, T. (2006). Braess's paradox in large random graphs. In *Proceedings of the 7th ACM conference on Electronic commerce* (pp. 296–305). Ann Arbor, MI: ACM.
- Youn, H., Gastner, M.T., & Jeong, H. (2008). Price of anarchy in transportation networks: Efficiency and optimality control. *Physical Review Letters*, *101*(12), 128701(1)–128701(4).