

Network Architecture in a Route Choice Experiment ^{*}

by

John Morgan^a, Henrik Orzen^b and Martin Sefton^c

June 2006

Abstract

This paper presents theory and experiments to investigate how network architecture influences route-choice behavior. We consider network changes illustrating the Pigou-Knight-Downs and Braess Paradoxes. We show that these paradoxes are examples of general network change properties that we term the “least congestible route” and “size” principle, respectively. We find that technical improvements to networks induce traffic flow adjustments in the direction predicted by equilibrium theory. In the case of the Pigou-Knight-Downs Paradox, these adjustments undermine short-term payoff improvements. In the case of the Braess Paradox, they reinforce the counter-intuitive, but theoretically predicted, effect of reducing payoffs to network users.

Acknowledgements

We thank seminar participants at the University of Leicester and at the Amsterdam 2004 International Meeting of the Economic Science Association for helpful comments.

a. Haas School of Business and Department of Economics, University of California, Berkeley, CA 94720-1900. e-mail: morgan@haas.berkeley.edu

b. School of Economics, University of Nottingham, Nottingham, NG7 2RD, United Kingdom. e-mail: henrik.orzen@nottingham.ac.uk

c. School of Economics, University of Nottingham, Nottingham, NG7 2RD, United Kingdom. e-mail: martin.sefton@nottingham.ac.uk

1 Introduction

Individuals and businesses are often presented with situations where they must decide the fastest way to get information, materials, or simply themselves from point A to point B.

A business in Seattle needs to ship a prototype of some new product for a client meeting in Shanghai. How should it do it? What transportation network should it use? How long will it take for the prototype to arrive? If the timing of the meeting is around a major holiday, such as Christmas, the anticipated time along certain networks might be slower because of others wishing to use the network.

A driver commuting from Antioch, California to San Francisco on a weekday morning faces a myriad of route choices in deciding how to proceed. Should the driver head down interstate 680 and through the Caldecott Tunnel, which could be crowded? Or should the driver come down interstate 80 and avoid the tunnel but face what is often a congested corridor around Berkeley?

A packet of information associated with a streaming audio file needs to get from the Verizon server to a user's computer in Walnut Creek, California. Should it go via a switch in Oakland or through San Ramon? Will it get stuck in the copper wires connecting the "last mile" of the Internet to the user's house? Will it get stuck inside the house when the user's wife makes a large file transfer over the wireless connection?

In all of these situations, a user's experience in the speed of arriving at his or her desired destination is affected by the actions of others. Most consider the time spent traveling, whether it be waiting for a page to load while surfing the web or fuming while sitting in a traffic jam, to be an undesirable cost of getting from where they are to where they want to go. Such people should make careful decisions about which route to take: choosing the right route might save considerable time. The best route could be the most direct one, but in general the best route will also depend on the congestion on alternative routes. Thus, the optimal route choice for an individual will depend on the route choices made by other individuals. In such a setting patterns of traffic flow may differ from socially efficient patterns, as individual travelers attempt to minimize their own travel cost without taking into account the effect of their actions on other travelers. For providers and managers of networks, attempts to improve traffic flow by modifying the network may also have unintended consequences, as technical improvements in

the network can induce changes in journey planning that, while sensible for each individual concerned, increase overall travel times.

This paper uses theory and laboratory experiments to investigate how network architecture influences route-choice behavior by comparing outcomes across several different networks. In the theory section, we derive two principles for planners. In what we term the **least congestible route principle** we show that the planner most effectively reduces travel time by first improving the route that is *least sensitive* to network congestion. The least congestible route principle takes an extreme form when one of the routes along the network is non-congestible. In that case, the benefits from any improvement to a congestible route are completely dissipated—this observation is the essence of the Pigou-Knight-Downs Paradox. That paradox has been described as being “so enshrined in transportation planning that it is often called the ‘fundamental law of traffic congestion’” (Arnott and Small, 1994).


The essence of the least congestible route principle and the Pigou-Knight-Downs Paradox is that as one link is improved, the volume of traffic on that link will increase until it is no more attractive than alternatives. Imagine a population of individuals who can choose between two routes into town. One route is non-congestible and takes one hour, regardless of the number of people using it. The other route is congestible, so that the travel time increases with the number of users. In equilibrium the number of individuals using the congestible route must be just such that the travel time on it is one hour. Now imagine that the congestible route is improved. This will encourage people to switch from the non-congestible to the congestible route until, once again, the travel times on both routes is one hour. Thus, the benefits from improving the route are completely dissipated.¹

Next, we introduce the **size principle**. This principle states that efficiency gains from adding network links always outweigh losses from user externalities once the number of network users is sufficiently large. The converse of the size principle is that in small networks user externalities from new choices can lead to worse outcomes. A striking example of this possibility is the Braess Paradox (Braess, 1968), which illustrates that adding a link on a congestible network can increase travel times for all users.

¹ This will be true so long as both options are used in equilibrium. If the congestible route takes less than an hour when the entire population uses it, then further improvements will reduce travel times.

A particularly simple version of the size principle/Braess Paradox is shown in Figure 1. Suppose 60 individuals want to get from O (origin) to D (destination), and minimize the time spent doing so. The left panel shows the two possible routes: O-H-D or O-L-D. It takes 60 minutes to get from O to L, or from H to D. It takes n_{OH} minutes to get from O to H, and n_{LD} minutes to get from L to D, where n_i is the number of travelers using link i . In equilibrium, half of the population takes the route O-H-D, half takes the route O-L-D, and each individual spends 90 minutes traveling. Now, if a link is added between H and L that takes no time at all (right panel), this will induce all individuals to use the route O-H-L-D, and each individual will spend 120 minutes traveling. This worsening result is the Braess Paradox.

Figure 1: The Braess Paradox



In contrast if the population, N , grows sufficiently large (i.e. $N > 120$) the efficiency of the additional link overcomes the adverse strategic effects. This is because each user will still take 120 minutes with the link, but her travel time will be $60 + N/2$ minutes without. This reversal illustrates the size principle.

While the least congestible route and size principles and their accompanying paradoxes reflect theoretical results stemming from equilibrium behavior, their behavioral plausibility, and thus their applicability as planning tools, is far from clear. Moreover, because of the very interconnectedness of traffic, data, and transportation networks, there is scant hope of finding “natural experiments” to cleanly and convincingly test these predictions. As a result, we turn to controlled laboratory experiments to examine individual behavior on networks.

In our experiments individuals make route choices on shared and congestible networks. Our subjects repeatedly choose routes over networks, receiving a reward for reaching their destination and incurring travel costs for each journey. We do not inform subjects of the travel cost functions associated with each road, rather, at the end of a journey we inform them of the travel cost they actually incurred and the number of other subjects using each road that they used. Thus, subjects can become familiar with travel costs from their travel experiences. We ask whether stable patterns of traffic flow emerge as individuals interact repeatedly on a network, whether traffic flows adjust in response to network changes, and if so how these relate to equilibrium and efficient traffic flows and adjustments.

We find that subjects do adjust to changes in network structure, altering travel patterns in an attempt to reduce travel times. In our sessions examining the Pigou-Knight-Downs Paradox we improve a congestible route and observe a significant increase in the amount of traffic using that route. In our sessions examining the Braess Paradox we find that adding a new road results in a shift of travelers towards the congestible links. These results thus reinforce a theoretical point about network design: when considering the consequences of modifying a network, one should pay careful attention to induced changes in the behavior of network users.

For the Pigou-Knight-Downs networks both equilibrium and efficiency entail a shift in traffic towards the improved congestible route. However, the two concepts deliver different point predictions, and actual adjustments are closer to those implied by equilibrium. Improvements to the congestible route on the Pigou-Knight-Downs network leads to a small reduction in travel times, and an even smaller reduction when we focus on the travel choices of experienced subjects. For the Braess networks equilibrium entails a shift of traffic in the direction we observe, whereas efficiency would entail a shift in the opposite direction. The observed shift in route choices has the effect of increasing travel times, and the increase in travel times is even greater when we focus on the choices of experienced subjects. Thus the Braess Paradox is observed in our laboratory setting.

Although aggregate route choice behavior is close to that predicted by equilibrium, we see some systematic deviations from equilibrium predictions. One important deviation we observe is persistent variability in route choices. In general, variability has the effect of increasing average travel times. For the unimproved Pigou-Knight-Downs network in particular, variability has critical consequences for network performance. While we see the congestible road

under-utilized relative to equilibrium, and other things equal this would result in lower travel times than equilibrium theory predicts, the adverse effect of variability in traffic flows outweighs this effect, and travel times are in fact even higher than in equilibrium. Moreover, even though the congestible route is faster, on average, than the non-congestible route, the average subject would have been worse off had they used the faster congestible route all the time, and would have been better off had they used the slower, non-congestible, route all the time.

We investigated further the discrepancies between observed and equilibrium behavior. Using a simple reinforcement learning model we are able to capture the qualitative features of these discrepancies. In particular, the learning model results in heterogeneous behavior among players that generates similar patterns of variability in route choice behavior to those observed in the experiment. Furthermore, we also observe a systematic tendency toward a more even distribution of traffic than equilibrium predicts, and this too is an implication of the simple learning model.

Experimental economists have only recently begun studying route choice behavior. The papers most closely related to ours are Selten et al. (forthcoming) and Rapoport et al. (2005a, 2005b). Selten et al. examine how subjects choose between two congestible routes on a fixed network and find that traffic flows do not converge to equilibrium, but rather fluctuate around the equilibrium prediction throughout the experiment. However, their focus is on the effect of information on route choices over a fixed network, whereas our focus is on the effect of network changes. Rapoport et al. compare network changes corresponding to the Braess Paradox. Like us, they find that equilibrium does a good job of predicting mean traffic flows, that fluctuations about the mean are persistent, and that the paradox is observed in the laboratory. Unlike us, they use a full information design (i.e. all subjects are fully informed of the cost functions on alternative routes), as well as different network structures and larger groups (18 subjects per network).

As far as we are aware, the Pigou-Knight-Downs Paradox has not been previously studied in the lab. However, the network game on which it is based resembles market entry games in which players can decide to enter, in which case they receive a payoff that is decreasing in the number of entrants, or not enter, in which case they receive a fixed outside option. In the first experimental market entry games Kahneman (1987) found that the number of entrants was very close to the number predicted by theory. Camerer (2003) provides a useful review of

subsequent experiments, and notes slight tendencies toward excess entry when equilibrium predicts few entrants, and under-entry when equilibrium predicts many entrants.² Also related are experiments on political participation games (see Goeree and Holt, forthcoming). In a variety of experimental games involving simultaneous binary decisions Goeree and Holt report variability in decisions and a bias toward a more even employment of the two options than equilibrium predicts.³ These results are broadly consistent with ours; in our experiment equilibrium predicts aggregate choices well, and where we see deviations from equilibrium, they tend to involve a more even allocation of subjects to the two routes than equilibrium predicts.

The remainder of the paper proceeds as follows. In section 2, we derive and discuss the least congestible route and size principles. Our experimental design and procedures are described in section 3, and the results in section 4. Finally, section 5 relates our results to other studies and offers concluding comments.

2 Theory

Consider a network in which there are N travelers each of whom wishes to get from some origin, O , to a destination, D , to which there are k possible routes. Other travelers are faced with the same decision and each makes his choice simultaneously. Suppose further that the reward from reaching the destination is sufficient that all individuals find it in their interest to travel regardless of the actions of others. A traveler's problem, of course, is to minimize the travel time in getting to D . Let the travel time along route i be given by $T = \alpha_i + \beta_i n_i$ where $\alpha_i, \beta_i \geq 0$ are parameters of the model and n_i denotes the number of other drivers using route i . The parameter α_i should be interpreted as the fixed time factor to traverse a road in the absence of any congestion. The parameter β_i is a measure of the marginal effect on travel times of congestion. Routes with high values of β_i are more impacted by an additional driver than are low β_i roads.

² There are exceptions to this pattern. For example, Fischbacher and Thöni (forthcoming) observe a systematic bias toward excess entry in experimental "winner-take-all" markets. Although their game has a similar structure to a market entry game, it has two features not shared by market entry games that they argue may account for the different results. First, in their game the prize increases in the number of entrants; second, rather than all entrants receiving the same payoff, one entrant wins the prize.

³ Goeree and Holt use the quantal response equilibrium concept (McKelvey and Palfrey, 1995) to explain the variability and bias.

An *equilibrium* in this game is a vector of the number of travelers using each of the k routes (n_1, n_2, \dots, n_k) such that no individual can profitably deviate by changing his or her route choice. We simplify the analysis by ignoring the integer constraint for the number of travelers; that is, technically, we study the notion of quasi-equilibrium as defined in Ellison *et al.* (2004). A quasi-equilibrium will be arbitrarily close to a Nash equilibrium for N sufficiently large. We further simplify by restricting attention to *viable routes*—those that are taken by a strictly positive numbers of travelers in equilibrium. Formally, we assume that

$$N \geq \sum_{i \neq j}^k \frac{\alpha_i - \alpha_j}{\beta_j} \quad \text{for all } i.$$

Finally, we will sometimes be interested in the case where one of the routes, which shall always be designated as route 1, is not congestible, i.e. $\beta_1 = 0$. Together with route viability, this requires that $\alpha_j < \alpha_1$ for all j .

An equilibrium is a solution to the system of equations

$$\begin{aligned} T &= \alpha_i + \beta_i n_i \quad \text{for all } i \\ \sum_{i=1}^k n_i &= N. \end{aligned}$$

Rewriting the representative equation i above for all i , we obtain

$$n_i = \frac{T - \alpha_i}{\beta_i}.$$

Summing over all i

$$N = T \sum_{i=1}^k \frac{1}{\beta_i} - \sum_{i=1}^k \frac{\alpha_i}{\beta_i}.$$

Solving for T then yields the equilibrium travel time of

$$T^* = \frac{\sum_{i=1}^k \frac{\alpha_i}{\beta_i} + N}{\sum_{i=1}^k \frac{1}{\beta_i}} \quad (1)$$

and the equilibrium number of travelers on route i is

$$n_i^* = \frac{T^* - \alpha_i}{\beta_i} . \quad (2)$$

Route Improvements:

Suppose that there is funding for a small road improvement on one of the k routes. Which route should be improved? One approach might be to improve the road where the marginal driver is having the greatest impact on travel times. In effect, this marginal driver is determining the equilibrium travel time for all drivers (since if he or she had a better route, it would contradict the idea that the given configuration of route choices comprises an equilibrium). Thus, it would seem that reducing the marginal impact of this driver would be helpful. As we will show below, this intuition fails to account for general equilibrium readjustment of driving patterns across all routes and is, in fact, the worst possible use of the funds for road improvement. Instead, we show that the **least congestible route principle** is optimal:

Least congestible route principle: Improvements should be made on the route **least** sensitive to congestion.

To formalize the least congestible route principle, we consider the optimal route in which to make an infinitesimal change in the α_j parameter. The equilibrium effect on travel times of such a change is

$$\frac{\partial T^*}{\partial \alpha_j} = \frac{1}{\beta_j \sum \frac{1}{\beta_i}} . \quad (3)$$

Notice that the magnitude of this change is inversely related to β_j ; hence, the smaller is β_j the greater is the reduction in equilibrium travel times of the improvement to route j .⁴

When $\beta_1 = 0$, the *Pigou-Knight-Downs Paradox* emerges as a special case of the least congestible route principle. From equation (1), using L'Hopital's rule, we obtain $T^* = \alpha_1$ and it follows from equation (3) that improvements to routes other than to route 1 produce **no reduction** in travel times whatsoever. Improving route 1, on the other hand, is optimal and leads to a one for one reduction in equilibrium travel time.

Next, consider the impact of a change of β_j on equilibrium travel times. Here we have

$$\frac{\partial T^*}{\partial \beta_j} = \frac{\frac{1}{\beta_j^2} \left(N + \sum \frac{\alpha_i}{\beta_i} \right) - \frac{\alpha_j}{\beta_j^2} \sum \frac{1}{\beta_i}}{\left(\sum \frac{1}{\beta_i} \right)^2} = \frac{1}{\beta_j^2} \frac{T^* - \alpha_j}{\sum \frac{1}{\beta_j}} . \quad (4)$$

The least congestible route principle holds for this type of change as well if we add the condition that the ordering of α_i is the same as the β_i ordering of the parameters; that is, if we assume $\alpha_i < \alpha_j$ if and only if $\beta_i < \beta_j$.


More Complicated Network Structures

Next, we slightly complicate the network structure. Suppose there are k links, each of which leads from the origin O to a midpoint M_i , $i = 1, \dots, k$. As usual, the time cost for link i is $T = \alpha_i + \beta_i n_i$. Furthermore, suppose there are another k links, one from each midpoint M_i to the destination, D . The time cost for one of these links, link j , is given by $T = \gamma_j + \delta_j n_j$.

Finally, suppose that one of the links from the origin to a midpoint and one of the links from a midpoint to the destination are non-congestible. Without loss of generality, we let these be the first among the k links from the origin and the last among the k links to the destination: $\beta_1 = 0$, $\delta_k = 0$. Figure 2 shows a representative network with this structure.

⁴ While we have shown the improvement principle for networks with linear congestion effects, the result holds more generally since even for networks with nonlinear congestion, the comparative static property under consideration will represent a linearization.

Figure 2: A network with k midpoints




Notice that the original network is nested as a special case of this network. As usual, we restrict attention to parameter values where all routes are used. When there are no further connections in the network, it follows from equation (1) that the equilibrium travel time is

$$T^* = \frac{N + \sum \frac{\alpha_i + \gamma_i}{\beta_i + \delta_i}}{\sum \frac{1}{\beta_i + \delta_i}}.$$

Costless Connection: Now consider how the performance of the network changes when we add costless links connecting node M_i with M_j for all i, j . That is, with the presence of these links, any traveler can costlessly switch from any node M_i to any other node M_j . The key question is what the addition of such links does to the performance of the network. A representative network with this structure is shown in Figure 3 (where, with a slight abuse of notation, for each link we use n to denote the number of users of that link).

Figure 3: Costless switching between the k midpoints



As usual we assume that all links are viable. We offer the following:

The Size Principle: Adding costless links reduces travel time if and only if there are a sufficiently large number of travelers.

The possibility that the addition of costless links could adversely affect outcomes stems from the externality that “midpoint switching” by one traveler might impose on other travelers. At the same time, the prospect of free midpoint switching effectively increases the number of possible paths in the network and this improves the prospect of finding a more efficient route. The size principle suggests that, as the number of travelers grows large, the efficiency effect always dominates the externality effect.

To see the “size principle” formally, notice that the equilibrium in the network with costless midpoint connections essentially decomposes into two simple network problems. In equilibrium, the travel times between O and all nodes M_i must be equal, and similarly the travel time between all nodes M_i and D must be equal. Since $\beta_1 = \delta_k = 0$, this immediately implies that the equilibrium travel time on the “improved” network is $T^{**} = \alpha_1 + \gamma_k$.

To determine the effect of the new links on equilibrium travel times we must compare T^* to T^{**} . For the improvements to actually reduce travel time requires that

$$N > \sum_{i=1}^k \frac{(\alpha_1 - \alpha_i) + (\gamma_k - \gamma_i)}{\beta_i + \delta_i}. \quad (5)$$

Equation (5) implies that the efficiency effect dominates the externality effect if and only if N is sufficiently large. When N is small the externality effect dominates—the *Braess Paradox* is a simple illustration of this.

These general theoretical results motivate our experiment, which we describe in the next section.


3 Experiment

While the previous section derived properties of varying network architectures using quasi-equilibrium and ignoring integer constraints, implementing networks in a laboratory setting necessitates selecting particular parameter values and dealing with consequent integer issues associated with the number of network users. Specifically, in all the networks we study there are eight network users and, as we show below, we select parameter values such that the associated quasi-equilibrium number of travelers on each route is an integer.

Pigou-Knight-Downs (PKD) Treatments

Figure 4 displays the first two networks used in our experiment. For the road between L and D the cost is nine times the number of travelers using that road. For the other roads the costs are fixed. We will refer to the route O-H-D as the “high road” and the route O-L-D as the “low road”. In the PKD Baseline network (left panel) efficiency, in the sense of minimizing average travel time, requires that two travelers use the low road and six use the high road. The average travel time with this configuration is 52.5. The PKD Improved network (right panel) has a lower cost of getting from O to L, and allows a planner to reduce the travel time of network users. In an efficient profile of choices three travelers use the low road and five use the high road, giving an average travel time of 46.875. Thus if travelers use the network efficiently the improvement will reduce travel times.

Figure 4: PKD networks implemented in the experiment



However, these efficient outcomes do not constitute equilibria, as in each case there is an incentive for high road users to switch to the low road. From equation (1), equilibrium travel time is determined by travel time on the high road: $T^* = 57$. Substituting into equation (2), we find that four travelers will take the high road and four the low road in PKD Baseline.⁵

Turning to the PKD Improved network, once again from equation (1), we have that the equilibrium travel time is $T^* = 57$. Substituting into equation (2), we find that 2 travelers take the high road while 6 travelers take the low road in equilibrium. This swing in the number of travelers towards the low road drives the Pigou-Knight-Downs Paradox—the benefits from the improvement are *completely dissipated* by a re-allocation of traffic from the high to the low road.

We summarize the theoretical properties of the PKD networks in Table 1.

Table 1: Efficient and equilibrium traffic on PKD networks


| | | Baseline | Improved |
|--------------------|------------------------------|----------|----------|
| Efficiency | <i>Travelers on low road</i> | 2 | 3 |
| | <i>Average travel time</i> | 52.5 | 46.875 |
| Equilibrium | <i>Travelers on low road</i> | 4 | 6 |
| | <i>Average travel time</i> | 57 | 57 |

⁵ There is an additional pure strategy equilibrium where five travelers take the high road and travel times on the two routes are not equalized. In this equilibrium, even though low road travelers enjoy faster travel times, high road travelers are indifferent between switching or not. The additional equilibrium is an inevitable consequence of choosing parameters so that the quasi-equilibrium where travel costs are equalized involves an integer number of users on any particular route. Note that the additional equilibrium does not survive a simple refinement. If there is an (arbitrarily small) ε chance that a subject does not travel, then the unique pure strategy equilibrium corresponds to the quasi-equilibrium.

Braess Treatments

Figure 5 illustrates our implementation of the Braess networks. In the Braess Baseline (left panel) efficiency requires that 3 travelers take the low road and 5 travelers take the high road, giving an average travel time of 55.5. In contrast, equation (2) shows that in equilibrium four travelers take the high road and four take the low road, giving each traveler a travel time of 57.

Figure 5: Braess networks implemented in the experiment



Now suppose a costless link is added between L and H (Figure 5, Right Panel). Since this allows for traffic to be diverted between the high and low roads it allows for an improved allocation of traffic on the network and lowers (slightly) the average travel time when traffic flows efficiently. The costless link in effect makes the network architecture equivalent to two PKD networks—one from O to the midpoint and one from the midpoint to D. Equilibrium on the left part of the network implies that the travel time from O to the midpoint will be 21 and consists of 7 travelers leaving on the high road. Similarly, the travel time from the midpoint to D will be 45 and consists of only 3 travelers arriving on the high road. The travel time increases to 67 per traveler. This is the essence of the Braess Paradox—the ability of each traveler to freely switch between the congestible routes worsens overall congestion on the network.

We summarize the theoretical properties of the Braess networks in Table 2.

Table 2: Efficient and equilibrium traffic on Braess networks

| | | Baseline | Improved |
|--------------------|---------------------------------------|-----------------|-----------------|
| Efficiency | <i>Travelers leaving on low road</i> | 3 | 4 or 5 |
| | <i>Travelers arriving on low road</i> | 3 | 2 or 3 |
| | <i>Average travel time</i> | 55.5 | 54.75 |
| Equilibrium | <i>Travelers leaving on low road</i> | 4 | 1 |
| | <i>Travelers arriving on low road</i> | 4 | 5 |
| | <i>Average travel time</i> | 57 | 67 |

Procedures

The experiment was conducted at the University of Nottingham in Spring 2004. Subjects were recruited from a university-wide subject pool comprised of undergraduates who had indicated a willingness to be paid volunteers in decision-making experiments. Six sessions with 16 subjects were conducted, with no subject participating in more than one session. Thus 96 subjects in total participated in the experiment.

In each session two groups of eight subjects were formed, and no interaction across groups took place. Thus each group of eight can be considered as an independent observation. Throughout a session, no communication between subjects was permitted, and all choices and information were transmitted via computer terminals. At the beginning of a session, the subjects were seated at computer terminals and given a set of instructions, which were then read aloud by the experimenter.⁶ The session then consisted of two phases of sixty rounds. In each round subjects could earn points, and at the end of the session subjects were paid based on their accumulated point earnings from all 120 rounds using an exchange rate of 1 penny for every 5 points earned. Earnings averaged £10.00 for sessions lasting between 50 and 60 minutes.⁷

In each round subjects had forty seconds to complete a route on a road map by clicking on roads. If they failed to complete a route in the forty seconds they received no points in that round. If they completed a route they received a reward of 100 points for arriving at their destination and paid a travel cost.

⁶ A copy of the instructions is included as Appendix A.

⁷ At the time of the experiment the exchange rate was approximately £1 = \$1.86.

The actual costs were those presented above. However, subjects were only told that costs would be less than 100, would be calculated in the same way in every round of a phase, and might depend on the number of other subjects using the road. At the end of a round each subject was informed of the cost they incurred from each road used, the number of subjects using those roads, and their own total travel cost and point earnings. In making this design choice we wanted to reproduce a natural scenario where network users learn about the costs of using routes from experience. At the end of a given journey, users discover how long the journey took, but not how long it would have taken had they chosen an alternative route. They also learn the volume of traffic sharing that route, so with sufficient experience of congestible routes they may learn how travel costs on these are related to traffic volume. From their long-term experiences involving journeys on alternative routes, users may even be able to estimate the cost functions associated with all routes. Even in this latter case, network users are unlikely to learn exactly which route will maximize their payoff, since this will depend on the choices of others and choices are made simultaneously.

In three sessions the PKD networks were used. One group of eight subjects experienced the PKD Baseline network in phase 1 and the PKD Improved network in phase 2, and the other group of eight subjects experienced the PKD Improved network in phase 1 and the PKD Baseline network in phase 2. Similarly, three sessions used the Braess networks, with one group experiencing the Braess Baseline network in phase 1 and the Braess Improved network in phase 2, with the other group experiencing the networks in the reverse order.

4 Results

In examining the results, we use several different performance metrics to compare the implications of equilibrium theory with two alternatives. The first alternative, which we term the “reduced form” heuristic, suggests that individuals do not respond to changes in the cost structure of the network. Under this heuristic, a social planner uses traffic flows from a given network and changes in cost parameters to determine the benefits of a modified network on travel times. The second alternative, which we term the “efficiency hypothesis,” postulates that individuals self-organize into efficient traffic flow patterns.

Throughout, we divide the analysis into the “short-run”—the first 30 periods under a given network structure—and the “long run”—the last 30 period under a given network structure.

Aggregate Route Choice and Travel Time per Route

The first performance metric we examine is aggregate route choice. To measure this, we compute the average number of travelers using a given route under a given network structure in the short-run and the long-run. Tables 1 and 2 display the predictions of the equilibrium and efficient hypotheses with respect to route choice. The reduced form theory offers no level prediction for route choice, but does hypothesize no change in route choices in response to a change in the network.

Our second metric is the average travel time associated with each route. Specifically, we compute the per period travel time for a given route in a given network. The equilibrium hypothesis under this metric predicts that that travel time per route will be equalized across routes. In contrast, efficiency requires that routes will experience unequal travel times. For instance, in the PKD networks, efficiency requires systematically shorter travel times along the congestible route compared to the non-congestible route.

Network Performance

The third performance metric we examine is the average travel time (or latency) occurring in the network as a whole. To measure this, we compute the average experienced travel time for a given network. Notice that either equilibrium or efficient traffic flows imply a one to one mapping between aggregate route choice and average experienced travel time. In general, however, this is not the case. To see this, let n_i be the number of users of road i , so that the journey time on that road is $\alpha_i + \beta_i n_i$. This structure can be used to describe all our experimental networks. For a stochastic distribution of route choices, the expected experienced travel time for a network user is

$$E \left[\frac{1}{N} \sum_{i=1}^k n_i (\alpha_i + \beta_i n_i) \right]$$

Running the expectations operator through this expression and simplifying yields

$$\frac{1}{N} \sum_{i=1}^k (E[n_i](\alpha_i + \beta_i E[n_i]) + \beta_i \text{Var}[n_i]). \quad (6)$$

Equation (6) shows that expected travel times depend on the variance of route choices as well as the route choice metric. Moreover, equation (6) implies that variability of route choice leads to longer experienced travel times.

Table 2 displays the predictions of the equilibrium and efficiency hypotheses with respect to this metric. The reduced form theory again offers no level prediction, but does offer a directional prediction of the effects of changing network structure relative to baseline behavior.

4.1 PKD Networks

4.1.1 Aggregate Route Choices and Travel Time per Route

For the PKD networks the data consist of route choices made over 5,760 subject-rounds (48 subjects x 120 rounds). In only 7 of these (about one tenth of one percent) did a subject fail to complete a route in the forty seconds allowed.⁸ Table 3 displays the average number of travelers choosing the low road for each group. The first three groups listed experienced the Baseline network first, while the second three experienced the Improved network first. Recall that both equilibrium and efficiency imply that the improvement in the network will lead to a shift in traffic onto the improved low road. In contrast, the reduced form heuristic predicts no shift in traffic. As Table 3 shows, the network improvement led to a shift in traffic to the low road—even in the short run—in every group. Thus, there is a significant shift in the direction predicted by equilibrium and efficiency (p-value = 0.031).⁹

Equilibrium and efficiency make specific predictions about the amount of traffic on each road. Table 3 shows that there are systematically fewer travelers on the low road than in equilibrium for both Baseline (prediction: 4) and Improved (prediction: 6) in every group in both the short-run and the long-run. Thus, we can reject the equilibrium hypothesis (p-value = 0.031).

⁸ On the Braess networks in only 3 of 5,760 subject-rounds did a subject fail to complete a route. Note that this small chance of failing to complete a journey could be interpreted as a “tremble” consistent with the refinement discussed in footnote 3.

⁹ We found no significant differences between groups who experienced the baseline or improved network first. Thus we pool the six independent groups and base statistical tests on two-sided Wilcoxon signed-rank tests applied to six independent matched pairs. Throughout the remainder of the paper p-values are based on the same procedure, since we found no significant order effects in any of our treatments.

At the same time, Table 3 also shows that there are systematically more travelers on the low road than under the efficiency hypothesis for both Baseline (prediction: 2) and Improved (prediction: 3) in every group in both the short-run and the long-run. Thus, we can also reject the efficiency hypothesis (p-value = 0.031). Note, however, that in all six groups the average number of travelers on the low road is much closer to the equilibrium number than the efficient number.

Table 3: Traffic flow in PKD networks
- Average number of travelers on the low road -

| <i>Group</i> | PKD Baseline | | PKD Improved | |
|--------------|---------------------|---------------------|---------------------|---------------------|
| | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> |
| PKD.1 | 3.50 | 3.93 | 5.17 | 5.37 |
| PKD.2 | 3.57 | 3.87 | 5.20 | 5.73 |
| PKD.3 | 3.97 | 3.73 | 5.27 | 5.50 |
| PKD.4 | 3.60 | 3.87 | 5.30 | 5.60 |
| PKD.5 | 3.77 | 3.70 | 5.07 | 5.13 |
| PKD.6 | 3.93 | 3.93 | 5.40 | 5.57 |
| ALL | 3.72 | 3.84 | 5.23 | 5.48 |

The central prediction of the equilibrium hypothesis is that the traffic flows organize themselves in such a way that travel times across routes will be equalized. However, because the number of travelers on the congestible low route is systematically lower than in equilibrium the travel time equalization hypothesis is rejected. To study this in more detail, Table 4 presents the difference in average travel times between the high road and the low road for each group. The congestible low route is significantly faster on both networks, whether we focus on the short-run or long-run (a two-sided Wilcoxon signed-rank test yields a p-value of 0.031 in all cases). On the Improved network there is a trend in the direction of equalization; the differences are significantly smaller in the long run than short run ($p = 0.031$). On the Baseline network, where differences are generally smaller, the trend is insignificant ($p = 0.281$).

Table 4: Investigating travel time equalization in PKD networks
- Travel time on high road minus travel time on low road -

| <i>Group</i> | PKD Baseline | | PKD Improved | |
|--------------|---------------------|---------------------|---------------------|---------------------|
| | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> |
| PKD.1 | 4.50 | 0.60 | 7.50 | 5.70 |
| PKD.2 | 3.90 | 1.20 | 7.20 | 2.40 |
| PKD.3 | 3.60 | 1.20 | 6.30 | 3.60 |
| PKD.4 | 0.30 | 2.40 | 6.60 | 4.50 |
| PKD.5 | 2.10 | 2.70 | 8.40 | 7.80 |
| PKD.6 | 0.60 | 0.60 | 5.40 | 3.90 |
| ALL | 2.50 | 1.45 | 6.90 | 4.65 |

Since the low road is significantly faster, could a randomly selected subject have done better by simply choosing the low road all the time? To answer this question, we calculated for each subject what their travel time would have been had they always chosen the low road, holding all other subjects' choices fixed. Table 5 presents the averages of these hypothetical travel times. On the Baseline network the average subject would be worse off. The reason for this is that the subject's presence on the congestible road would increase traffic and therefore journey times on that road. For the PKD Improved network the average subject would have been better off in the short run (average journey time would have been reduced by about one point). However, because of the tendency in the direction of equalization, these "arbitrage" opportunities on the Improved network disappear in the long run, even though the congestible road is still faster.

Table 5: Actual versus hypothetical (low road) travel times on PKD networks
(Average across individuals, keeping behavior of other network users constant)

| | <i>Rounds 1-30</i> | | <i>Rounds 31-60</i> | |
|----------|---------------------------|-------------------------------------|---------------------------|-------------------------------------|
| | <i>Actual travel time</i> | <i>Travel time if always on low</i> | <i>Actual travel time</i> | <i>Travel time if always on low</i> |
| Baseline | 57.63 | 59.31 | 57.47 | 60.23 |
| Improved | 54.18 | 53.21 | 54.98 | 55.18 |

4.1.2 Network Performance

Table 6 displays the experienced travel time (or latency) for each group. Recall that both the efficiency and reduced form hypotheses predict that the improvement in the network will lead to a reduction in experienced travel times. In contrast, the equilibrium hypothesis predicts no change in network performance with the improvement. As Table 6 shows, experienced travel times were reduced after the improvement—even in the short run—in every group. Thus, there is a significant shift in the direction predicted by the efficiency and reduced form hypotheses (p-value = 0.031).

Table 6: Latency in PKD networks
- Average travel time for network users -

| <i>Group</i> | PKD Baseline | | PKD Improved | |
|--------------|---------------------|---------------------|---------------------|---------------------|
| | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> |
| PKD.1 | 57.56 | 58.35 | 53.44 | 54.49 |
| PKD.2 | 56.89 | 57.38 | 54.15 | 55.58 |
| PKD.3 | 57.38 | 57.53 | 53.74 | 55.05 |
| PKD.4 | 57.49 | 57.30 | 54.53 | 54.71 |
| PKD.5 | 57.94 | 56.36 | 54.53 | 54.38 |
| PKD.6 | 58.50 | 57.90 | 54.68 | 55.69 |
| ALL | 57.63 | 57.47 | 54.16 | 54.98 |

That being said, the observed latencies are closer to the equilibrium hypothesis than the alternatives. Table 6 shows that travel times are systematically higher than under efficiency for both Baseline (prediction: 52.5) and Improved (prediction: 46.875) in every group in both the short-run and the long-run. Thus, we can reject the level predictions of the efficiency hypothesis in every case (p-value = 0.031). In contrast, travel times are systematically higher than those predicted in equilibrium in the Baseline (prediction: 57) but systematically **lower** under the Improved network (prediction: 57). Again, we can reject the level predictions of the equilibrium hypothesis in both the short-run and the long-run, but the observed behavior is much closer to these levels.

Finally, to examine the level predictions of the reduced form hypothesis we need a benchmark from which to compare. For the groups that first encounter the Baseline and then the

Improved network we take the last 30 rounds in Baseline as a benchmark. Relative to that benchmark there is a 7% reduction in travel time in the short run (the first 30 rounds of the Improved network) and a 5% reduction in the long run (the second 30 rounds of the Improved network). Using an analogous procedure for the groups that first encounter the Improved network, we find that the deletion of the network improvement results in a 6% increase in travel time in the short run and 4% increase in the long run. Of course, this is vastly less than the 15% reduction (22% increase), which would have been forecast had one simply used reduced form estimates to predict changes in performance.¹⁰ Moreover, the actual changes in travel time appear to be systematic: regardless of the order in which networks were presented to subjects, travel time is lower in the Improved than in the Baseline network in all groups ($p = 0.031$).

Comparing the results of Table 6 with those in Section 4.1.1 reveals an apparent puzzle for the PKD Baseline network. As seen in Section 4.1.1, low road utilization in Baseline is less than that predicted in equilibrium, which implies that the low road is significantly faster than in equilibrium, whereas the travel time along the high road is of course exactly as predicted, since it is fixed regardless of the number of travelers. However, as seen in Table 6, the average experienced travel time *exceeds* the equilibrium prediction. The variance term in equation (6) explains this apparent puzzle. Despite efficiency enhancing under-utilization of the low road, the variance in the route choices raises the experienced travel time in the Baseline. That is, the “variance effect” dominates the mean effect in this network structure. In the Improved network, the variance effect is also present but the mean effect (under-utilization of the low road) dominates.¹¹ Intuitively, every time *many* travelers use the low road, the low road is slow and *many* travelers are adversely affected by this. On the other hand, every time *few* travelers use the low road, the low road is fast but *only those few* travelers benefit from this.

4.1.3 Summary

The improvement to the PKD network leads to a significant shift in traffic flows, in clear contrast to the reduced form heuristic. While the shift is in the direction consistent with equilibrium and efficiency, the point predictions of both the efficiency and equilibrium

¹⁰ For groups PKD.1, .2 and .3 the average travel time in the Improved network would have been 48.83 had they made the same choices as in the last 30 rounds of Baseline. For groups PKD.4, .5 and .6 the average travel time in Baseline would have been 67.32 had they made the same choices as in the last 30 rounds of Improved.

¹¹ In fact, variability is very similar on both networks. Averaging across groups, the standard deviation of n_i is 1.15 for the Baseline network and 1.13 for the Improved network.

hypotheses are rejected. In addition, network users are adversely affected by traffic flow fluctuations, which in Baseline cause the experienced travel times to rise above equilibrium levels despite the under-utilization of the congestible road. Furthermore, travel times are not fully equalized, in contrast to equilibrium predictions, and the PKD paradox—the complete dissipation of the benefits from the road improvement—is not borne out either in the short-run or the long-run. On the other hand, route choices and experienced travel times are closer to the predictions of equilibrium theory than the efficiency benchmark, and on average travelers could not have achieved a long-run reduction in travel time by always choosing the faster low road.

4.2 Braess Networks

4.2.1 Aggregate Route Choices and Travel Time per Route

Table 7 displays the average number of travelers choosing the low road for each group. The groups that first encounter the Baseline network and then the Improved network are listed first. Note that the Improved network can be viewed as the composition of two networks in each of which users make a binary choice between a congestible and a non-congestible road. In Table 7 we analyze these two components separately. Recall that the predicted equilibrium response to the improvement in the network is a shift in traffic onto the congestible roads. In contrast, efficiency requires a shift in traffic onto the non-congestible road in the first leg, and either no shift or a shift in traffic onto the non-congestible road in the second leg. As Table 7 shows, traffic shifted in the direction predicted by equilibrium. There is a significant increase in the amount of traffic leaving the origin on the congestible road and a significant increase in the amount of traffic arriving at the destination on the congestible road (the p-value is 0.031 in both cases and for both short and long run).

Again, as well as looking at directional effects, one can look at traffic flows relative to point predictions. On the Baseline network traffic flows do not differ significantly from the equilibrium prediction of 4 travelers on each road (p-value = 0.529 in both short run and long run). On the improved network there are systematic differences between equilibrium and actual traffic flows. For the first leg, equilibrium predicts just one traveler on the low road, and traffic is more evenly distributed across the two routes than this (p = 0.031 for both short and long run). For the second leg, equilibrium predicts a less extreme distribution of traffic (5 travelers on the

low road), but again there is a systematic tendency toward a more even distribution of traffic ($p = 0.031$ for the short run and $p = 0.063$ for the long run). Thus, the congestible roads are under-utilized on both legs of the Improved network.

Table 7: Traffic flow in Braess networks
- Average number of travelers on the low road -

| <i>Group</i> | Braess Baseline | | Braess Improved 1st Leg | | Braess Improved 2nd Leg | |
|--------------|------------------------|-------------------------|---|-------------------------|---|-------------------------|
| | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> |
| BRS.1 | 3.87 | 3.80 | 2.83 | 1.90 | 4.73 | 4.50 |
| BRS.2 | 3.77 | 4.07 | 2.30 | 1.47 | 4.33 | 4.13 |
| BRS.3 | 4.00 | 3.93 | 3.03 | 2.50 | 4.57 | 4.47 |
| BRS.4 | 4.13 | 4.20 | 2.77 | 2.23 | 4.60 | 4.90 |
| BRS.5 | 4.07 | 3.87 | 2.83 | 2.07 | 4.33 | 4.60 |
| BRS.6 | 3.87 | 3.83 | 2.73 | 1.63 | 4.53 | 5.10 |
| ALL | 3.95 | 3.95 | 2.75 | 1.97 | 4.52 | 4.62 |

Table 8 displays the difference in the average travel times between the high road and the low road. Since Baseline traffic flows do not differ significantly from the equilibrium prediction, we cannot reject the travel time equalization hypothesis. For the case of the Improved network, where fewer travelers use the congestible roads than predicted, these roads are significantly faster than the non-congestible roads, and the travel time equalization hypothesis is rejected. When we compare the travel time discrepancies in the short run and long run we observe a significant trend toward equalization on the first leg ($p = 0.031$), but not on the second leg ($p = 0.438$).

Table 8: Investigating travel time equalization in Braess networks*- Travel time on high road minus travel time on low road -*

| <i>Group</i> | Braess Baseline | | Braess Improved 1st Leg | | Braess Improved 2nd Leg | |
|--------------|------------------------|-------------------------|---|-------------------------|---|-------------------------|
| | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> | <i>Rounds 1-30</i> | <i>Rounds 31-60</i> |
| BRS.1 | 1.60 | 2.40 | -5.50 | -2.70 | 2.40 | 4.50 |
| BRS.2 | 2.80 | -0.80 | -3.90 | -1.40 | 6.00 | 7.80 |
| BRS.3 | -1.21 | 0.80 | -6.10 | -4.50 | 3.90 | 4.80 |
| BRS.4 | -1.60 | -2.40 | -5.30 | -3.70 | 3.60 | 0.90 |
| BRS.5 | -0.80 | 1.60 | -5.50 | -3.20 | 6.00 | 3.60 |
| BRS.6 | 1.60 | 2.00 | -5.20 | -1.90 | 4.20 | -0.90 |
| ALL | 0.40 | 0.60 | -5.25 | -2.90 | 4.35 | 3.45 |

Again, we ask whether the differences in travel times could have allowed subjects to increase their payoff if they had employed a different strategy. In Table 9 we conduct a similar analysis to that presented previously for the PKD networks. We ask whether a randomly selected subject could have benefited had she always chosen the road that was fastest overall, holding other subjects' choices fixed. For the Baseline and second leg of the Improved networks the difference in travel times is insufficient to provide easy opportunities to increase payoffs. The most interesting case is the first leg of the Improved network, where we observed the largest travel time discrepancies in the short run, and a significant tendency to reduce this discrepancy in the long run. Here, a randomly selected subject would have increased their short run payoff by 1.41 points and their long run payoff by 0.30 points had she always chosen the congestible road.

Table 9: Actual versus hypothetical* travel times on Braess networks*- Average across individuals, keeping behavior of other network users constant -*

| | Rounds 1-30 | | Rounds 31-60 | |
|--------------------------------|-------------------------------|-------------------------------------|-------------------------------|-------------------------------------|
| | <i>Actual travel time</i> | <i>Hypothetical travel time</i> | <i>Actual travel time</i> | <i>Hypothetical travel time</i> |
| Baseline | 59.29 | 61.11 | 58.59 | 61.11 |
| Improved - 1 st Leg | 18.19 | 16.78 | 19.14 | 18.84 |
| Improved - 2 nd Leg | 44.31 | 44.57 | 44.29 | 45.36 |

* Baseline and Improved (2nd Leg): we replace a subject's choices with the low road. Improved (1st Leg): we replace a subject's choices with the high road.

4.2.2 Network Performance

Table 10 displays latency on the Braess networks for each group. As with the PKD networks, the externalities road users place on one another generate inefficient travel flows. In the Baseline network, an efficient flow of traffic would generate an average travel time of 55.5, and actual travel times are 6% greater than this. Indeed, the efficiency loss is even greater than is predicted by equilibrium. For the Improved network average travel times are 62.97, lower than equilibrium predicts (67), but still 15% above efficient levels.

Table 10: Latency in Braess networks
- Average travel time for network users -

| <i>Group</i> | Braess Baseline | | Braess Improved | |
|--------------|------------------------------|-------------------------------|------------------------------|-------------------------------|
| | <i>Rounds</i> <i>1-30</i> | <i>Rounds</i> <i>31-60</i> | <i>Rounds</i> <i>1-30</i> | <i>Rounds</i> <i>31-60</i> |
| BRS.1 | 59.50 | 58.30 | 63.03 | 63.65 |
| BRS.2 | 58.65 | 59.20 | 62.33 | 61.58 |
| BRS.3 | 60.20 | 58.80 | 61.90 | 62.18 |
| BRS.4 | 60.10 | 59.00 | 63.00 | 64.03 |
| BRS.5 | 58.90 | 58.40 | 61.90 | 63.18 |
| BRS.6 | 58.40 | 57.85 | 62.88 | 65.95 |
| ALL | 59.29 | 58.59 | 62.50 | 63.43 |

In the Baseline network, even though traffic flows do not differ systematically from equilibrium, travel times are significantly higher than the equilibrium prediction ($p = 0.031$ for both the short run and long run). The reason for this seemingly paradoxical finding is again the adverse effect of variability in traffic flows (equation 6).

On the Improved network travel times are significantly lower than equilibrium predicts ($p = 0.031$ for both short run and long run). Although the variance effect is also present here, the overall outcome is dominated by the under-utilization of the congestible links relative to the equilibrium predictions.

While the additional link in the Improved network in principle allows more efficient traffic flows than are possible in the Baseline, equilibrium theory predicts an *increase* in travel times after the technical improvement. In the data, even though the performance of the Baseline network is worse than equilibrium predicts, and the performance of the Improved network is

better than equilibrium predicts, the table shows that all groups experience an increase in travel times after the improvement. Thus there is a significant increase in travel times in both the short run ($p = 0.031$) and the long run ($p = 0.031$), and the Braess Paradox is observed in our data.

To get an idea of the quantitative effect, we consider the groups that first encounter the Baseline network and then the Improved network, and again take the last 30 rounds in Baseline as a benchmark. Relative to that benchmark there is a 9% increase in travel time in both the short run (the first 30 rounds of the Improved network) and the long run (the second 30 rounds of the Improved network). For the groups that first encounter the Improved network, we find that eliminating the costless link results in a 15% reduction in travel time in the short run and a 17% reduction in the long run. Had one simply used the reduced form heuristic and assumed that the addition of a costless link would have no effect on travel flows, there would of course be no change in travel times.

4.2.3 Summary

In the Braess Baseline network average traffic flows do not diverge systematically from the equilibrium prediction (equal usage of each road) and there is no significant difference in travel times between the alternative routes. Nevertheless, variability in choices across rounds raises travel times above equilibrium levels. For the Improved network, equilibrium predicts a shift in traffic onto congestible links, and although we see a shift in this direction, it is not as marked as predicted. As a result, congestible links are under-utilized relative to equilibrium and travel times are not equalized in the Improved network. Overall, the Improved network performs better than predicted. Even so, traffic flows are far from efficient and we observe the Braess Paradox: the Improved network is significantly slower than the Baseline network.

4.3 Baseline Neutrality

4.3.1 Aggregate Route Choices

Our Braess Baseline network is derived from the PKD Baseline network by modifying the cost parameters on the high road so that it is no longer non-congestible, but in such a way as to preserve equilibrium travel flows and times. Since the previous sub-sections reported some systematic deviations from equilibrium, it is unclear whether this change in network architecture will indeed be neutral.

First we look at the distribution of travel choices. As seen in Tables 3 and 7 there is less traffic on the low road in the PKD Baseline than the Braess Baseline; however, the difference is only marginally significant in the short run (two-sided Wilcoxon rank-sum test $p = 0.078$) and insignificant in the long run ($p = 0.379$). For each group we also computed the standard deviation of the number of low-road users over the 30 rounds in each half of a session. A two-sided Wilcoxon rank-sum test fails to reject the null hypothesis that the variability in traffic flows is identical in PKD and Braess Baseline (p -value = 0.936 for rounds 1-30, p -value = 0.576 for rounds 31-60).

4.3.2 Network Performance

Second, we look at network performance. From equation (6) it can be shown that even if the distribution of choices is identical on the two networks, expected travel times will differ. Specifically, assuming a common expected number of low road users and a common variance in the number of low road users (here denoted by n), the difference in expected travel times is:

$$\text{Expected Travel Time}_{\text{Braess}} - \text{Expected Travel Time}_{\text{PKD}} = \frac{(8 - E[n])(12 - 3E[n]) + 3 \text{Var}(n)}{8}$$

In equilibrium $E[n] = 4$ and $\text{Var}[n] = 0$, and the network performances are the same. In the data, because the average number of low road users is slightly below 4 on both networks, and because the variability of the number of low road users is positive, this expression suggests that average journey times will be faster on the PKD network. Indeed this is the case: average travel times differ significantly across networks (p -value = 0.008 for rounds 1-30, p -value = 0.020 for rounds 31-60), and are higher on the Braess network.

5 Variability in Route Choice

As we have seen, an important determinant of network performance in the experiment is persistent variability in congestion. Even in the long-run, subject route choices continued to display considerable variability and, as a consequence, drove latency higher than would be the case had route choices converged. None of our existing hypotheses offers an explanation for variability in route choice, so, in this section, we explore two other alternative hypotheses. The

first of these supposes that the variability in route choices is the result of mixed strategy equilibrium play. As we will show, mixed strategy equilibria exist in all of our networks and, of course, these equilibria entail variability in route choice. The second hypothesis uses a learning model approach of Erev and Roth (1998) to generate variability in route choices as a result of the path of reinforcement received by a given subject.

5.1 Mixed Strategies

In addition to the equilibrium identified by equation (2), there exist mixed strategy equilibria in the networks we study. Specifically, we examine completely mixed equilibria.¹² Let n be the number of low road users, such that the journey time on the low and the high road are $\alpha_L + \beta_L n$ and $\alpha_H + \beta_H(N - n)$, respectively. Suppose now that player i ($i = 1, \dots, N$) chooses the low road with probability p_i . For this to constitute an equilibrium the expected travel time for player i must be the same on either route, taking into account her own presence on each route. Thus, for all i ,

$$\alpha_L + \beta_L \left(1 + \sum_{j \neq i} p_j \right) = \alpha_H + \beta_H \left(1 + \sum_{j \neq i} (1 - p_j) \right)$$

which can be rewritten as

$$\sum_{j \neq i} p_j = \frac{\alpha_H + \beta_H N - (\alpha_L + \beta_L)}{\beta_L + \beta_H}.$$

Summing over all i gives

$$(N-1) \sum_{i=1}^n p_i = \frac{\alpha_H + \beta_H N - (\alpha_L + \beta_L)}{\beta_L + \beta_H} N$$

and therefore

$$p_i = \frac{\alpha_H + \beta_H N - (\alpha_L + \beta_L)}{(N-1)(\beta_L + \beta_H)} \equiv p^* \text{ for all } i. \quad (7)$$

¹² At the end of section 5.2, we consider (and dismiss) asymmetric mixed strategy equilibria as a possible explanation.

Thus, equation (7) characterizes the unique symmetric mixed strategy equilibrium arising under each network. In such an equilibrium all players choose the low road with the same probability p^* , and this leads to a standard deviation in the number of travelers on each road of $\sqrt{Np^*(1-p^*)} > 0$. In principle, this could account for the observed variability in route choices.

Table 11 compares the symmetric mixed strategy equilibrium (sMSE) predictions with the experimental results using three metrics. The first is the usual route choice metric. The other two reflect predictions relating to variability in route choices. We computed the standard deviation of the number of travelers on the low road treating the number of travelers on the low road in a given period and a given network structure as an observation. The switching propensity looks at the fraction of times a given traveler changes his or her route choice in consecutive rounds for a given network structure. The statistics for each of these metrics are computed only for the long-run.

Table 11: Actual and predicted network choices under the mixed strategy hypothesis
- Data from last 30 rounds of a phase -

| | | PKD Baseline | PKD Improved | Braess Baseline | Braess Improved 1 | Braess Improved 2 |
|--------------------------------------|-------------|-------------------------|-------------------------|----------------------------|------------------------------|------------------------------|
| Avg. number on low road | <i>sMSE</i> | 3.43 | 5.71 | 3.71 | 1.14 | 4.57 |
| | <i>Data</i> | 3.84 | 5.48 | 3.95 | 1.97 | 4.62 |
| Stdev. of no. on low road | <i>sMSE</i> | 1.40 | 1.28 | 1.41 | 0.99 | 1.40 |
| | <i>Data</i> | 1.02 | 1.02 | 1.08 | 0.93 | 1.07 |
| Switching Propensity | <i>sMSE</i> | 49% | 41% | 50% | 25% | 49% |
| | <i>Data</i> | 18% | 21% | 27% | 20% | 25% |

As Table 11 shows, the symmetric mixed strategy equilibrium correctly predicts the average under-utilization, relative to our earlier theoretical predictions, of the low road in PKD, Braess Baseline, and Braess Improved 2, as well as the relative over-utilization of the low road in Braess Improved 1. However, the predictions are too extreme in this regard in the two Baseline treatments ($p = 0.031$ in both cases) and insufficient in PKD Improved ($p = 0.059$) and Braess Improved 1 ($p = 0.031$); Braess Improved 2 is the only case where the point prediction of the

symmetric mixed strategy equilibrium is not rejected. In all other networks, the average number of low road users is closer to four than predicted by the sMSE solution.

Turning to the standard deviation metric, while the symmetric mixed strategy equilibrium predicts persisting variability in route choices, its point predictions are too high compared to the variability observed in the data (p-value = 0.031 in PKD Baseline, Braess Baseline and Braess Improved 2; p-value = 0.142 in PKD Improved and Braess Improved 1). The third metric, switching propensity, illustrates this effect even more starkly. The mixed strategy hypothesis implies far more switching behavior than we actually observe.

Finally, note that with one exception the mixed strategy predictions regarding network latency are the same as the predictions of the pure strategy equilibrium. Thus, here we observe the same discrepancies as before: higher-than-predicted network latency in PKD Baseline and lower-than predicted latency in PKD Improved and Braess Improved. The exception is the Braess Baseline network, where the mixed strategy hypothesis predicts longer travel times than the pure strategy equilibrium. Although travel times observed in the data indeed systematically exceed the pure strategy prediction, they remain significantly below 59.25, the expected travel time under the mixed strategy hypothesis (p-value = 0.031). Thus, the mixed strategy equilibrium prediction is again too extreme.

Summary: While the mixed strategy hypothesis can account for the presence of persistent variability in congestion and for the direction of other observed deviations from the pure strategy equilibrium, it predicts too much variability and too frequent route switching relative to the data, and does not improve the accuracy of point predictions with respect to route choices and network latency.

5.2 Learning

What causes the recurring, systematic discrepancies between the predictions of the mixed strategy hypothesis and the patterns observed in data? The key difficulty under any homogenous probabilistic choice model is that the more “balanced” the road usage (i.e. the closer the average number of travelers on the low road is to 4), the higher is the predicted variability of road usage. The data, however, displays a tendency towards a more balanced road usage than predicted by the mixed strategy hypothesis *and* a more concentrated distribution of road usage than predicted

by the mixed strategy hypothesis. This can be explained by heterogeneity of play among travelers.

One way in which such heterogeneity could arise is as a result of path-dependent learning trajectories. To illustrate this, we ran simulations based on a simple reinforcement learning algorithm. We assume that in round 1 each player i has an initial propensity $q_{ij}(1) > 0$ to play her j th pure strategy (which in Braess Improved is “high-high”, “high-low”, “low-high” or “low-low”, and in the other networks is simply “high” or “low”). In each round, player i then picks strategy k with probability

$$p_{ik}(t) = \frac{q_{ik}(t)}{\sum_j q_{ij}(t)} .$$

Players update their propensities to play strategies simply by adding the payoff arising from the chosen strategy to the propensity for that strategy. Thus, if player i plays her k th strategy in round t and obtains a payoff $\pi_i(t)$, then the propensity to play her j th strategy in round $t + 1$ is

$$q_{ij}(t) = \begin{cases} q_{ij} + \pi_i(t) & \text{if } j = k \\ q_{ij} & \text{else} \end{cases} .$$

In our experiment subjects were not informed about network payoff structures. They only knew that they would receive 100 points for completing a journey and that there would be travel costs, which would be deducted from, but never exceed, the 100 points. Since subjects begin the experiment without any hints about which strategy might be better, we restrict initial propensities to be identical across strategies: $q_{ij}(1) = q_i(1)$ for all j . For a given player this setup describes the “basic” reinforcement learning model in Erev and Roth (1998). It has one parameter, the value of $q_i(1)$, which influences the speed of the adjustment process: agents make greater adjustments to the probability of choosing a particular strategy, the higher the experienced payoff is relative to current propensity levels. The parameter may thus be interpreted as reflecting expectations about how much one might earn in this game. Since players begin our experiment without knowledge

of payoff structures we let $q_i(1)$ be determined randomly for each player according to a uniform distribution on the interval $[1, 100]$.

Each simulation ran for 60 periods on a particular network, and we ran 1000 simulations for each network. Table 12 summarizes our findings (averaging over the 1000 runs) using the same metrics as described above.

Table 12: Actual and simulated network choices under reinforcement learning
- Data from last 30 rounds of a phase -

| | | PKD Baseline | PKD Improved | Braess Baseline | Braess Improved 1 | Braess Improved 2 |
|--------------------------------------|-------------------|-------------------------|-------------------------|----------------------------|------------------------------|------------------------------|
| Avg. number on low road | <i>Simulation</i> | 3.64 | 4.66 | 3.79 | 3.42 | 4.24 |
| | <i>Data</i> | 3.84 | 5.48 | 3.95 | 1.97 | 4.62 |
| Stdev. of no. on low road | <i>Simulation</i> | 1.19 | 1.21 | 1.15 | 1.33 | 1.28 |
| | <i>Data</i> | 1.02 | 1.02 | 1.08 | 0.93 | 1.07 |
| Switching Propensity | <i>Simulation</i> | 30% | 28% | 30% | 35% | 35% |
| | <i>Data</i> | 20% | 25% | 30% | 20% | 25% |
| Avg. network latency | <i>Simulation</i> | 57.10 | 51.60 | 58.40 | 60.74 | |
| | <i>Data</i> | 57.50 | 55.00 | 58.60 | 63.43 | |

First, the simulation reliably produces a more balanced distribution of travelers on the high and the low road than the mixed strategy hypothesis. Second, with one exception (Braess Improved 1st leg), we find that the standard deviation in the simulated networks is lower than predicted by the sMSE but not close to zero, again the same type of deviation from theory that we observe in the data.¹³ Third, the switching propensities, while still higher than those in the data, are much closer than the mixed strategy hypothesis. Finally, the discrepancies between the theoretically expected network latency and the average latency arising in the simulation are always in the same direction as observed in the experimental data. For instance, it is interesting to see that, just like in the experimental data, travel times are slightly greater than in equilibrium in PKD Baseline and lower than in equilibrium in PKD Improved. The reason for this—in both

¹³ The reason that the standard deviation is relatively high in Braess Improved 1st Leg is that the simulated propensity adjustments in this segment of the network are comparatively slow. This is because travel costs on the second leg (up to 72) tend to be much higher than those on the first leg (up to 24), and this means that the learning process in the model is more sensitive to costs incurred on the second leg.

the experimental and the simulation data—lies in the combined effect of, relative to the mixed strategy hypothesis, lower variability in road usage and more balanced traffic flows. While the lower variability reduces latency in both networks relative to the sMSE, the greater balance in road usage implies *more* traffic on the congestible road in PKD Baseline, thereby increasing network latency (and overturning the benefits from lower variability), and implies *less* traffic on the congestible road in PKD Improved, thereby reducing network latency even further. Similar effects are also observed in the Braess networks.

There is further evidence that path-dependent learning trajectories may play an important role in the dynamics of choices. For example, subjects display a tendency to “lock in” to certain strategies in all networks. As before, we consider the second-half data of each condition. Within this time frame there is no systematic trend in the overall propensities to choose the low road: the propensities in the third quarter (rounds 31-45) do not differ significantly from the propensities in the last quarter (rounds 46-60). However, both subjects who choose the low road more often than the high road in the third quarter and subjects who choose the high road more often are disproportionately more likely to again choose their “preferred” road more often in the fourth quarter. For example, in the third as well as the fourth quarter of PKD Baseline roughly half the subjects (48%) chose the low road more often than the high road. However, 87% of those who display a tendency for the low road in the third quarter display the same tendency in the last quarter, and 84% of those who choose the high road more often in the third quarter also do so in the last quarter. In equilibrium, where players make independent random draws in each round, this should not occur: based on a time frame of 15 rounds the expected proportion of players displaying a tendency towards the low road is 29% (high road 71%), regardless of the history of play. In the simulations, on the other hand, we find a pattern very similar to that observed in the data: about 45% of our simulated agents display a tendency towards the low road in the third quarter, and 85% of these (and 88% of those with a tendency towards the high road) then also display a tendency towards the low road (high road) in the last quarter. This “lock in” effect is not specific to PKD Baseline: we get qualitatively similar results in all our networks, both in the experimental data and in the simulations.

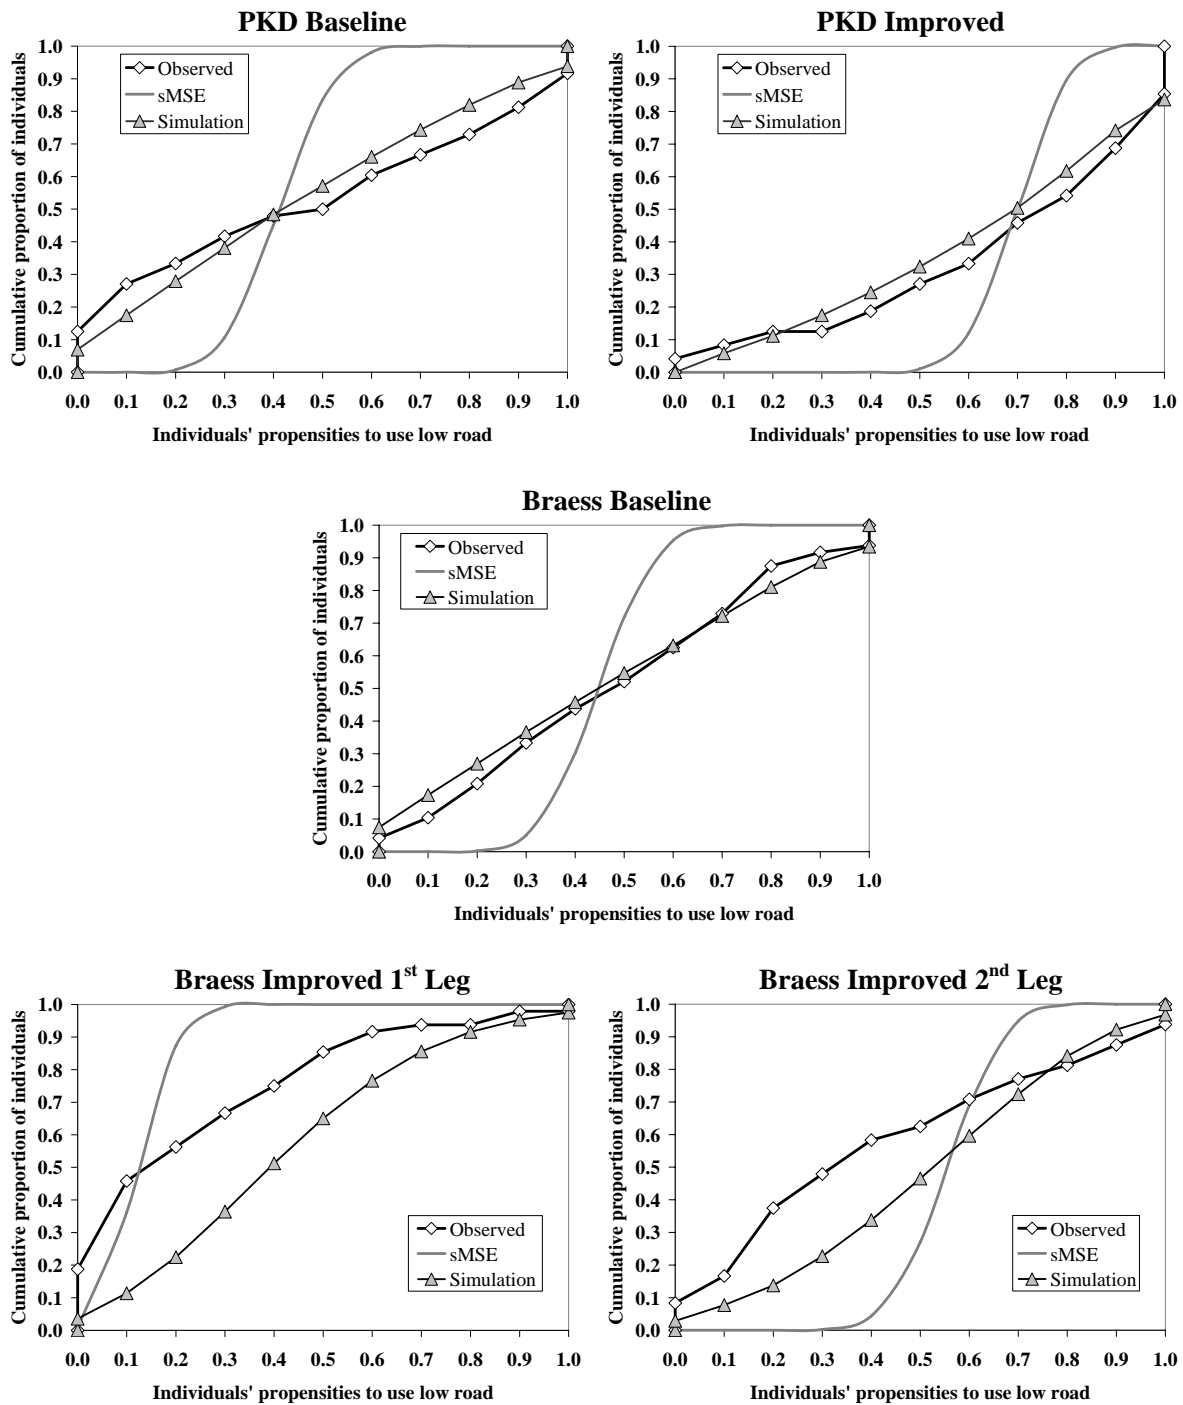
It should be pointed out, though, that the learning model does not provide systematically “better” predictions than the equilibrium hypothesis. Comparing Table 2 to Table 12, notice that the predictions of the equilibrium hypothesis concerning route choices are often closer to the

empirical averages than the results of the simulation. Of course, one could fit parameter values from the data, in which case simulated behavior will more closely approximate the data. Indeed, one could also improve fit between simulated behavior and data by considering more complex learning models with additional parameters. However, our point here is simply to illustrate how broad patterns observed in the data may result from reinforcement-type learning, rather than to provide a fully-fledged model of the data generation process.

Finally, an alternative approach is to consider asymmetric mixed strategy equilibria. In such an equilibrium, some players take their preferred route with certainty while others mix over route choices. These equilibria give scope for some heterogeneity amongst players and, as a result, would for example imply lower switching levels than in the symmetric mixed strategy equilibrium--in some cases similar to those we find in the data. But note that although several of these asymmetric mixed-strategy equilibria exist for each of our networks, the degree of heterogeneity they allow is limited, because all players that randomize between routes have to do so with the *same* probability.¹⁴ Thus, in these equilibria there are at most three distinct “types” of travelers: constant low road users, constant high road users, and switchers. In contrast, actual heterogeneity between subjects is much more diffused than that. To show this, we counted the proportion of times a subject used the low road using data from the last thirty rounds on a network. Cumulative histograms of these proportions are displayed in Figure 9. (For comparison we have also plotted the corresponding distributions that emerge in the symmetric equilibrium and in the simulations.) Almost all subjects switch between roads but there is substantial variability across subjects in the propensity to use the low road, which in some of the networks leads to nearly uniform distributions of traveler types. This pattern is not consistent with any of the asymmetric mixed strategy equilibria.

¹⁴ Consider two players, i and j , who both mix between routes and let \hat{n} be the expected number of low road users excluding i and j . Then, to make player i indifferent between routes, player j must choose the low road with probability p_j so that it solves $\alpha + \beta(\hat{n} + p_j + 1) = \gamma + \delta(2 + N - \hat{n} - p_j)$. However, to make player j indifferent at the same time, player i must also choose the low road with that probability.

Figure 9: Cumulative distributions of players' propensities to use low road
- Data from last 30 rounds of a phase -



5. Conclusion

Equilibrium theory offers important principles for network design. We identify two such principles; the least congestible route principle states that improvements should be made on the route *least* sensitive to congestion, and the size principle states that adding costless links reduces travel time if and only if there are a sufficiently large number of travelers on the network. Two paradoxes from the transportation literature are closely related to these principles. The Pigou-Knight-Downs Paradox can be viewed as an extreme violation of the least congestible route principle: here the network has a non-congestible route, and an improvement to an alternative, congestible, route has no effect on equilibrium journey times. The Braess Paradox is a manifestation of the size principle: with a sufficiently small number of travelers on a network adding a costless link increases equilibrium journey times.

We then conduct an experiment to study how technical improvements to networks affect travel costs. In our experiment subjects repeatedly made route choices, receiving a reward for reaching their destination and incurring travel costs which could depend on the number of other subjects taking that route. Subjects were only informed of travel costs and congestion associated with each journey at the end of that journey; they were not told how travel times were determined on alternative routes.

In our PKD sessions we improve a congestible route on a network that also has a non-congestible route. Here the Pigou-Knight-Downs Paradox is predicted: equilibrium travel times are the same with and without the improvement. In our experiment we observe that, as predicted by equilibrium theory, the improvement causes network users to switch from non-congestible to congestible roads, eroding efficiency gains, particularly in the long run after subjects had acquired experience with the new network structure. However, even in the long run the dissipation of the gains is not quite complete, and travel times are 4% lower in the improved network.

In our Braess sessions we add a new road to a network so that the Braess Paradox is predicted: equilibrium travel times are higher with the additional road. Here we found that travel times were indeed higher on the “improved” network. Focusing on the long run, we find that the network ‘improvement’ raises travel times by about 8%.

For all our treatments we observe adjustments in traffic flows in response to exogenous changes in network architecture. Often these adjustments are in the opposite direction to that

commensurate with attaining efficiency; indeed, we see no evidence that subjects are able to coordinate and attain efficient outcomes. Thus, travelers are unable to overcome the inefficiencies associated with externalities on congestible roads. One feature of the network games studied here that may make efficient flows difficult to achieve is that they require inequitable travel times across individuals. Traffic flows that minimize average travel times mean longer travel times for some travelers than others, and without a mechanism for sharing the benefits of efficient traffic patterns, individuals may not be willing to suffer long travel times while others enjoy short travel times. In our networks a focus on private travel costs is perhaps even more marked because we only give subjects feedback on own travel times.

In contrast, there is a reliable tendency for distributions of route choices to shift in the direction predicted by equilibrium. Thus, our experiment reinforces the theoretical point that accounting for induced changes in behavior is a crucial element of successful network design, and provides evidence of the usefulness of equilibrium theory for predicting the direction in which traffic flows will change. Nevertheless, some systematic departures from equilibrium predictions are observed, and variability in route choice behavior is observed, even in the long run. A contributing factor to this may be that equilibrium in these networks implies indifference (or near indifference) between alternative routes, and so noisy traffic flows around the equilibrium may occur because each subject can cheaply deviate from equilibrium. We note that this feature of our network games is shared by a vast array of models used in economics, where equilibrium involves equalization of returns from alternative choices. As our experiment demonstrates, this variability in route choices has systematic implications for network performance.

In order to further investigate deviations from equilibrium we conducted simulations based on a simple reinforcement learning model. Although, like the equilibrium model, this learning model does not give an exact explanation of subject choices, it generates similar patterns in traffic flows to those observed in our experimental data. For example, the higher reinforcement of more successful strategies tends to promote usage of routes that offer higher payoffs, and so generates adjustments in the direction of travel time equalization. At the same time, the model generates heterogeneity in individual strategies and variability in congestion over time. Thus, learning plays a role in both equilibrating travel times and generating deviations from equilibrium. Adjustments in the direction of travel time equalization and variability in

congestion are also key features of our data, and they have important impacts on network performance. For instance, in the PKD network the congestible road is under-utilized relative to equilibrium. On the improved network this results in better network performance than in equilibrium, and over time subjects shift toward the faster congestible route, reducing the disparity between travel times on the two routes. The unimproved network performs worse than in equilibrium, since the adverse effect of variability in traffic flows outweighs the effect of under-utilization of the congestible route.

These results demonstrate both the force and limitations of equilibrium incentives. While attempts to minimize own travel costs undermine network performance, they do not lead to complete convergence to equilibrium, even in the relatively simple traffic networks examined in the laboratory. Moreover, even when the deviation from equilibrium involves, on average, less travel on congested routes, this does not improve network performance because the accompanying variability in route choices, as subjects persistently strive to find the best route, provides an offsetting source of network inefficiency.

References

- Arnott, R. and K.A. Small (1994). The economics of traffic congestion. *American Scientist* **82** Sept.-Oct., 446-455.
- Braess, D. (1968). Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* **12**, 258-268.
- Ellison, G., D. Fudenberg and M. Möbius (2004). Competing Auctions. *Journal of the European Economic Association* **2(1)**, 30-66.
- Erev, I. and A.E. Roth (1998). Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria. *American Economic Review* **88**, 848-881.
- Fischbacher, U. and C. Thöni (forthcoming). Inefficient Excess Entry in an Experimental Winner-Take-All Market, *Journal of Economic Behavior and Organization*.
- Goeree, J.K. and C.A. Holt (forthcoming). An Explanation of Anomalous Behavior in Models of Political Participation, *American Political Science Review*.
- Kahneman, D. (1987). Experimental Economics: A Psychological Perspective. In R. Tietz, W. Albers and R. Selten (Eds.), *Modeling Bounded Rationality*, 11-20.
- McKelvey, R.D. and T.R. Palfrey (1995). Quantal Response Equilibria for Normal Form Games, *Games and Economic Behavior* **10**, 6-38.
- Rapoport, A., Kugler, T., Dugar, S. and E. Gisches (2005a). Choice of Routes in Congested Traffic Networks: Experimental Tests of the Braess Paradox. *Mimeo*, University of Arizona.
- Rapoport, A., Kugler, T., Dugar, S. and E. Gisches (2005b). Braess Paradox in the Laboratory: An Experimental Study of Route Choice in Traffic Networks with Asymmetric Costs. *Mimeo*, University of Arizona.
- Selten, R., T. Chmura, T. Pitz, S. Kube and M. Schreckenberg (forthcoming). Commuters route choice behaviour. *Games and Economic Behavior*.

Appendix A. Instructions

General Rules

This is an experiment in the economics of decision making. You will be paid in private and in cash at the end of the experiment. The amount you earn will depend on your decisions, so please follow the instructions carefully.

There are sixteen people participating in this experiment. All participants are reading the same instructions as you are and have been recruited in the same fashion. Likewise, all participants are participating in this experiment for the first time, as you are.


It is important that you do not talk, or in any way try to communicate, with the other participants during the experiment. If you have a question, raise your hand and a monitor will come over to where you are sitting and answer your question in private.

The experiment will consist of two phases of 60 rounds. In every round of the experiment you will be matched with seven other participants – the same seven others in every round.

In each round you will have an opportunity to earn points. At the end of the experiment, you will be paid an amount in cash that will be determined by the total number of points you earn from all rounds.

Description of a Round

At the beginning of each round you will see a roadmap on your screen. An example is given below (the ones you will see in the experiment will be different):



The map has cities and roads. The cities are denoted by letters, the roads by the lines connecting them. You begin at city O (origin). In each round you will receive 100 points for making a trip to city D (destination). There are a variety of routes to D.

You will have 40 seconds to choose a route from O to D by clicking on roads between pairs of cities. If you click on a road between two cities its colour will change from black to red. (You can change it back by clicking again). When you have created a complete route from O to D you will be given an opportunity to confirm your route. If you do not complete a trip to D in 40 seconds you will earn zero points for that round.

You will also pay a cost of travelling along each road you use, the cost varying from road to road. On some roads the cost will depend on how many other participants choose to travel along that road. The total cost of travel from city O to city D will always be less than 100 points regardless of the route travelled.

When all participants have completed their trips, or after 40 seconds have passed, the round will end. The computer will then show you the cost you incurred from each road you used, the number of participants using those roads, and your total travel cost. The computer will also show you your point earnings for the round: 100 points for reaching D minus the travel costs, or zero points if you did not complete a route to D.

For the 60 rounds of phase one, the way in which travel costs are determined will remain the same. Likewise, for the 60 rounds of phase two, the way in which travel costs are determined will remain the same. However the costs in phases one and two may differ. In between phases there will be a screen reminding you that phase one has ended and phase two is about to begin.

Ending the Experiment

At the end of the last round you will be informed of your total point earnings. This will be the sum of your point earnings from the 120 rounds. A monitor will come to your desk and pay you in private and in cash. You will receive 1 penny for every 5 points earned.