Route vs. Segment:
An Experiment on Real-Time Travel Information in Congestible Networks

Vincent Mak

University of Cambridge
Cambridge Judge Business School, Trumpington Street,
Cambridge CB2 1AG, United Kingdom
Email: v.mak@jbs.cam.ac.uk

Eyran J. Gisches

University of Arizona
1130 E. Helen St. Tucson AZ, 85716
Email: eyran@email.arizona.edu

Amnon Rapoport

University of California Riverside
900 University Ave. Riverside CA, 92521
Email: amnonr@ucr.edu

May 16, 2014

Acknowledgement: The authors gratefully acknowledge financial support from the National Science Foundation grant SES-0752662 awarded to the University of Arizona, as well as financial support from an IFREE Small Grant awarded to the University of California, Riverside. IRB approval has been obtained for the data collection described in this paper.

Route vs. Segment:

An Experiment on Real-Time Travel Information in Congestible Networks

Abstract

Research on the impact of travel information on route choice decisions has mostly focused on the effects of *exogenous* environmental uncertainties. However, it has been rarely investigated as to whether such information has significant impact on route choice under *endogenous* strategic interactions. We report the results of a computerized experiment that focuses on these issues. In our experiment, large groups of subjects independently made routing decisions in a congestible network that defies analysis by common users. Equilibrium predictions were tested in a between-subject design: in one condition, players committed themselves to a route at the origin; in the other, they made segment choices sequentially at each decision node upon receiving complete en route information. Despite the complexity of the network, traffic patterns in both conditions converged rapidly to equilibrium predictions. We account for the observed results by a Markov adaptive learning model postulating regret minimization and inertia. We find that subjects' learning behavior was similar across conditions, except that they exhibited more inertia with en route information. However, the additional route switching in the condition without en route information mutually cancelled out to render route choice patterns non-significantly different between conditions.

Keywords: transportation; behavioral operations; experiments; game theory

1. Introduction

The study of complex congestible networks with decentralized decision making, such as choice of routes in traffic networks, falls in the intersection of transportation research, operations management, computer science, and behavioral economics. Agents traversing such networks have to choose among multiple routes, each of which consisting of a sequence of segments (or links) connecting a common origin to a common destination. Each route segment is associated with a latency (or cost) function that serves as a proxy for the effects of delay in travel. The travel cost along each segment is characterized by negative externalities in that it is increasing in the total number of agents traversing that segment. Further complications arise from the fact that in some cases agents have to commit to a *route choice* from the origin to the destination at the beginning of their traversal, whereas in other cases they make *segment choice* decisions sequentially at each network junction predicated on their information about the current traffic conditions. The latter context has become ever more pertinent in recent years, as real-time travel information has become conveniently accessible to users of transportation networks through the development of infrastructure such as Advanced Traveler Information System (ATIS) or Advanced Traffic Management System (ATMS) (Chorus, Molin, and van Wee 2006a, Zhang and Levinson 2008).

Accompanying such advance in communications technology is a growing literature on route choice decisions with real-time travel information (see Chorus et al. 2006a for comprehensive reviews of earlier studies, and Ben-Elia et al. 2013a, Section 2, for an update). Research in this stream is often approached by individual decision making laboratory experimentation, in which route conditions would be presented as *exogenous* random variables, which can be used to capture environmental uncertainties such as accidental network disruptions due to bad weather,

so that travel information could help the individual subject resolve the uncertainties (see e.g., Chorus et al. 2006a, 2006b, Zhang and Levinson 2008, Gao, Frejinger, and Ben-Akiva, 2008, Ben-Elia and Shiftan 2010, Lu, Gao, and Ben-Elia 2011, and Ben-Elia et al. 2013a, 2013b). However, the interaction between travelers is largely ignored in this line of research with the exception of Lu et al. (2011), who studied how real-time information regarding exogenous uncertainties affects route choices in a simple congestible network.

The present study adopts the laboratory experimentation approach in research on real-time travel information. We depart from most related studies by focusing on real-time information regarding *endogenous* uncertainties that arise from strategic interactions in a game-theoretic non-cooperative context. As Ben-Elia et al. (2013, p.148) put it, this is a realistic scenario in which “each of the travelers choose a route and these result in endogenous volumes and travel times which are dependent on the complete set of choices.” It is obviously a major formative component of real-time travel information, and yet the influence (if any) of its provision on route choices is poorly understood.

A major purpose of the present study is to investigate whether real-time information about network congestion has a significant impact on route choice in complex networks. Specifically, our design compares subject decisions in two route choice conditions that differ in their information structure. In one condition, players are required to commit to a route choice at the origin; in the other, they make segment choices sequentially at each decision node upon receiving en route congestion information. We examine the *strategic* decisions that our laboratory network users make when they repeatedly and independently choose routes in a computerized network for payoff contingent on their performance. Using equilibrium analysis, we generate benchmark predictions of route choices and test them experimentally.

Relatedly, a second purpose of our experiment is to test whether seemingly impossible tacit coordination in achieving equilibrium traffic might or might not be attainable in the laboratory, and whether en route information provision is necessary or sufficient for such coordination. The theoretical analysis that informs our experimental predictions can be traced back to the pioneering work by Knight (1924), Vickrey (1969), and others. Well-known in that stream of results are traffic paradoxes such as the Pigou-Knight-Downs Paradox and Downs-Thomson Paradox (Downs 1962), the Bottleneck Paradox (Arnott, Palma, and Lindsey 1993), and the Braess Paradox (Braess 1968). One might indeed ask whether, in complex networks, economic theorizing based on strong assumptions about rationality and strategic sophistication hold water at all. While Rosenthal (1973) proved the existence of pure-strategy Nash equilibria in a large class of congestible network games, such games often exhibit an enormous multiplicity of equilibria, mostly involving asymmetric decisions among agents. Tacit coordination in achieving a single equilibrium in such networks seems utterly impossible. If the aggregate route choices of independent players over time could ever exhibit regularities that approximate an equilibrium outcome, then that would be like “magic”, as Kahneman (1988) remarked on similar observations in a considerably simpler market entry game class experiment.

An especially relevant analytical concept, which is also susceptible to tacit coordination problems, can be found in Wardrop (1952) (see also Beckmann, McGuire, and Winsten 1956; Arnott and Small 1994; Correa and Stier-Moses 2011). Wardrop’s widely adopted notion of traffic equilibrium is a steady state in which all users bear identical individual congestion costs that are not higher than the cost of switching to any unused route. Wardrop equilibrium can be seen as a special case of Nash equilibrium, and the two become equivalent when the number of drivers is large (Morgan, Orzen, and Sefton 2009). Wardrop’s concept is the starting point in

many subsequent studies, as reviewed in Correa and Stier-Moses (2011); see Cominetti, Correa, and Stier-Moses 2009, and Feldman and Tamir 2012, for two recent applications in operations. There is a growing experimental literature on route choice in simple traffic networks (e.g. Selten et al. 2007; Rapoport, Mak, and Zwick 2006; Morgan et al. 2009; Rapoport et al. 2009; Daniel, Gisches, and Rapoport 2009; Gisches and Rapoport 2012; Rapoport, Gisches, and Mak 2013; Rapoport et al. 2014) that offers strong support for the descriptive validity of equilibrium concepts for such networks. This is corroborated by many market entry game experiments with binary decisions that offer similar evidence (e.g. Rapoport, Seale, and Winter 2002, also cf. Camerer 2003, Chapter 7.3, Erev, Ert, and Roth 2010).

Our objective is to study network games with considerably more complex topology than in previous related studies in the laboratory, and determine whether the “magic” of large-scale coordination may still occur. Importantly, we examine if the provision of en route travel information might lead to large-scale coordination or, conversely, whether large-scale coordination can be achieved even without such information. Previous research suggests that travelers have low willingness to pay for real-time information (see Chorus et al. 2006b). It might be the case that, at least for commuters who collectively traverse the same congestible network, users could establish equilibrium coordination with or without en route travel information. Therefore, our experiment which suggests such evidence could help rationalizing the underutilization of observations in previous research.

1.1. Overview of the Experiment

To fix ideas, consider a situation where every day, after completing their shift, a fixed and commonly known number of commuting workers decide independently which route to choose from a common workplace (origin) to an Interstate highway (destination). To simulate this

situation experimentally, we construct a traffic network (Figure 1) with several routes from a common origin to one of several alternative entries to a common destination (terminal nodes). Each route consists of the same number of segments, many of which are shared between several routes. Importantly, all the segments are susceptible to congestion. We model this situation as a non-cooperative n -person game with complete information, which is iterated for a fixed number of T rounds in the session. Because our focus is on strategic uncertainty, we have deliberately left out the modelling of environmental uncertainties arising from accidental network disruptions (e.g., accidents, road construction, bad weather).

– Insert Figure 1 about here –

Specifically, we examine the decisions players make under two information regimes, using equilibrium predictions as benchmarks. In the route choice condition (Condition RCC), only ex-post information is provided at the end of each round t ($t=1, 2, \dots, T$) about the arrival distribution of the players over the available routes connecting origin to terminal nodes, and the associated costs of travel. This can be likened to daily traffic information that is commonly available in the media or online. Given this information, which is commonly known, each player in Condition RCC commits herself to a single route on round $t+1$. In the segment choice condition (Condition SCC), players receive en route information about the arrival distribution of the n players across the intermediate decision nodes connecting the segments (hence, the requirement that all routes have the same number of segments). This can be likened to the provision of real-time travel information via an ATIS, for example, which might provide users with comprehensive traffic conditions in road networks. Each player in Condition SCC only commits herself at each intermediate node to a single road segment. Condition RCC has the characteristics of a simultaneous game, while Condition SCC, with a richer information structure,

is of a sequential nature. Thus, our experiment encompasses a test of equilibrium concepts on two major distinctions of games in the same congestible network. Our design specifically isolates the influence of en route travel information on route choices: both conditions provide ex-post travel information at the end of every round, but only Condition SCC supplies this information *en route* during the round.

The rest of the paper is organized as follows. Section 2 introduces terminology and describes the network presented to the subjects in our experiment, and Section 3 presents the equilibrium solutions for both Conditions RCC and SCC. Section 4 presents the design and procedures of the experiment, and describes our major findings. A major question concerns the dynamics of play over iterations of the stage game. We account for the observed results by a Markov adaptive learning model postulating regret minimization and inertia. Section 5 concludes with a discussion of the significance and generality of our findings.

2. The Network

The network in our study (Figure 1) is depicted as a directed graph $G = (V, E)$ with V being a set of vertices (also called *nodes*) and E a set of directed edges (also called *links* or *segments*), together with a set of routes that share a common origin and common destination. Specifically, the network contains 10 nodes and 12 segments. In correspondence with the terminology in the subject instructions (cf. Appendix B), the nodes are labeled by the Roman letters A, B, ..., J. The four nodes in the set $\Delta = \{D, G, I, J\}$ are *terminal* (destination) nodes, all leading to the same destination, and the remaining six nodes $\{A, B, C, E, F, H\}$ are *decision* nodes. (This is, in fact, a single origin-destination network as nodes D, G, I, and J may be connected to a single terminal node with zero-cost edges.) We consider the case of $n=18$ players who are required to choose independently and anonymously a route from the origin A to any of the nodes in Δ with the

incentivized objective of minimizing cost of travel. A segment from node i to node j , denoted as ij , is assigned a cost function $c_{ij}(f_{ij})$, also called a “latency function,” which specifies the cost of travel incurred by a player traversing that segment as a function of the number of players, f_{ij} ($1 \leq f_{ij} \leq n$) who choose the same segment. Mostly for experimental reasons, the cost functions all have the linear (affine) form of $c_{ij}(f_{ij}) = a \times f_{ij} + b$, $a > 0$, $b \geq 0$. For example, if four commuters traverse segment AB, then each of them incurs a travel cost of $7 \times 4 + 2 = 30$. The cost functions are separable over segments; that is, the cost of traversing segment ij only depends on the number of players traversing it, and not on the number of players traversing any other segment.

– Insert Table 1 about here –

Altogether there are eight alternative routes from the origin to the destination set; they are denoted by ABCD, ABCG, ABFG, ABFI, AEFG, AEFI, AEHI, and AEHJ (e.g., ABFG denotes the route which passes through nodes A, B, and F, to arrive at destination G), and are listed in the first column of Table 1. Each route includes three segments (e.g., route ABFG includes the segments AB, BF, and FG). We refer to each choice of segment as a *move*. In Condition RCC, each of the n players commits to a single route (a sequence of connected segments) at the beginning of the game, whereas in Condition SCC each player only commits to a single move. Thus, upon reaching decision node j , a player in Condition SCC has to make a binary choice between the two segments emanating from node j , whereas a player in Condition RCC has to choose one of eight routes.

In every session, the stage game was repeated 50 times, each repetition being referred to as a *round*. After each round in Condition RCC, players were accurately informed about the distribution of route choices (i.e., number of players choosing each route) in that round and the associated costs of travel. Within each round in Condition SCC, after all the players chose a

segment on the k -th move ($k=1, 2, 3$) they also were informed about the arrival distribution (i.e., how many players had reached which node) up to that move.

In sharp contrast with experiments on route choice in the transportation literature, but in conformity with reality, pairs of routes in Figure 1 share 0, 1, or 2 segments. For example, routes ABCD and AEHJ share no common segments, routes ABCD and ABFJ share a single segment, and routes ABCD and ABCG share two segments. Because of the multiplicity of routes and high degree of segment sharing, players in both conditions face a formidably difficult task.

3. Equilibrium Predictions

In Appendix A, we outline the procedures used to compute the equilibrium traffic or route choice distribution. It turns out that both conditions have the same equilibrium solution, which is presented in Table 1: the eight routes are listed in the first column and the route choice distribution in the second column, while the third column lists the associated individual travel costs. Column 2 in Table 2 presents the equilibrium traffic flow (i.e., how many players traverse which segment); it can be shown that this traffic flow is consistent with only one, unique distribution of route choices among players, which is the same as the distribution in column 2 in Table 1 (see Appendix A for more details).

– Insert Table 2 about here –

In equilibrium, only six of the eight routes are chosen by at least a single player. The equilibrium traffic has Wardrop equilibrium characteristics in that the congestion cost per player in equilibrium is identically 100 (column 3 in Table 1), and it is less than the cost of choosing any of the unused routes. There are over 3.3 billion ($\approx 18!/(3! \times 1! \times 2! \times 2! \times 8! \times 2!)$) pure-strategy equilibria, each a distinct mapping from the set of players to the set of routes, that result in the same route choice distribution. All of these equilibria have the same total cost, i.e., the sum of

travel costs over all the players. Therefore, the equilibria are all Pareto unrankable. The equilibria only differ from one another in the identities of the individual players choosing each route, and require high asymmetric coordination to achieve. This level of multiplicity suggests that the probability that a group of players will stumble upon any particular equilibrium is negligible. In fact, even if the equilibrium route choice distribution is disclosed to the players at the beginning of the experiment, the coordination problem is essentially unsolvable if there is no central authority to coordinate route choices. One may readily reach the conclusion that the equilibrium solutions in the network under study are devoid of any descriptive power. However, as shown below, the equilibrium solutions are in fact highly predictive of the observed aggregate traffic patterns in the experiment.

The equilibrium traffic is inefficient, which is typical of decentralized congestible networks (Mak and Rapoport 2013). In terms of route choice distribution, there are two efficient solutions that minimize the total cost of travel across the 18 players (columns 4 to 7 in Table 1), both resulting in the same traffic flow (column 3 in Table 2).

4. The Experiment

4.1. Method

4.1.1. Subjects. One hundred and eighty undergraduate and graduate students at the University of Arizona volunteered to participate in a computer-controlled experiment on decision making for payoff contingent on performance. Male and female students participated in almost equal numbers. The subjects were divided into ten groups each participating in a single session. Five groups of 18 subjects were assigned to Condition RCC and five others to Condition SCC. Payments were cumulative and stated in tokens that at the end of the session were converted to

US dollars at the exchange rate of US\$1=100 tokens. Each session lasted about two hours. The mean payment across all 10 sessions was US\$20.9 including a show-up fee of US\$5.

4.1.2. Procedure. The experimental sessions were conducted in a computerized laboratory with multiple terminals located in separate cubicles. Subjects were handed the instructions that they read at their own pace (see Appendix B for the instructions for Condition RCC). Questions about the procedures were answered individually by the experimenter.

The instructions described the network game and the procedures for choosing among the routes. They exhibited the network in Figure 1, explained the cost functions, and illustrated the computation of the travel cost with several detailed examples. In Condition RCC, a round was concluded after all 18 players in the group typed in their route choices; in Condition SCC, the players received feedback after every move, and a round was concluded after all players typed in their third segment choice. In both conditions, after a round was concluded, a new screen was displayed with full information presented in two tabular forms that showed the number of players choosing each route as well as each segment, and the associated costs of travel.

Each subject was provided with a reward (endowment) of $R=140$ tokens in each round for completing her travel. Importantly, individual payoff for the round was calculated by *subtracting* the travel cost on the round from the fixed reward. The value of R was chosen so that, in equilibrium, a player in Conditions RCC and SCC earns $140-100=40$ tokens. Examination of the payoff function over different profiles of choices shows that the function is anything but flat, so that deviations from equilibrium play result in considerable changes in individual payoffs. With the value of the reward set at $R=140$, on any particular round players could either gain or lose money. For example, if at the end of a round in any information condition, 10 subjects choose route ABCG, 6 choose route AEFG, and 2 choose route AEHJ, then the respective earnings in

that round for members of each subgroup are -71, -12, and 68, respectively. All the 180 subjects ended the session with positive payoffs.

4.2. Results

4.2.1. Preliminary Analysis. Table 1 (columns 8 and 9) and Table 2 (columns 4 and 5), respectively, present the mean choice frequencies of routes and segments, i.e., the mean observed number of players (averaged over sessions and rounds) traversing each route and each segment in each condition. Figure 2 exhibits the mean choice frequency of each of the eight routes in every round for each condition separately. Both numerical and graphical presentations indicate that the observed route and segment choice frequencies approached theoretical predictions from early on in the experimental session in both conditions.

– Insert Figure 2 about here –

It is problematic in our case to statistically test for convergence to equilibrium. This is because the expected equilibrium frequencies of two of the routes (ABFI and AEFI) and, by implication, one of the segments (FI), are zero, so that noisy deviations in choice frequencies from equilibrium for those routes/segments could not be a zero mean random variable. Here, we report some indirect statistical evidence. First, we conduct chi-squared tests for each group in each condition using observed mean route choice frequencies of the group but leaving out the two routes with zero expected frequencies; the tests could not reject the hypothesis that the frequency distributions equal the equilibrium predictions ($p > 0.9$ in all tests). Second, chi-square tests comparing the mean choice distributions of all the routes in the two conditions (with the group as the unit of observation) could not reject the hypothesis that the distributions are the same ($p > 0.7$), supporting the claim that subject behavior approached the same pattern with or without en route information.

As another check of equilibrium predictions, we calculate the coefficient of variation (CV) of congestion costs among players (defined as the standard deviation of costs divided by the average cost) for every group in every round. In general, the CV did not evolve to a large extent over rounds, but there was a gradual drift towards lower values (see Figure A1 in Appendix C); aggregating over sessions and rounds, its mean is 0.17 for Condition RCC and 0.15 for Condition SCC, versus the equilibrium predictions of zero in both conditions. We conclude that, overall, players bore similar congestion costs in both conditions in good agreement with Wardrop equilibrium predictions.

The similarity in traffic patterns in both conditions in their approximation to equilibrium predictions might suggest that en route information was not utilized by the subjects in Condition SCC. However, additional analysis shows otherwise. To proceed, for the data from Condition SCC, we conduct logistic regressions on how a segment choice was influenced by the realized traffic up to that point. For example, conditioned on a subject having just arrived at B, we conduct a logistic regression in which the dependent variable is whether the subject next chose BC or BF, and the independent variable is the number of players who had also arrived at B (the remaining subjects must have arrived at E). We then conduct a similar regression for observations of subjects who had just arrived at E. To account for possible correlations among decisions by the same player over rounds, as well as among players in the same group, we obtain our estimations using the generalized estimating equations (GEE) approach (cf. Hardin and Hilbe 2003). Both estimations yielded significant dependence between the choice of the next segment and the number of players who had also arrived at the same node ($p < 0.01$ in both cases). Similar logistic regressions for segment choices further down the network have also yielded significant dependences on realized traffic.

Our analysis shows that the en route or “real-time” travel information provided in Condition SCC did have an impact on decisions. This is in line with intuition and also previous research (e.g., Ben-Elia et al. 2013a). However, aggregate traffic pattern could converge sufficiently quickly even in Condition RCC – which provided only ex-post but not en route travel information – that the additional information in Condition SCC did not have a significant impact on overall traffic in the experimental setting. *Even though subjects’ decisions could be influenced by the en route information they received, observed aggregate traffic in both conditions rapidly approached equilibrium play.* Hence, in the following sections, we shall focus our analysis on round-to-round route choice.

4.2.2. Route Switching Over Successive Rounds. On average, subjects switched their routes between successive rounds 28.5 times and 24.4 times within a 50-round session in Conditions RCC and SCC, respectively. Two-sample *t*-test (with group as the unit of observation) shows that this difference between conditions is significant ($p < 0.01$); subjects in Condition SCC tended to be *relatively* more stable possibly because of receiving more information updates about the other players’ decisions within every round (see also relevant discussion in Section 4.2.5).

– Insert Figure 3 about here –

Switching frequencies decreased with experience: as shown in Figure 3, the percentage of route-switching players within every condition decreased with round reaching around 40% by the end of the session. The mean number of switches per player dropped from 16.0 in the first half (round 1-25) to 12.5 in the second half (round 26-50) of the session in Condition RCC (a switch between routes from round 25 to 26 is counted as a switch within the *second* half of the session); the corresponding drop for Condition SCC was 14.3 to 10.1. Paired *t*-tests (with group as the unit of observation) that compared mean number of switches in the two halves yielded

significant results ($p < 0.01$ for both conditions). This can partly be explained by the subjects' realization that switching did not increase payoff. In fact, statistical analysis shows that switching had a negative impact on payoffs. To proceed, we regress every player's payoff in every round after round 1 on a dummy variable that is equal to 1 if the player had switched route from the previous round, and 0 otherwise. To account for correlations among decisions by the same player over rounds, as well as among players in the same group, we use a GEE model in the linear regression estimations. As it turns out, the estimations yield a significant negative dependence ($p < 0.01$) between payoff and switching frequency in either of the two conditions.

– Insert Figure 4 about here –

Switching patterns were highly heterogeneous among subjects. Figure 4, which is a histogram of the number of switches among subjects, exhibits smooth, unimodal distributions for both conditions with no distinct segments. Further examination of individual data confirms this heterogeneity, as we can identify subjects who had a preferred route that was chosen more than half of the time since early in the session, as well as subjects who switched frequently between several routes throughout the session.

4.2.3. Inefficiencies. Even though observed traffic approached equilibrium rapidly, switching routes between rounds, which is critical for adaptation, necessarily came at some costs to the subjects. Since the sum of travel costs borne by the users of any particular segment is quadratic in that segment's traffic flow, small deviations in traffic flows from equilibrium tend to be magnified at the level of total cost over all players. Consequently, subjects performed worse than in equilibrium in terms of total costs. On average, the total costs were 2037.8 per round in Condition RCC and 2001.2 in Condition SCC. These values exceed the equilibrium level (which is 1800) by 13.4% in Condition RCC and 11.3% in Condition SCC. The extent of inefficiency

(measured as percentage deviation of total cost from the social optimum of 1774) is 14.7% in Condition RCC and 12.6% in Condition SCC; these values exceed the corresponding 1.5% with equilibrium traffic.

4.2.4. Learning Model: Setup. Despite persistent switching of route choices by most of the subjects, Figure 2 suggests that aggregate behavior stabilized around the equilibrium route choice distribution rapidly in both conditions. We propose a simple learning mechanism to account for this rapid convergence. Our purpose is *not* to competitively test learning models (among the wide range of possibilities that are available, cf. Camerer 2003, Chapter 6, and Erev et al. 2010) that may best account for data, but rather propose a single model that provides insights into what behavioral factors might have impacted convergence. We model subjects in either condition *as if* they committed to their route choices at the beginning of every round. Our rationale for doing so, even for Condition SCC, is that the observed aggregate traffic in both conditions rapidly approached equilibrium play. Thus, it seems that, from the perspectives of understanding rapid equilibrium convergence, and for the sake of parsimony, we consider it unnecessary to consider within-round responses to en route information in Condition SCC.

– Insert Table 3 about here –

Previously, rapid convergence to equilibrium outcomes in market entry games has been successfully accounted for by a modified form of the reinforcement learning model (Erev and Roth 1998), which can be seen as belonging to the larger family of experienced weighted learning models (Ho, Camerer, and Chong 2007). The modification typically includes some dependence of a player's current choice on what other players have done in the past (e.g., Rapoport et al. 2002). Recent developments (cf. Erev et al. 2010) are generally consistent with this view. We construct our learning model along a similar line, but also in consideration of

possible differences due to the larger strategy set in our setting compared with the binary strategy sets in market entry games. The larger strategy set suggests that the average player would pay much more attention to the utilization of every route (whether it was her chosen route or not) from round to round, compared with the simpler market entry games. As a result, the dependence on what other players have done, which is already useful in accounting for learning in market entry games, should take on a bigger role in our case.

On the other hand, assuming that subjects used simple rules or heuristics (which should be the case with such rapid convergence) to determine their choices, the increased attention to other players' choices should be traded off by a decreased attention to the decision maker's own history of choices and payoffs, compared with previous reinforcement learning-based models. Thus, we assume that a decision maker's current choice in our study was explicitly dependent on what that decision maker and other players did in the *previous round only*. Indeed, Erev et al. (2010)'s large-scale market entry game experiments also showed that subjects were highly influenced by what happened in the most recent round. This recency effect could only be more pronounced in our complex congestible network game in which the outcomes of each round would require much of the player's attention to digest. To sum up, we assume that subjects reconsidered their choices from round to round according to a Markov adaptive learning model.

The most natural specification of such a model would be one in which the attraction of a route in round $t+1$ increases with its (actual or counterfactual) payoff in round t , had the player chosen that route in round t , given the route choices of all other players in that round. Thus, we begin with a simple "baseline" specification: suppose the players are at the beginning of round $t+1$, $t = 1, 2, 3 \dots 49$. Let $C_{a,j,t}$ be the travel cost that player a would have incurred had she chosen route

j in round t , given the choices of all other players in that round. Then, define the “attraction” of (or “propensity” towards choosing) this route as:

$$A_{a,j,t}^{baseline} = -\lambda C_{a,j,t},$$

where λ is a response sensitivity parameter to be estimated and is expected to be positive. The probability of player a choosing route j in round $t + 1$ is given by the multinomial logit function:

$$\Pr^{baseline}(R_{a,t+1} = j) = \exp(A_{a,j,t}^{baseline}) / \sum_k \exp(A_{a,k,t}^{baseline}),$$

where $R_{a,t+1}$ is the route choice of player a in round $t+1$ and the summation in the denominator is over all eight routes.

Our model may be interpreted as a logistically smoothed version of Cournot dynamics (Camerer 2003, Chapter 6.3); this is because as λ (which can alternatively be interpreted as an inverse noise parameter as well as a response sensitivity parameter) tends to infinity, decision-making in the baseline learning model becomes best responding to other players’ choices in the previous round. The baseline model is also conducive to convergence to equilibrium choices in our network. This is because routes with low travel costs in a round hold stronger attraction to players compared with routes with high travel costs in the same round. As a result, routes with low costs in round t would see a larger increase in the number of users from round t to round $t+1$, compared with routes with high costs in round t , which might even suffer a decrease in the number of users in the next round. Since a route’s travel cost increases in the number of users, the net effect would be an equalizing of travel costs among routes, which would bring aggregate traffic towards the equilibrium prediction.

The model that we propose to estimate the data is built on the baseline model. In the estimated model, we allow for further stabilizing effects due to reference dependence in players’ choices. Specifically, we postulate that a decision maker’s actual choice and payoff in round t would

serve as reference points for her decision making in round $t+1$. Formally, in addition to the previous notations, denote player a 's actual route choice in round t as i , so that $C_{a,i,t}$ is the congestion cost that a incurred in round t . We propose the following model to be estimated:

$$A_{a,j,t} = \begin{cases} (\lambda + \lambda_+)(C_{a,i,t} - C_{a,j,t}) & \text{if } j \neq i \text{ and } C_{a,i,t} \geq C_{a,j,t}, \\ \lambda_0 & \text{if } j = i, \\ \lambda(C_{a,i,t} - C_{a,j,t}) & \text{if } j \neq i \text{ and } C_{a,i,t} < C_{a,j,t}, \end{cases}$$

where λ , λ_0 , and λ_+ are parameters to be estimated, and are all expected to be non-negative.

Similar to the baseline model, the probability of player a choosing route j in round $t + 1$ is the multinomial logit function:

$$\Pr(R_{a,t+1} = j) = \exp(A_{a,j,t}) / \sum_k \exp(A_{a,k,t}).$$

The interpretation of the parameters, when they are *positive*, is as follow: λ is the baseline response sensitivity of the player to the actual and counterfactual payoffs with no reference dependence. The other two parameters represent two postulated forms of reference dependence:

- (1) The parameter λ_+ represents a *regret effect* due to the player's inclination to minimize regret – which we assume to be forgone payoffs by proxy – so that routes that could have led to higher payoffs in round t than the actual choice hold extra attraction compared with otherwise. If, as we generally postulate, realized play in the previous round was effective proxy representation of players' anticipation of what would happen in the present round, then the parameter reflects the impact of players' anticipated or expected regret on their route choice. In this sense, our approach would be consistent with the classic treatment of regret by Loomes and Sugden (1982), who proposed their theory as an alternative to prospect theory (Kahneman and Tversky 1979, 1991). Also consistently, research in network route choices

with exogenous uncertainties has found that anticipated regret plays an important role in those choices (e.g., Chorus et al. 2006a, Ben-Elia et al. 2013b).

- (2) The parameter λ_0 represents *inertia* with respect to the actual route choice in round t , over and above payoff considerations. If it turns out that $\lambda_+ = \lambda_0 = 0$ and $\lambda > 0$, then the estimated model is effectively identical to the baseline model, since the logit choice function is invariant up to a constant shift ($\lambda C_{a,i,t}$ in this case) in all the attraction terms. Thus, if our estimations yield λ_+ and λ_0 that are non-significantly different from zero, then the learning process could be accounted for by the baseline model without the postulated reference-dependent effects.

Significant sensitivity to forgone payoffs as well as inertia have both been observed as behavioral regularities in previous market entry game experiments (e.g., Erev et al. 2010). Thus, the postulated reference effects could possibly be exhibited by our dataset too. These effects have been conceptualized also to accelerate and stabilize convergence to equilibrium in our case. With the regret effect, low-cost routes would attract switches towards them at the expense of high-cost routes to a larger extent than only under the baseline model (controlling for λ); this would then further speed up the equalization of travel costs among routes. With inertia, once players converge towards an approximation of the equilibrium choice distribution, they would have an extra tendency to stick with it compared with when there is no inertia. An especially rapid and stable convergence would occur if baseline learning and regret effect are prominent from early on, while inertia increases in relative magnitude later in the session. This is feasible in the present model since the first two effects are large when there are considerable cost differences among routes, which happen when players have not yet converged near equilibrium traffic; however, after players converge near equilibrium traffic and cost differences among

routes become small, the impact of inertia will become relatively prominent and can serve to stabilize convergence effectively. We shall examine such possibility in subsequent analysis.

Even with homogeneous learning parameters among players (but also cf. footnote 2), the present learning model can give rise to highly heterogeneous switching behavior that helps us to understand the switching patterns discussed earlier. This is because: (1) ours is a stochastic choice model with probabilistic choices, so that there is an element of randomness in the choices which would result in heterogeneous behavior; (2) the learning model allows different players to experience different attractions for the same route in different rounds. In particular, the attraction of a route to a specific player is a function of whether it was chosen on the last round and also how its payoff compared with those of other routes had the player chosen those routes. This mechanism could lead to very different choice probability distributions and switching probabilities for different players even within the same group.

The formulation of our model, where the attraction of each alternative depends on a sensitivity parameter that interacts with forgone payoff, and the choice probabilities have a multinomial logistic dependence on the attractions, is common in the experimental literature on learning in games (see e.g., Camerer 2003, Chapter 6, and references therein). In the spirit of that literature, in establishing the present learning model our objectives are exactly to account for the impact of (session- and individual-specific) strategic interactions and history of play on choice by the model itself through the attractions. Thus, we find it in principle legitimate to assume that the remaining noise components are independent across data points. As such, we proceed to use standard multinomial logit estimation techniques in our analysis.

4.2.5. Learning Model: Results. We first use maximum likelihood to estimate our model for the two conditions separately with data from all 50 rounds, followed by an estimation of pooled

data from both conditions.¹ The results are displayed in Table 3. $\text{Log } L$ is the natural logarithm of the likelihood function of the model i.e.,

$$\text{Log } L = \sum_a \sum_{t=1}^{49} \text{Log Pr}(R_{a,t+1} = j_{a,t+1} \mid \lambda, \lambda_0, \lambda_+, H_{a,t}),$$

where $j_{a,t+1}$ is the observed route choice of player a in round $t+1$, $H_{a,t}$ is a 13-dimension vector consisting of the route choice of a in round t and the number of players traversing each of the 12 segments in round t , while summation over a means summation over all players in the relevant condition (RCC or SCC).²

Table 3 shows that the data exhibit significant presence of regret- and inertia-driven motivations in route choices. The regret effect (captured by λ_+) is at least 50% of that of the baseline response sensitivity (captured by λ) in both conditions and the pooled data, and could notably add to the speed of convergence to equilibrium. Note that, as discussed before, although the inertia coefficient, λ_0 , looks much larger than the other two coefficients, it is not weighted by any cost difference terms in the learning model, as the other two coefficients are. The implication is that: (1) at the beginning of the session, when there were considerable cost differences between the routes, convergence would be less hindered by inertia compared with later rounds; (2) as the session proceeded and the cost differences between the routes became small, inertia then became a strong stabilizing force.

¹ Note that we have not found any significant group learning effect: chi-square tests comparing mean route choice frequency distributions in the first vs. second half of the session in each condition revealed no significant differences ($p > 0.9$ in both tests). Hence, we pool data within each condition in the estimations. Further analysis confirms that our qualitative conclusions remain the same if the estimations were carried out separately for each group.

² We have also conducted a mixed logit analysis (Train 2003, Chapter 6; see Ben-Elia et al. 2013b for a related example) to investigate the heterogeneity of learning parameters among subjects. In the model, we assume that the learning parameters are normally distributed among the subjects; for the estimation that pools all data, we find relatively significant heterogeneity with the inertia parameter λ_0 (mean = 1.70, s.d. = 1.42), some heterogeneity with the regret parameter λ_+ (mean = 0.0094, s.d. = 0.0017), and very little heterogeneity with the response sensitivity parameter λ (mean = 0.014, s.d. = 0.001). As such, all estimations have approximately the same population mean as reported in the main text. Estimations with the two conditions separately yield similar population means as reported in the main text, and the same qualitative conclusions about heterogeneity as the pooled estimates. Hence, for simplicity, we focus on reporting the analysis of a model assuming homogeneous learning parameters.

Overall, the magnitudes of the estimates appear to be similar across the separate conditions as well as the pooled data. Further likelihood ratio tests reveal that λ_0 is significantly larger in Condition SCC than in Condition RCC ($\chi^2(1) = 31.8, p < 0.01$), but not the other two estimates ($\chi^2(1) = 3.32, p > 0.05$, for λ ; and $\chi^2(1) = 2.64, p > 0.1$, for λ_+).³ Further analysis with each half of the session shows that subjects exhibited more inertia in Condition SCC in both halves. Thus, it seems that the learning processes were qualitatively similar across the two conditions, but subjects had stronger inertia with their choices in Condition SCC than in Condition RCC. The difference in inertia is, indeed, in line with the switching analysis reported earlier. On the other hand, our preliminary analysis also shows that subjects' choices seemed to be sensitive to en route information in Condition SCC. It thus appears that the en route information in Condition SCC made it easier for subjects to approach segment-by-segment coordination, which helped each player adopt a "habitual" route approximating a pure-strategy equilibrium. However, the additional route switching in Condition RCC mutually cancelled out to render traffic non-significantly different across conditions.

Lastly, we examine how players' learning behavior evolved over the course of the session by estimating how the learning model parameters in each condition might have changed from the first half (round 1-25) to the second half (round 26-50) of the session. The procedures are similar to the comparison of the parameters across conditions, but instead of defining the dummy variable by whether the observation comes from Condition RCC or Condition SCC, we define the dummy variable by whether the observation comes from the first or the second half of a session. In both Conditions RCC and SCC, relevant likelihood ratio tests show that λ_0 increased

³ Formally, we first define a dummy variable d_C that is equal to 1 when an observation in the likelihood function comes from Condition SCC, and to 0, otherwise. We then replace λ in the model by $\lambda + \lambda_C \cdot d_C$ and similarly with the other two parameters, where the new coefficients such as λ_C are estimated together with the original coefficients. We then test hypotheses such as $\lambda_C = 0$ using likelihood ratio tests for each parameter separately.

significantly ($p < 0.01$) from the first to the second half of the session, suggesting that with more experience in playing the game subjects developed more inertia with their choices. No other statistically significant differences are observed with the parameters ($p > 0.1$ in all other cases), in particular with the regret effect λ_+ . Our observations are also consistent with the switching analysis reported earlier, where we found that switching became less frequent as the session progressed. We conclude that the rapid convergence was fueled from early on by the baseline learning and the regret effect, and then stabilized by an increase in inertia in later rounds.

5. Conclusions

Regardless of whether or not subjects were provided with en route information as they traversed the network, traffic in our experiment rapidly approached the equilibrium pattern. We find that subjects' learning behavior was similar across conditions, except that they exhibited more inertia in Condition SCC. We also find that subjects were sensitive to en route information in that condition. It appears that the en route information in Condition SCC made it easier for subjects to approach segment-by-segment coordination, which helped each player adopt a "habitual" route approximating a pure-strategy equilibrium. However, the additional route switching in Condition RCC mutually cancelled out to render traffic non-significantly different across conditions. We conclude that the differences in information provision and inertia did not translate into significant differences in the rapid convergence towards equilibrium in both conditions.

The lack of significant differences between conditions is pertinent to the design of systems such as ATIS and ATMS, which are introduced with the aim of reducing road users' uncertainties about route conditions. Our observations suggest that, with or without en route travel information, traffic might attain similar pattern, if (1) travelers have experience with each

other, as in a commuting context, and are informed about ex-post traffic (in the form of, for example, daily traffic information that is commonly available in the media or online); and (2) endogenous strategic interactions, rather than exogenous uncertainties, create the main uncertainties in payoffs. Previous research suggests that travelers have very low willingness to pay for real-time information (Chorus et al. 2006b). This is consistent with our experimental results, which suggest that, at least for commuters who collectively traverse the same congestible network repeatedly, users could establish equilibrium traffic with or without en route travel information. Our experiment provides a clue to understanding the underutilization of observations in previous research.

Our experimental results provide evidence in support of equilibrium concepts in congestible network analysis, and help strengthening the empirical validity of a central tenet in this important field of research. Ours is the first experimental test of network traffic equilibrium employing a non-trivial network that is complex beyond analysis to ordinary users. As a rule, real-life networks are more complex and involve larger groups of users than in our setup. But from the perspective of individual users, our experimental setting might only be *marginally* different from real-life networks in terms of the extent to which the decision context is complex beyond analysis. Nevertheless, additional experimental and field studies need to be performed with larger groups, different networks, and different information structures, in order to further examine the predictive validity of theoretical analysis of congestible network traffic.

Another avenue for further research is the introduction of exogenous uncertainties about delays or occasional incoming flow from other parts of the networks. As explained earlier, because our focus is on strategic interaction among network users, we have deliberately refrained from incorporating environmental uncertainties in our study. Future experimentation could relax

these limitations, which could further increase the difficulty of traffic approaching relevant equilibrium predictions. An interesting possibility is the interaction of environmental uncertainties and strategic interactions; for example, bad weather or other accidental disruptions might make the congestion externalities of each user for other users higher than otherwise (see, e.g., Rapoport et al. 2014). This can be translated as scenarios in which the congestion cost *structure* (i.e., cost functions) of specific segments is uncertain to travelers, which has been receiving research attention in recent years (see Unnikrishnan and Waller 2009, and a simple version of it in Lu et al. 2011). As proposed by Wijayaratna (2014), experimentation can indeed help us understanding the behavioral implications of this type of uncertainties.

References

- Arnott, R., A. de Palma, C. R. Lindsey. 1993. Properties of dynamic traffic equilibrium involving bottlenecks, including a paradox and metering. *Transportation Science* **27** 148-60.
- Arnott, R., K. A. Small. 1994. The economics of traffic congestion. *American Scientist* **82** 446-55.
- Beckmann, M. J., C. G. McGuire, C. B. Winsten. 1956. *Studies in the Economics of Transportation*. New Haven, CN: Yale Univ. Press.
- Ben-Elia, E., R. Di Pace, G. N. Bifulco, Y. Shiftan. 2013a. The impact of travel information's accuracy on route-choice. *Transportation Research Part C* **26** 146-159.
- Ben-Elia, E., R. Ishaq, Y. Shiftan. 2013b. "If only I had taken the other road ...": Regret, risk and reinforced learning in informed route-choice. *Transportation* **40** 269-293.
- Ben-Elia, E., Y. Shiftan. 2010. Which road do I take? A learning-based model of route-choice behavior with real-time information. *Transportation Research Part A* **44** 249-264.
- Braess, D. 1968. Über ein Paradoxon der Verkehrsplanung. *Unternehmensforschung*, **12** 258-68.
- Camerer, C. F. 2003. *Behavioral Game Theory*. Princeton, NJ: Princeton University Press.
- Chorus, C. G., T. A. Arentze, E. J. E. Molin, H. J. P. Timmermans, B. Van Wee. 2006a. The value of travel information: Decision strategy-specific conceptualizations and numerical examples. *Transportation Research Part B* **40** 504-519.
- Chorus, C. G., E. J. E. Molin, B. Van Wee. 2006b. Use and effects of advanced traveller information services (ATIS): A review of the literature. *Transport Reviews* **26** 127-49.
- Cominetti, R., J. R. Correa, N. E. Stier-Moses. 2009. The impact of oligopolistic competition in networks. *Operations Research* **57** 1421-1437.

- Correa, J. R., N. E. Stier-Moses. 2011. Wardrop equilibria. In *Wiley Encyclopedia in Operations Research and Management Science*, J. J. Cochran, ed. John Wiley & Sons.
- Daniel, T. E., E. J. Gisches, A. Rapoport. 2009. Departure time in Y-shaped traffic networks with multiple bottlenecks. *American Economic Review* **99** 2149-76.
- Downs, A. 1962. The law of peak-hour expressway congestion. *Traffic Quarterly* **16** 393-409.
- Erev, I., A. E. Roth. 1988. Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* **88** 848-81.
- Erev, I., E. Ert, A. E. Roth. 2010. A choice prediction competition for market entry games: An introduction. *Games* **1** 117-136.
- Feldman, M., T. Tamir. 2012. Conflicting congestion effects in resource allocation games. *Operations Research* **60**(3) 529-540.
- Gao, S., E. Frejinger, M. Ben-Akiva. 2008. Adaptive route choice in stochastic time-dependent networks. *Transportation Research Record* **2085** 136-43.
- Gisches, E. J., A. Rapoport. 2012. Degrading network capacity may improve performance: private versus public monitoring in the Braess paradox. *Theory and Decision* **73** 267-93.
- Hardin, J. W., J. M. Hilbe. 2003. *Generalized Estimating Equations*. Boca Raton, FL: Chapman & Hall/CRC.
- Ho, T. H., C. F. Camerer, J.-K. Chong. 2007. Self-tuning experience weighted attraction learning in games. *Journal of Economic Theory* **133** 177-98.
- Kahneman, D. 1988. Experimental economics: A psychological perspective. In *Bounded Rational Behavior in Experimental Games and Markets*, R. Tietz, W. Albers, R. Selten, eds. Berlin: Springer-Verlag.

- Kahneman, D., A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* **47** 263-291.
- Kahneman, D., A. Tversky. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* **5** 297-323.
- Knight, F. H. 1924. Some fallacies in the interpretation of social cost. *Quarterly Journal of Economics* **38** 592-606.
- Loomes, G., R. Sugden. 1982. Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal* **92** 805-824.
- Lu, X., S. Gao, E. Ben-Elia. 2011. Information impacts on route choice and learning behavior in a congested network. *Transportation Research Record* **2243** 89-98.
- Mak, V., A. Rapoport. 2013. The price of anarchy in social dilemmas: Traditional research paradigms and new network applications. *Organizational Behavior and Human Decision Processes* **120** 142-53.
- Morgan, J., H. Orzen, M. Sefton. 2009. Network architecture and traffic flows: experiments on the pigou-knight-downs and Braess paradoxes. *Games and Economic Behavior* **66** 348-72.
- Rapoport, A., E. J. Gisches, T. Daniel, R. Lindsey. 2014. Pre-trip information and route choice with stochastic travel conditions: Experiment. Working paper.
- Rapoport, A., E. J. Gisches, V. Mak. 2013. Distributed decisions in networks: Laboratory study of routing splittable flow. *Production and Operations Management*, in press.
- Rapoport, A., T. Kugler, S. Dugar, E. J. Gisches. 2009. Choice of routes in congested traffic networks: experimental tests of the Braess paradox. *Games and Economic Behavior* **65** 538-71.

- Rapoport, A., V. Mak, R. Zwick. 2006. Navigating congested networks with variable demands: Experimental evidence. *Journal of Economic Psychology* **27** 648-66.
- Rapoport, A., D. A. Seale, E. Winter. 2002. Coordination and learning in large groups with asymmetric players. *Games and Economic Behavior* **39** 111-36.
- Rosenthal, R. W. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* **2** 65-7.
- Selten, R., T. Chmura, T. Pitz, M. Schreckenberg. 2007. Commuters' route choice behavior. *Games and Economic Behavior* **58** 394-406.
- Unnikrishnan, A., S. T. Waller. 2009. User equilibrium with recourse. *Networks and Spatial Economics* **9** 575-593.
- Vickrey, W. S. 1969. Congestion theory and transport investment. *American Economic Review* **59** 251-60.
- Wardrop, J. G. 1952. Some theoretical aspects of road traffic research. *Proc. Inst. Civil Eng. Part II* **1** 325-78.
- Wijayaratna, K. 2014. Investigating the Braess Paradox within the user equilibrium with recourse modelling framework. 32nd Conference of Australian Institutes of Transport Research.
- Zhang, L., D. Levinson. 2008. Determinants of route choice and value of traveler information. *Transportation Research Record* **2086** 81-92.

Figure 1 The traffic network in the experiment presented here in a highway analogy. All users start from the origin A and need to drive through the network to reach the “Interstate” via one of the four entrances D, G, I, or J. Every segment allows traffic in only one direction indicated by arrow. The congestion (delay) cost along each segment is a linear function of the number of players who traverse that segment, denoted by “ f ” in the figure.

Figure 2 Mean choice frequency of each route (i.e., number of players choosing the route) by round and condition with Condition RCC in black and Condition SCC in gray. Equilibrium predictions are indicated by dashed lines.

(a) Routes that pass by B

(b) Routes that pass by E

Figure 3 Percentage of players who switched routes, by round and condition (total number of players per condition = 90).

Figure 4. Histogram of route switching frequencies for each condition (total number of players per condition = 90).

Table 1 Equilibrium predictions, social optima, and mean observed route choice frequencies in Condition RCC. Also listed at the bottom are the total costs (summed over all players) in each case including mean observed total costs. For the observed means, the standard deviations (with group as the unit of observation) are indicated in parentheses.

<i>Route</i>	<i>Equilibrium predictions</i>		<i>Social optimum (1)</i>		<i>Social optimum (2)</i>		<i>Mean observed route choice frequency</i>	
	<i>No. of players</i>	<i>Cost per player</i>	<i>No. of players</i>	<i>Cost per player</i>	<i>No. of players</i>	<i>Cost per player</i>	<i>Condition RCC</i>	<i>Condition SCC</i>
ABCD	3	100	2	84	2	84	2.70 (0.18)	2.97 (0.21)
ABCG	1	100	1	93	1	93	1.27 (0.25)	1.01 (0.22)
ABFG	2	100	2	102	3	102	1.70 (0.25)	1.95 (0.16)
ABFI	0	NA	1	112	0	NA	0.64 (0.07)	0.34 (0.14)
AEFG	2	100	2	100	1	100	1.41 (0.15)	1.83 (0.26)
AEFI	0	NA	0	NA	1	110	0.73 (0.15)	0.31 (0.17)
AEHI	8	100	9	101	9	101	7.55 (0.43)	7.69 (0.31)
AEHJ	2	100	1	88	1	88	1.99 (0.17)	1.90 (0.17)
Total cost	1800		1774		1774		2038 (28.4)	2001 (20.2)

Table 2 Predicted and observed mean traffic flow in every segment in both conditions. The socially optimal traffic flow is also presented for comparison. Note that both of the socially optimal route choice distributions in Table 1 result in the same traffic flow as listed here. For the observed means, the standard deviations (with group as the unit of observation) are indicated in parentheses.

<i>Segment</i>	<i>Predicted (equilibrium) traffic flow</i>	<i>Social Optimum</i>	<i>Mean observed traffic flow</i>	
			<i>Condition RCC</i>	<i>Condition SCC</i>
AB	6	6	6.31 (0.23)	6.27 (0.16)
BC	4	3	3.98 (0.20)	3.98 (0.15)
CD	3	2	2.70 (0.18)	2.97 (0.21)
AE	12	12	11.69 (0.23)	11.73 (0.16)
BF	2	3	2.34 (0.30)	2.29 (0.27)
CG	1	1	1.27 (0.25)	1.01 (0.22)
EF	2	2	2.14 (0.26)	2.14 (0.18)
FG	4	4	3.11 (0.15)	3.78 (0.13)
EH	10	10	9.54 (0.31)	9.59 (0.17)
FI	0	1	1.37 (0.18)	0.65 (0.30)
HJ	2	1	1.99 (0.17)	1.90 (0.17)
HI	8	9	7.55 (0.43)	7.69 (0.31)

Table 3 Estimations of the learning model (standard error in parentheses). All parameter estimates are significantly positive at $p < 0.05$.

<i># Obs.</i>	<i>Condition RCC 4410</i>	<i>Condition SCC 4410</i>	<i>Pooled 8820</i>
λ	0.013 (0.0013)	0.016 (0.0016)	0.014 (0.0010)
λ_0	1.63 (0.041)	1.96 (0.042)	1.80 (0.029)
λ_+	0.0069 (0.0021)	0.012 (0.0026)	0.0094 (0.0017)
Log L	-7736	-7005	-14778

APPENDICES

Appendix A: Computation of Equilibrium Traffic Flow

In computing theoretical predictions, we focus on pure-strategy Nash equilibria for simplicity and tractability. The equilibria we obtain are also exactly Wardrop equilibria in that all users bear the same individual congestion cost in equilibrium, and any unused route incurs a higher than equilibrium cost to any user who deviates to that route unilaterally. It turns out that each of the two experimental conditions has a unique, pure-strategy equilibrium traffic flow (i.e. a complete specification of how many players traverse each segment) that, in fact, is approached by observed data, thereby providing additional ex-post justification to our analysis.

1. Segment choice condition (Condition SCC)

As described in the main text, the game in Condition SCC involves three moves. We begin our analysis at the third and last move when all the players have already arrived at either one of the nodes C, F, and H. For this last move, players at node C need to choose between segments CD and CG, players at node F need to choose between segments FG and FI, and players at node H need to choose between segments HI and HJ, all simultaneously and independently. Moreover, every player is informed about the exact number of players at the node that he/she is currently at. Hence for every node $i \in \{C, F, H\}$, and for every possible positive number of players that may be at i immediately before the third move (i.e. every element of the set $\{1, 2, 3 \dots 18\}$), we compute all the feasible Nash equilibrium traffic flow in which, given the number of players choosing either of the two segments from i , no player would be better off by unilaterally deviating his/her choice from one segment to another. The results of the computation are listed in Table A1(a), (b), and (c). For example, Table A1(a) shows that, if the number of players at C immediately before the third move is 17, the equilibrium traffic flow can either be: (a) seven

players traversing CD and 10 players traversing CG, or (b) eight players traversing CD and nine players traversing CG. This is also an example where there are multiple feasible equilibrium traffic flows; but in most cases, the equilibrium traffic flow is unique – although any equilibrium traffic flow with, say, n players traversing one segment and m traversing another, would imply $(m+n)!/(m!n!)$ equilibria each being a distinct mapping from the set of players to the set of segments.

We next proceed to the second move, when all the players have already arrived at either one of the nodes B and E. For this move, players at node B need to choose between segments BC and BF, and players at node E need to choose between segments EF and EH, all simultaneously and independently. Every player knows the exact number of players at both nodes B and E (note that the two numbers must add up to 18, the total group size). In any equilibrium, the cost for a player at node $j \in \{B, E\}$ in choosing a segment to node $i \in \{C, F, H\}$ for this move involves two additive components: (i) the cost of traversing segment ji , and (ii) the prospective cost in the third move when the player will be at node i . In equilibrium, the number of players traversing ji is known to all players – therefore (i) is straightforward – and so is the number of players at node i . Thus, for (ii), we assume that the traffic flow from i (i.e. the traffic “further down the network”) is an equilibrium traffic flow as listed in Table A1(a), (b), and (c), and the associated cost is the maximum among the costs of the traffic from i in the associated third-move equilibrium. For example, the equilibrium cost of using BF per player is equal to the equilibrium cost of traversing BF (given all players’ choices in equilibrium) + Max{Equilibrium costs in the traffic from F given the total number of players reaching F in equilibrium}. For a numerical example, when the number of players at B is 12, so that the number of players at E is 6, it turns out that the equilibrium cost of using BF per player is $(2 \cdot 7 + 26) + \max\{42, 40\} = 82$, where $2 \cdot 7 + 26 = 40$ is the

cost of traversing BF (given that a total of 7 players use BF in equilibrium) and $\max\{42,40\}$ is taken from Table A1(b) when the number of players at F is 7 (as will happen in equilibrium). With this assumption, for every possible number of players that may be at B immediately before the second move (i.e., every element of the set $\{0, 1, 2, 3 \dots 18\}$; the number at E must be 18 minus the number at B), we compute all the feasible Nash equilibrium traffic flow in which no player has an incentive to unilaterally deviate. Our results are listed in Table A1(d).

We now proceed to the first move, where all 18 players start at A and need to choose between segments AB and AE simultaneously and independently. As with the second move, the equilibrium cost of a player who choose segment A_j ($j \in \{B, E\}$) is the cost of traversing A_j plus the prospective cost of making the second move from node j . For the latter component of the cost, as with the corresponding procedure described above for computing second-move equilibrium traffic flow, we assume equilibrium traffic further down the network, i.e. the traffic from j is an equilibrium traffic flow as listed in Table A1(d), and the associated cost is the maximum among the costs of the traffic from j in the associated second-move equilibrium. It turns out that there is only one feasible equilibrium traffic flow, which is listed in Table A1(e). Overall, there is only one traffic flow which is in equilibrium all through the three moves, which is highlighted in gray throughout Table A1(a) to (e), summarized in column 2 in Table 2 in the main text, and forms our theoretical prediction for Condition SCC.

2. Route choice condition (Condition RCC)

It can be deduced from the network structure that the equilibrium traffic flow in Condition SCC is consistent with only one distribution of route choices among players, which is the distribution listed in column 2 in Table 1 in the main text. For example, because the equilibrium traffic flow in Condition SCC has three players traversing CD, exactly three players must have chosen

ABCD in any consistent route choice distribution; similarly, exactly one player must have chosen ABCG in any consistent route choice distribution; other similar deductions can be worked out accordingly. It is also straightforward to see that, for players in Condition RCC, if their route choices result in the distribution in column 2 in Table 1, then no player has an incentive to unilaterally deviate – so that the distribution is indeed an equilibrium distribution in Condition RCC.

It remains to show that the distribution in column 2 in Table 1 is the unique feasible pure-strategy equilibrium distribution in Condition RCC. We first prove that any such distribution of route choices in Condition RCC, if re-presented as equilibrium traffic flow (for example, the traffic flow along BC would be the number of players choosing ABCD + the number of players choosing ABCG; the rest can be worked out similarly), must also form a feasible equilibrium in Condition SCC. To do so, first observe that any unilateral deviation among such traffic flow at the third move is obviously equivalent to a unilateral deviation in route choice (e.g. a unilateral deviation from CD to CG is equivalent to a unilateral deviation from ABCD to ABCG). Next, consider a unilateral deviation at the second move, so that the number of players at nodes H, F, and G immediately before the third move will differ by only one compared with the original traffic flow. If, after this deviation, all subsequent third-move traffic follows the equilibria in Table 1S(a) to (c), then it can be confirmed by inspection of those tables that the resulting change in third-move traffic will only involve one player more in one segment and one player fewer in another. Since only one player is involved, we conclude that a unilateral deviation in the second move is also equivalent to a unilateral deviation in route choice. A similar reasoning, using Table 1S(d), can show that a unilateral deviation in the first move is also equivalent to a unilateral deviation in route choice. But since the original route choice distribution is supposed to be an

equilibrium in Condition RCC, so that no player has an incentive to deviate unilaterally, when it is re-presented as traffic flow in Condition SCC, no player in that condition should have an incentive to unilaterally deviate at any move either. Therefore, the traffic flow must be an equilibrium one in Condition SCC.

Since we already know that there is only one feasible equilibrium traffic flow in Condition SCC, the equilibrium route choice distribution in Condition RCC must also be uniquely the distribution listed in column 2 in Table 1, which then forms our theoretical prediction for Condition RCC.

Table A1 Numerically computed pure-strategy equilibrium traffic flow from every decision node in Condition SCC. Although there is only one overall equilibrium traffic flow as indicated in Table 2 (column 2) and Table A1(e), deviations from the overall equilibrium when players choose segments from A, B, or E could lead to various scenarios further down the network for which the equilibrium traffic flows are tabulated below. Traffic along the overall equilibrium path is highlighted in gray. See the main text of this Appendix for more details.

(a) Traffic from C

<i>Number of players at C</i>	<i>Equilibrium traffic: Number of players</i>		<i>Equilibrium costs: Cost per player</i>	
	<i>Along CD</i>	<i>Along CG</i>	<i>Using CD</i>	<i>Using CG</i>
1	1	0	12	NA
2	2	0	21	NA
3	2	1	21	30
3	3	0	30	NA
4	3	1	30	30
5	3	2	30	35
6	4	2	39	35
7	4	3	39	40
8	4	4	39	45
9	5	4	48	45
10	5	5	48	50
11	5	6	48	55
12	6	6	57	55
13	6	7	57	60
14	6	8	57	65
15	7	8	66	65
16	7	9	66	70
17	7	10	66	75
17	8	9	75	70
18	8	10	75	75

(b) Traffic from F

<i>Number of players at F</i>	<i>Equilibrium traffic: Number of players</i>		<i>Equilibrium costs: Cost per player</i>	
	<i>Along FG</i>	<i>Along FI</i>	<i>Using FG</i>	<i>Using FI</i>
1	1	0	12	NA
2	2	0	18	NA
3	3	0	24	NA
4	4	0	30	NA
5	5	0	36	NA
6	5	1	36	40
7	6	1	42	40
8	6	2	42	48
8	7	1	48	40
9	7	2	48	48
10	8	2	54	48
11	8	3	54	56
12	9	3	60	56
13	9	4	60	64
14	10	4	66	64
15	10	5	66	72
15	11	4	72	64
16	11	5	72	72
17	12	5	78	72
18	12	6	78	80

(c) Traffic from H

<i>Number of players at H</i>	<i>Equilibrium traffic: Number of players</i>		<i>Equilibrium costs: Cost per player</i>	
	<i>Along HI</i>	<i>Along HJ</i>	<i>Using HI</i>	<i>Using HJ</i>
1	0	1	NA	18
2	1	1	23	18
3	2	1	24	18
4	3	1	25	18
5	4	1	26	18
6	5	1	27	18
7	6	1	28	18
8	7	1	29	18
9	7	2	29	30
9	8	1	30	18
10	8	2	30	30
11	9	2	31	30
12	10	2	32	30
13	11	2	33	30
14	12	2	34	30
15	13	2	35	30
16	14	2	36	30
17	15	2	37	30
18	16	2	38	30

(d) Simultaneous traffic from B and E

<i>No. of players at B</i>	<i>No. of players at E</i>	<i>Equilibrium traffic: No. of players</i>				<i>No. of players reaching F in equilibrium</i>	<i>Equilibrium costs: Total cost per player</i>			
		<i>Along BF</i>	<i>Along BC</i>	<i>Along EF</i>	<i>Along EH</i>		<i>Using BF</i>	<i>Using BC</i>	<i>Using EF</i>	<i>Using EH</i>
0	18	0	0	3	15	3	NA	NA	67	75
1	17	0	1	3	14	3	NA	21	67	72
2	16	0	2	3	13	3	NA	37	67	69
3	15	0	3	3	12	3	NA	53	67	66
3	15	1	2	2	13	3	52	37	54	69
4	14	1	3	2	12	3	52	53	54	66
5	13	1	4	2	11	3	52	60	54	63
5	13	2	3	2	11	4	60	53	60	63
6	12	2	4	2	10	4	60	60	60	60
7	11	3	4	1	10	4	62	60	47	60
8	10	4	4	1	9	5	70	60	53	58
9	9	4	5	1	8	5	70	72	53	55
10	8	5	5	0	8	5	72	72	NA	55
11	7	6	5	0	7	6	78	72	NA	52
12	6	7	5	0	6	7	82	72	NA	49
13	5	7	6	0	5	7	82	83	NA	46
14	4	8	6	0	4	8	90	83	NA	43
15	3	8	7	0	3	8	90	91	NA	40
16	2	9	7	0	2	9	92	91	NA	37
17	1	10	7	0	1	10	100	91	NA	30
18	0	10	8	0	0	10	100	103	NA	NA

Note:

- (1) The total number of players at B and E must be equal to the group size i.e., 18.
- (2) The total cost per player is calculated assuming pure-strategy equilibrium traffic further down the network; see the main text of this Appendix for more details.

(e) Traffic from A (the whole game)

<i>Number of players at A</i>	<i>Equilibrium traffic: Number of players</i>		<i>Equilibrium costs: Total cost per player</i>	
	<i>Along AB</i>	<i>Along AE</i>	<i>Using AB</i>	<i>Using AE</i>
	6	12	100	100

Note:

- (1) The number of players at A must be equal to the group size i.e. 18.
- (2) The total cost per player is calculated assuming pure-strategy equilibrium traffic further down the network; see the main text of this Appendix for more details.

Appendix B: Instructions for Condition RCC

Introduction

Welcome to an experiment on route choice in traffic networks. During this experiment, you will be asked to make a large number of decisions and so will the other participants. Your decisions, as well as the decisions of the other participants, will determine your monetary payoff according to the rules that will be explained shortly.

Please read carefully the instructions below. If you have any questions, raise your hand and one of the supervisors will come to assist you.

Note that hereafter communication between the participants is prohibited. If the participants communicate with one another in any shape or form, the experiment will be terminated.

The Route Selection Task

The experiment is computerized. You will make your decisions by clicking on the appropriate buttons that will appear on your screen.

There are **18** participants in the lab today who will play together as one group for the duration of this experiment.

You will participate in **50** identical rounds of a routing game. On each round, you will be presented with a diagram of a traffic network with a given origin and **four** possible destinations, and you will have to choose a route from the origin to one of these destinations.

Description of the Game

Consider the traffic network that is exhibited in the diagram below. You are required to choose one of **eight** routes that connect the origin, denoted by the letter A (START), to one of the final destinations, denoted by J, I, G and D (END). All four destinations are equivalent and it does not matter which one is reached. The eight alternative routes in this traffic network are denoted in the diagram by the letter combinations [A-B-C-D], [A-B-C-G], [A-B-F-G], [A-B-F-I], [A-E-F-G], [A-E-F-I], [A-E-H-I] and [A-E-H-J].

(Please study this diagram.)

Traversing a network is always costly in terms of the time needed to complete each segment of the road, gas, tolls, etc. The travel costs are shown near each segment of the routes that you may choose. For example, consider segment [A-B] (the bottom-left segment); each participant that chooses this segment will be charged a total cost of $(6 \times ab + 4)$ where 'ab' (in the shaded box) indicates the number of participants (that ranges from 1 to 18) that choose this segment as part of their route. A similar cost structure applies to all the other segments in the network.

Please note that the cost charged for choosing each segment has two components. The first cost component is susceptible to congestion; it increases proportionally to the number of participants who choose that segment. The second cost component is not susceptible to congestion; it is constant and as such it is not affected by the number of participants choosing that segment.

For example, if only a single participant chooses segment [A-B], then his cost for this segment is $6 \times 1 + 4 = 10$. If, instead, 5 participants choose the same segment, then the cost of this segment for each of them is $6 \times 5 + 4 = 34$.

All the participants choose routes independently of one another, leaving the origin A and traveling to one of the destinations J, I, G, or D at the same time.

Example:

Suppose that on a certain round, you choose route [A-B-F-G] and altogether, out of the 18 participants:

2 participants choose [A-E-H-J],

2 participants choose [A-E-H-I],

1 participant chooses [A-E-F-I],

3 participants choose [A-E-F-G],

3 participants choose [A-B-F-I],

2 participants (you and one other participant) choose [A-B-F-G],

2 participants choose [A-B-C-G],

3 participants choose [A-B-C-D],

At the end of such a round the feedback screen shown below will be presented to you:

On the left side of the screen you can see the eight different routes. For each route, the computer informs you the number of participants that chose that route (**purple**), the cost of the route for each of them (**red**), and their payoff (**black**). At the bottom of the screen, in **blue**, you can view your personal payoff.

On the right side of the screen you can view the entire network. The route you have chosen is highlighted in **blue** ([A-B-F-G]). The number of participant that chose each segment is shown in

pink (e.g., 10 for segment [A-B]). You can also view the resulting cost for each segment in **red** on white background (e.g., **64** for segment [A-B]).

To continue the current example, the total cost incurred for choosing route [A-B-F-G] was:

Segment [A-B]: $6 * \underline{10} + 4 = 64$ { 10 = 2 from [A-B-F-G] + 3 from [A-B-C-D] + 2 from [A-B-C-G] + 3 from [A-B-F-I] }

Segment [B-F]: $2 * \underline{5} + 26 = 36$ { 5 = 2 from [A-B-F-G] + 3 from [A-B-F-I] }

Segment [F-G]: $6 * \underline{3} + 6 = 24$ { 10 = 2 from [A-B-F-G] + 1 from [A-E-F-G] }

for a total cost of $64 + 36 + 24 = 124$ for route [A-B-F-G] in this round.

At the end of *each round*, you will receive a reward of **140** points for reaching the destination. Your earnings will be determined by subtracting your cost for the round from this reward.

Back to the example (*see the screen above*), the earnings for the different routes are:

for participants that chose route [A-E-H-J], payoff = $140 - 76 = 64$ points

for participants that chose route [A-E-H-I], payoff = $140 - 70 = 70$ points

for participants that chose route [A-E-F-I], payoff = $140 - 164 = -24$ points

for participants that chose route [A-E-F-G], payoff = $140 - 108 = 32$ points

for participants that chose route [A-B-F-I], payoff = $140 - 180 = -40$ points

for participants that chose route [A-B-F-G], payoff = $140 - 124 = 16$ points

for participants that chose route [A-B-C-G], payoff = $140 - 136 = 4$ points

for participants that chose route [A-B-C-D], payoff = $140 - 131 = 9$ points

Notice that it is possible for some players to end up with negative points for a certain round as shown in the example above for players who chose routes [A-E-F-I] and [A-B-F-I] in the current example.

All 50 rounds will have exactly the same structure.

Procedure

At the beginning of each round, the computer will display a network diagram similar to the one present before. You will then be asked to choose one of the eight possible routes. To choose a route, using your mouse, simply click on the segments comprising that route. The color of the segments that you choose will change to blue indicating your choice. To change your choice, please click on a chosen segment again and its color will change back to black. Once you are satisfied, press the "Confirm" button. You will be asked to verify your choice.

Once all the participants confirm their choice of routes, you will receive a feedback screen similar to the one presented in the example before.

Payments

At the end of the experiment, you will be paid for your cumulative earnings in all the 50 rounds. You will be paid in cash for your earnings with an exchange rate of \$1=**100** points. In addition, you will receive a show up bonus of **\$5** for attending the experiment.

Once you are certain that you understand the task, please place the instructions on the table in front of you to indicate that you have completed reading them. If you have any questions, please raise your hand and one of the supervisors will come to assist you.

The experiment will begin shortly. Thank you for your participation.

Appendix C: Supplementary Figure

Figure A1 The coefficient of variation (CV) of travel costs among players (defined as the standard deviation of the costs divided by the average cost) by session and round in both conditions. The equilibrium prediction for the CV is zero.