

Network Architecture and Traffic Flows: Experiments on the Pigou-Knight-Downs and Braess Paradoxes

by

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

July 2007

Abstract

This paper presents theory and experiments to investigate how network architecture influences route-choice behavior. We consider changes to networks that, theoretically, exhibit the Pigou-Knight-Downs and Braess Paradoxes. We show that these paradoxes are specific examples of more general classes of network change properties that we term the “least congestible route” and “size” principles, respectively. We find that technical improvements to networks induce adjustments in traffic flows. In the case of network changes based on the Pigou-Knight-Downs Paradox, these adjustments undermine short-term payoff improvements. In the case of network changes based on the Braess Paradox, these adjustments reinforce the counter-intuitive, but theoretically predicted, effect of reducing payoffs to network users. Although aggregate traffic flows are close to equilibrium levels, we see some systematic deviations from equilibrium. We show that the qualitative features of these discrepancies can be accounted for by a simple reinforcement learning model.

Acknowledgements

We are grateful for the extremely helpful comments of three anonymous referees as well as the Action Editor. We also thank seminar participants at the University of Leicester, the Institute for Transport Studies, University of Leeds, and at the Amsterdam 2004 International Meeting of the Economic Science Association for their 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

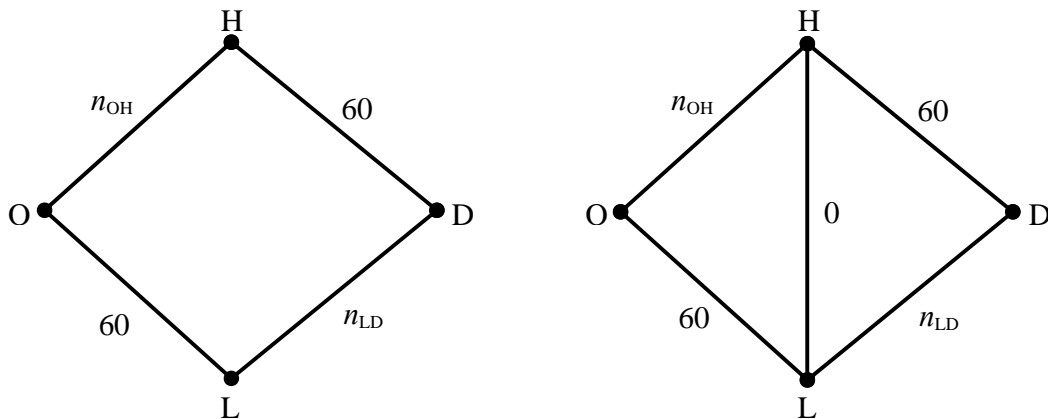
Economic agents are often presented with situations where they must decide how to get information, materials, or simply themselves from point A to point B. In such situations the agents often interact on congested networks, and an individual's travel cost will typically depend on the decisions made by other network users. The externalities that network users impose on each other have important consequences for network performance and design. One consequence is that traffic flows may not be socially efficient, as each user attempts to minimize his own travel cost without taking into account the effect of his actions on others. A second consequence, of practical importance to providers and managers of networks, is that technical improvements to a network may have unintended consequences, as they can induce behavioral changes that, while sensible for each individual, negatively affect overall network performance.

This paper investigates how network architecture influences route-choice behavior by comparing outcomes across different networks. In the theory section, we present two principles for planners. In what we term the **least congestible route principle** we show, employing the notion of Wardrop equilibrium (Wardrop, 1952, Beckmann, McGuire and Winsten, 1956), that the planner most effectively reduces travel time by first improving the route that is *least* sensitive to network congestion. This principle takes an extreme form when one of the routes along the network is non-congestible: Imagine a population of individuals who can choose between two routes. One route takes one hour, regardless of the number of people using it, while 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 will ensure that the travel time 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 are equalized—the travel time remains at one hour even after the improvement. In this case, the benefits from improving the route are *completely* dissipated.¹ This observation is the essence of the Pigou-Knight-Downs Paradox (Downs, 1962), which 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).

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.

Another example of how network architecture influences route choice behavior in counterintuitive ways is the Braess Paradox (Braess, 1968). A version of this 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 and the travel times along each link where, for example, n_{LD} represents the number of travelers using the L-D link. 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. 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. As a result, travel times increase to 120 minutes—strategic effects in response to network improvements leave all travelers worse off.

Figure 1: The Braess Paradox



If, however, the population, N , grows sufficiently large (i.e. $N > 120$) the efficiency of the additional link overcomes the adverse strategic effects, because, while each user will still take 120 minutes with the link, her travel time will be $60 + N/2$ minutes without. This reversal illustrates the **size principle**—the efficiency gains from adding network links always outweigh losses from user externalities once the number of network users is sufficiently large.

While the two 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. Alternatives to equilibrium models, based on adaptive models of route choice behavior often imply significant differences between choices and equilibrium, even in the long run. For example, Horowitz (1984) studies several models of adaptive route-choice behavior on simple two-link networks and finds that traffic flows sometimes converge to values

that differ substantially from equilibrium. Moreover, Bazzan and Klügl (2003) study a network where, in theory, the Braess paradox applies, and report simulations in which artificial agents employing a simple heuristic algorithm to make route choices are able to avoid the adverse effects of the paradox.

Examining the behavioral relevance of equilibrium in network environments is notoriously difficult. As Horowitz (1984) notes, the substantial discrepancies that are often found between measured traffic volumes and those computed with equilibration models could reflect model misspecification rather than failure of equilibrium. 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 equilibrium predictions. As a result, a number of researchers have recently turned to controlled laboratory experiments to examine route-choice behavior. In these experiments subjects make choices between routes and receive financial rewards that are linked to the routes chosen.²

Along these same lines, we study the behavioral relevance of the paradoxes identified above by varying the network architecture using controlled laboratory experiments. Unlike many existing studies which use a full information design, we do not inform subjects of the travel cost functions of all links in the network. Rather, as is commonplace in transportation networks, our subjects get feedback about the travel cost of journeys they undertake as well as the number of other travelers sharing each link in the journey. 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, even in this limited information setting, 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

² We will discuss some of these experimental studies in the next section. A number of other experimental studies from the transportation literature use a rather different methodology in which subjects respond to descriptions of hypothetical route choice scenarios, but are not paid according to their decisions. See, for example, Avineri and Prashker (2006), Iida, Akiyama and Uchida (1992), as well as the experiments reviewed in Mahmassani and Srinivasan (2004).

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 magnitude of actual adjustments is closer to equilibrium than to efficient route utilization. Improvements to the congestible route on the Pigou-Knight-Downs network lead to a small reduction in travel times, and an even smaller reduction when we focus on travel choices in later rounds after subjects had time to gain experience with the network. For the Braess networks, equilibrium entails a shift of traffic in the direction we observe, whereas efficiency entails 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 choices in later rounds. Thus, the Braess Paradox is observed in our laboratory setting.

Although aggregate traffic flows are close to equilibrium levels, we see some systematic deviations from equilibrium. 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, which would result in lower travel times all else equal, the adverse effect of variability in traffic flows outweighs the mean effect, and travel times are, in fact, even higher than in equilibrium.

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 across routes than is consistent with equilibrium, and this too is an implication of the learning model.

Related Experimental Literature

Selten, Chmura, Pitz, Kube and Schreckenberg (2007) conduct a route-choice experiment where subjects repeatedly choose between two congestible routes on a fixed network. The information structure of their experiment is similar to ours: Subjects are not informed of the travel cost functions associated with each route. They find that although aggregate choices are close to equilibrium proportions, traffic flows fluctuate around equilibrium rather than converging to it.

Providing subjects with post-round information on the travel cost of the *non-chosen* route dampens, but does not eliminate, traffic flow fluctuations. Helbing (2004) replicates and extends the Selten *et al.* setup. He finds that by providing subjects with additional feedback either in the form of potential payoffs from alternative routes or individual recommendations as to route choices further reduces the volatility of traffic flows. Unlike our study, neither of these studies varies network architecture.³

As far as we are aware, the Pigou-Knight-Downs Paradox has not been previously studied in the lab. However, the structure of the underlying game somewhat resembles a market entry game where the profits of entrants decline with the number of entrants. In the first experimental study of market entry games, Kahneman (1988) found that the number of entrants was 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. These results are broadly consistent with ours; in our experiment aggregate choices are close to equilibrium, and where we see deviations from equilibrium, they tend to involve a more even allocation of subjects to the two routes.

The Braess Paradox has recently been studied in experiments by Rapoport, Kugler, Dugar and Gisches (2005, forthcoming). There are some procedural differences between these experiments and ours. Rapoport *et al.* use a full information design, larger groups, and different network parameters.⁴ Their results are similar to ours even though the information conditions of subjects are quite different. Like us, they find that subjects' route choices change in the direction predicted by equilibrium when changes in the network occur, that the paradox is observed in the laboratory, that mean traffic flows are close to equilibrium (in particular in the long run), and that behavior is not affected by the order in which the baseline and the improved networks are presented. Similar to our findings, they also observe persistent fluctuations in traffic flows and route switching, as well as considerable heterogeneity in behavior across individuals.

The remainder of the paper proceeds as follows. In section 2 we present the background theory, and introduce the least congestible route and size principles. Our experimental design and

³ Schneider and Weimann (2004) and Ziegelmeyer *et al.* (forthcoming) also report experiments on traffic networks. However, their focus is on bottleneck models where subjects choose departure times.

⁴ Rapoport, Mak and Zwick (2006) also use a full information design to study the effects of variation in the number of travelers on an augmented Braess network.

procedures are described in section 3, and the results in section 4. In section 5 we explore alternative explanations of the variability in traffic flows we find in the data. Section 6 concludes.

2 Theory

Consider a network where N individual travelers wish to get from some origin, O , to a destination, D , to which there are k possible routes. All travelers face the same decision and make their choices 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 $T_i = \alpha_i + \beta_i n_i$ where $\alpha_i, \beta_i > 0$ are parameters of the model and n_i denotes the number of travelers using route i . The free-flow parameter α_i should be interpreted as the fixed time factor to traverse a road in the absence of any congestion. The congestion 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 traveler than are low β_i roads.

A standard framework for analyzing equilibrium traffic flows in network problems, and the one we adopt in this section, is the notion of Wardrop Equilibrium (Wardrop, 1952; Beckman *et al.*, 1956). Such an equilibrium consists of a vector of the number of travelers using each of the k routes (n_1, n_2, \dots, n_k) such that the travel times on any used route (i.e. where $n_i > 0$) are less than or equal to the travel time on any other route. When the vector (n_1, n_2, \dots, n_k) is integer valued, then a Wardrop equilibrium is also a Nash equilibrium. When it is not, a Wardrop equilibrium corresponds to the notion of quasi-equilibrium as defined in Ellison, Fudenberg and Möbius (2004). For large N , Haurie and Marcotte (1985) show that Nash equilibrium converges to Wardrop equilibrium.

We will restrict attention to the case where all routes are used in equilibrium. Denoting the least congestible road as route 1 (i.e. $\beta_1 \leq \beta_i$ for all i), the following *route viability conditions* are sufficient for all routes to be used in equilibrium:

$\alpha_i < \alpha_1$ for all $i \neq 1$, and

$$N > \sum_{i=1}^k \frac{\alpha_1 - \alpha_i}{\beta_i} .$$

The first set of conditions ensures that no route is weakly dominated by route 1 (i.e. for no route i is it the case that $\beta_1 \leq \beta_i$ and $\alpha_1 \leq \alpha_i$). The latter condition merely ensures that there are a sufficient number of travelers so that the least congestible route is used.

A Wardrop equilibrium is a solution to the system of equations

$$T = \alpha_i + \beta_i n_i \quad \text{for all } i$$

$$\sum_{i=1}^k n_i = N.$$

Rewriting the representative equation above for all i , we obtain $n_i = (T - \alpha_i)/\beta_i$. Summing over all i and solving for T then yields the equilibrium travel time of

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

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 (where an improvement means a reduction in either the free-flow parameter or the congestion parameter). Which route should be improved?

A common intuition suggests that the optimal strategy is to improve the road that is most heavily used. Underlying this intuition is the following argument: Since total travel time is $\sum_{i=1}^k n_i T_i$, if traffic flows do not adjust in response to changes in the network (i.e. reductions in α_j or β_j) then the marginal reduction in travel time is greatest for the most traveled route since $\partial \left(\sum_{i=1}^k n_i T_i \right) / \partial \alpha_j = n_j$ and $\partial \left(\sum_{i=1}^k n_i T_i \right) / \partial \beta_j = n_j^2$. Of course, missing from this intuition is the idea that travelers respond (at all) to changes in the network.

A slightly more sophisticated 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. Much like the

first approach, this intuition also fails to account for equilibrium readjustment of travel patterns following a change in network structure. Surprisingly, this approach is 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

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

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 .⁵

When $\beta_1 \rightarrow 0$, the *Pigou-Knight-Downs Paradox* emerges as a special case of the least congestible route principle. From equation (1), using L'Hospital'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.

A modification of the least congestible route principle is required for technical improvements embodied in changes to the congestion parameter β_j . Differentiating equilibrium travel times with respect to β_j gives

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

⁵ 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.

Note that the least congestible route principle does not hold *exactly* for this type of change because this expression involves both the volume of traffic using a route and the congestion parameter for that route. However, for any two equally traveled routes, the better route to improve is the one that is less congestible.

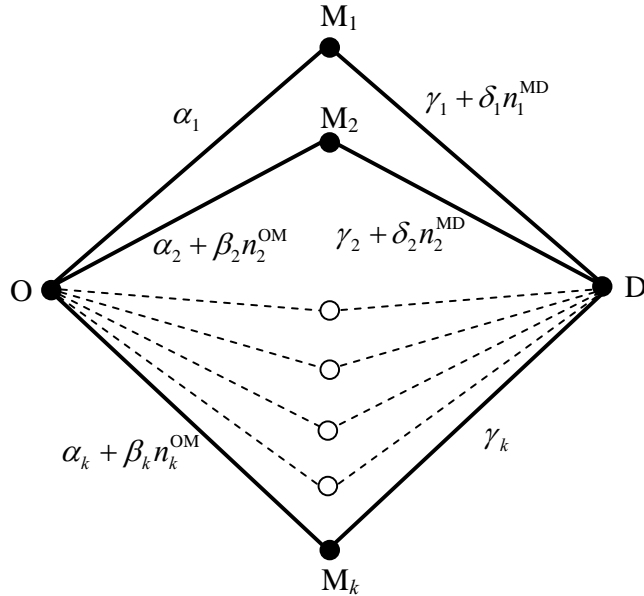
More Complicated Network Structures

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 before, the time cost for link i is linear:

$T_i = \alpha_i + \beta_i n_i^{\text{OM}}$. 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 $T_j = \gamma_j + \delta_j n_j^{\text{MD}}$.

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.⁶ Figure 2 shows a representative network with this structure.

Figure 2: A network with k midpoints



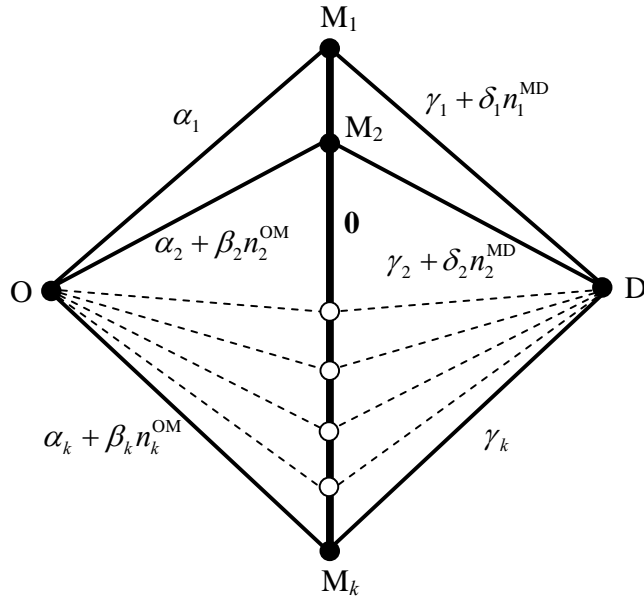
⁶ Formally, we study the limit as $\beta_1, \delta_k \rightarrow 0$.

Notice that the original network is nested as a special case of this network. Again, 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

$$T^* = \left(N + \sum \frac{\alpha_i + \gamma_i}{\beta_i + \delta_i} \right) / \left(\sum \frac{1}{\beta_i + \delta_i} \right).$$

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 . With 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.

Figure 3: Costless switching between the k midpoints



We offer the following:

The Size Principle: Adding costless links reduces travel time if there are sufficiently many travelers.

The implied 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. 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 when N is sufficiently large. When N is small and the route viability condition holds for the improved network, the externality effect dominates—the *Braess Paradox* is a simple illustration of this. Of course, when N becomes very small, the route viability condition for the improved network will no longer be satisfied. In particular, for N sufficiently small, the noncongestible links will no longer be used in equilibrium in the improved network, and the efficiency effect will again dominate. Returning to the example in the introduction, notice that, if the number of travelers was reduced from 60 to 20, then the Braess “improvement” would *reduce* travel costs from 70 minutes to 40 minutes. The reason is that, when there are sufficiently few travelers, the noncongestible links cease to be viable in the improved network.⁷

Previous theoretical research has also highlighted versions of the size principle. Specifically, Pas and Principio (1997) as well as Penchina (1997) provide versions of the principle for the case of an “anti-symmetric” network where $k = 2$.⁸ The size principle is also implicit in the complete characterization of the $k = 2$ case offered by Frank (1981).⁹

⁷ We thank the action editor for pointing this out.

⁸ In our notation, an anti-symmetric network is one with the restriction that, for all i , $\alpha_i = \gamma_{k+1-i}$ and $\beta_i = \delta_{k+1-i}$.

⁹ Kameda, Altman, Kozawa and Hosokawa (2000) study Braess-like Paradox results occurring in a Nash equilibrium of a $k = 2$ network for parameter values where the Wardrop equilibrium does not yield Braess outcomes.

3 Experiment

While the previous section derived properties of varying network architectures using Wardrop 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 Wardrop equilibrium number of travelers on each route is an integer.¹⁰

Pigou-Knight-Downs (PKD) Treatments

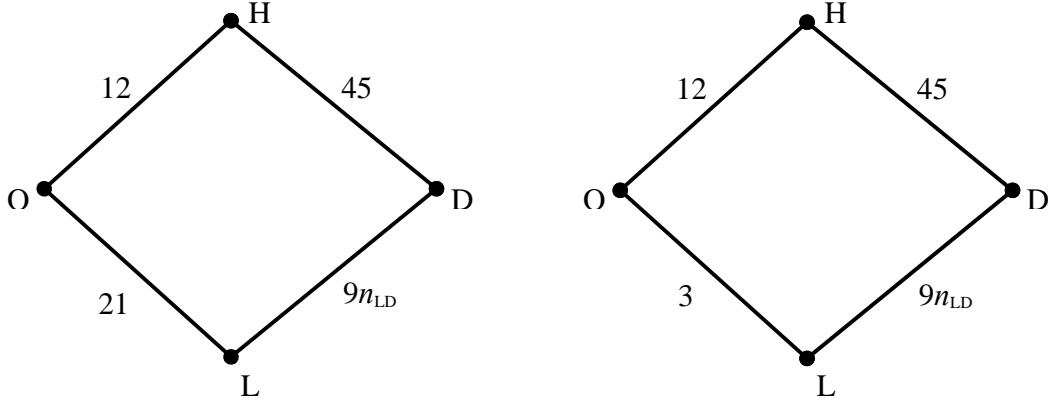
Figure 4 displays the first two networks used in our experiment. For the road between L and D the cost is nine times the number of travelers using that road. For the other roads the costs are fixed. We will refer to the route O-H-D as the “high road” and the route O-L-D as the “low road”. In the PKD 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.

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. Equilibrium travel time is determined by travel time on the high road: $T^* = 57$. When four travelers take the high road and four the low road travel time is equalized across routes in the PKD Baseline network.¹¹

¹⁰ If instead of Wardrop equilibrium, we use the notion of Nash equilibrium, then, in addition to pure strategy Nash equilibria, there are also mixed-strategy Nash equilibrium solutions for our networks. We discuss these in detail in Section 5.

¹¹ There is an additional pure strategy Nash 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 Wardrop 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 Nash equilibrium corresponds to the Wardrop equilibrium.

Figure 4: PKD networks implemented in the experiment



Turning to the PKD Improved network, equilibrium travel time is again $T^* = 57$. When two travelers take the high road and six take the low road travel time is equalized across the two routes. 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.

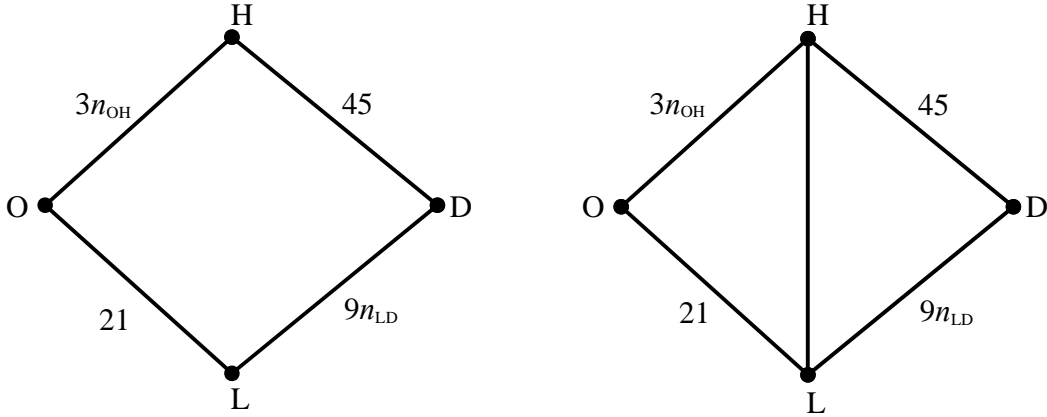
Table 1: Efficient and equilibrium traffic on PKD networks

		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

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 take the high road, giving an average travel time of 55.5. In contrast, 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 (right panel). By permitting traffic to be diverted between the high and low roads, an improved allocation of traffic on the network is possible and the average travel time decreases (slightly) 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. Thus, the overall travel time on the network increases to 66 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	66

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.¹² 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.¹³

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.

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 incurred from *each* link that they used, the number of subjects using those links, their own total travel cost and point earnings. Our rationale for this design decision was the following. First, we wanted to reproduce a scenario where network users learn about the properties of the network through their own experiences. Specifically, our design has the natural feature that, 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. Second, we wanted to examine the practical relevance of the least cost and size principles. One limitation of this design decision is

¹² A copy of the instructions is included as Appendix A.

¹³ At the time of the experiment the exchange rate was approximately £1 = \$1.86.

that, strictly speaking, the experimental design is disconnected from a formal game-theoretic model of route choice decision making. Thus, one must be cautious in drawing conclusions based on comparisons between the (Nash) equilibrium of such a game theoretic model and results from experiments where the assumptions of the model are not met. That being said, to the extent one views equilibrium analysis as offering guidance in practical settings where not all of the model assumptions are likely to be met, it seems to us sensible to take equilibrium predictions as a benchmark against which to compare our experimental results.

By providing a subject with feedback about both the cost incurred and the volume of traffic sharing each link taken, a subject might learn the cost structure of the network. Specifically, any subject could, in principle, calculate a linearization of the cost function for each link taken so long as there was any variation in the volume of traffic on a given link. Since the cost functions for each link were, in fact, linear, this calculation would yield an exact estimate of the cost function of a given link.

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

We use several different performance metrics to compare the implications of Wardrop equilibrium with two alternative benchmarks. 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 at all. 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 benchmark,” 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 equilibrium and efficiency benchmarks with respect to route choice. While the reduced form heuristic offers no level prediction for route choice, it does predict that route choices are unresponsive 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. Wardrop equilibrium implies that that travel time per route will be equalized across routes. In contrast, efficiency requires unequal travel times across routes. 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 over a given time horizon (i.e. the short-run or the long-run). As we will see, the statistical properties of this measure are of some interest. For the PKD networks, given some distribution of route link choices by individuals, the expectation of the travel time experienced by 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 latency depends on the variance of route choices as well as the average traffic flows over each route.¹⁴ Moreover, equation (6) implies that higher variability of route choices leads to longer experienced travel times. The intuition is that every time *many* travelers use a congestible road, that road is slow and *many* travelers are adversely affected. On the other

¹⁴ While the explicit calculation of equation (6) was for the PKD networks, an analogous calculation reveals a variance term for the Braess networks as well.

hand, every time *few* travelers use a congestible road, that road is fast but *only those few* travelers benefit from the under-utilization. Tables 1 and 2 display the equilibrium and efficiency benchmarks with respect to the latency metric. The reduced form heuristic 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.¹⁵ 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 implied by equilibrium and efficiency (p-value = 0.031).¹⁶

Table 3 also shows that there are systematically fewer travelers on the low road than in Wardrop equilibrium for both Baseline (equilibrium: 4) and Improved (equilibrium: 6) in every group in both the short-run and the long-run. Thus, there are significant differences between observed and equilibrium travel flows (p-value = 0.031). At the same time, Table 3 shows that there are systematically more travelers on the low road than implied by efficiency for both Baseline (efficient: 2) and Improved (efficient: 3) in every group in both the short-run and the long-run. Thus, there are significant differences between observed and efficient travel flows (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.

¹⁵ On the Braess networks in only 3 of 5,760 subject-rounds did a subject fail to complete a route.

¹⁶ 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.

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

A key characteristic of Wardrop equilibrium is that travel times across routes are equalized. However, because the number of travelers on the congestible low route is systematically lower than in equilibrium we can reject the “travel time equalization hypothesis”. 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

4.1.2 Network Performance

Table 5 displays the experienced travel time (or latency) for each group. Recall that, according to both the efficiency and reduced form benchmarks, the improvement in the network will lead to a reduction in experienced travel times. In contrast, in equilibrium the improvement should have no impact on network performance. As Table 5 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 implied by the efficiency and reduced form benchmarks (p-value = 0.031).

Table 5: 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

However, latencies systematically deviate from both the efficiency and the equilibrium benchmarks. Table 5 shows that travel times are systematically higher than under efficiency for both Baseline (efficient: 52.5) and Improved (efficient: 46.875) in every group in both the short-run and the long-run (p-value = 0.031 in every case). Travel times are systematically higher than those implied by equilibrium in the Baseline (equilibrium: 57) and then systematically *lower* under the Improved network (equilibrium: 57). Having said this, the observed latencies are closer to the equilibrium benchmark than the efficient benchmarks in all groups in both the short run and the long run.

Finally, to examine the reduced form heuristic 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. 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.¹⁷ 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.

Comparing the results of Table 5 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 significantly less than in equilibrium, which implies that the low road is significantly faster than in equilibrium, whereas the travel time along the high road is the same as in equilibrium because it is fixed regardless of the number of travelers. However, as seen in Table 5, the average experienced travel time *exceeds* the equilibrium level. 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.¹⁸

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, there are significant deviations from the efficiency and equilibrium benchmarks. 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, 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

¹⁷ 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.

¹⁸ In fact, variability is very similar on the two networks. Averaging across groups, the standard deviation of n_i is 1.15 for the Baseline and 1.13 for the Improved network. Such traffic flow fluctuations are of course not compatible with Wardrop equilibrium. In contrast, in a mixed-strategy Nash equilibrium one *would* expect variability in traffic flows. In Section 5 we investigate the explanatory power of mixed-strategy solutions in more detail.

experienced travel times are closer to the equilibrium 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 6 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 6 we analyze these two components separately. Recall that equilibrium requires a shift in traffic onto the congestible roads after the improvement in the network. 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 6 shows, traffic shifted in the direction of 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).

Table 6: 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

On the Baseline network traffic flows are not significantly different from equilibrium flows 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. In equilibrium, one traveler uses the low road on the first leg, and five use the low road on the second leg. In fact, traffic is more evenly distributed across the two routes on both legs (first leg: $p = 0.031$ for both short and long run, second leg: $p = 0.031$ for the short run and $p = 0.063$ for the long run). Thus, relative to equilibrium, the congestible roads are under-utilized on both legs of the Improved network.

Table 7 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 equilibrium, we cannot reject the travel time equalization hypothesis. For the case of the Improved network, where the congestible roads are under-utilized relative to equilibrium, 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 7: 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

4.2.2 Network Performance

Table 8 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, these efficiency losses exceed those that would be

incurred if the system were in equilibrium. For the Improved network average travel times are 62.97, better than in equilibrium (66), but still 15% above efficient levels.

Table 8: Latency in Braess networks
- Average travel time for network users -

<i>Group</i>	Braess Baseline		Braess Improved	
	<i>Rounds 1-30</i>	<i>Rounds 31-60</i>	<i>Rounds 1-30</i>	<i>Rounds 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 in equilibrium ($p = 0.031$ 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 in equilibrium ($p = 0.031$ 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 equilibrium.

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

4.2.3 Summary

In the Braess Baseline network average traffic flows do not diverge systematically from equilibrium (equal usage of each road) and there is no significant difference in travel times between the alternative routes. Nevertheless, travel times are higher than in equilibrium, due to variability in choices across rounds. For the Improved network, equilibrium entails a shift in traffic onto congestible links. Although we see a shift in this direction, congestible links are under-utilized relative to equilibrium, and the Improved network performs better than in equilibrium. 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 6 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 traffic flows. 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 benchmarks offer an explanation for variability in route choice, so, in this section, we explore two other alternatives. 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 alternative uses a learning model approach 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 Wardrop equilibrium, there exist mixed strategy equilibria in the networks we study. Here we examine completely mixed equilibria.¹⁹ Let n be the number of low road users,

¹⁹ At the end of section 5.2, we consider the consistency between our results and asymmetric mixed strategy equilibria.

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

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

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

Table 9 uses three metrics to compare our experimental results with what would be expected in a symmetric mixed strategy equilibrium (sMSE). The first is the usual route choice metric. The other two metrics relate 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 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.

As Table 9 shows, average under-utilization, relative to the Wardrop equilibrium, of the low road in PKD, Braess Baseline, and Braess Improved 2, and the high road in Braess Improved 1 is exactly what we would expect in a symmetric mixed strategy equilibrium. In Braess

Improved 2, the average number of low road users does not differ significantly from the number we would expect in sMSE. However, in all other cases there are significant differences between the data and sMSE traffic flows, and these differences go in the direction of a more even distribution of traffic across the high and low roads (PKD Baseline: $p = 0.031$, PKD Improved: $p = 0.059$, Braess Baseline: $p = 0.031$, Braess Improved 1: $p = 0.031$).

Table 9: Comparison of observed and symmetric mixed strategy equilibrium traffic flows
- 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%

With respect to standard deviations in route choice, we observe less variability in route choices than would be expected in sMSE (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 even more starkly. The mixed strategy equilibrium implies far more switching behavior than we actually observe.

Finally, note that, for the PKD networks and Braess Improved network, latency is the same in either the Wardrop equilibrium or a sMSE (in either case latency is pinned down by the travel cost associated with non-congestible links). Thus, we observe the same discrepancies as before: Latency is above equilibrium levels in PKD Baseline and below equilibrium levels in PKD Improved and Braess Improved. In the Braess Baseline network, sMSE travel times are longer than Wardrop equilibrium travel times. Although observed travel times systematically exceed the Wardrop equilibrium level, they remain significantly below 59.25, the level expected in a symmetric mixed strategy equilibrium (p -value = 0.031). Thus, observed latencies are not consistent with symmetric mixed strategy equilibrium.

Summary: Persistent variability in congestion and the direction of other observed deviations from Wardrop equilibrium are consistent with a symmetric mixed strategy equilibrium.

However, we observe far less variability and route switching than one would expect in a symmetric mixed strategy equilibrium. Route choices and network latency are also significantly different from what one would expect in a symmetric mixed strategy equilibrium.

5.2 Learning

The discrepancies between mixed strategy equilibrium and the patterns observed in the data are of a fundamental nature: In any homogenous probabilistic choice model, the more “balanced” the road usage (i.e. the closer the probability of choosing the low road is to 0.5), the higher is the variability of road usage. However, in the data we find a tendency towards a more balanced road usage *and* a more concentrated distribution of road usage (relative to the sMSE). This indicates that there is heterogeneity of play among the travelers in our experiment.

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”). Players then update their propensities to play strategies 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

$$q_{ij}(t+1) = \begin{cases} q_{ij}(t) + \pi_i(t) & \text{if } j = k \\ q_{ij}(t) & \text{else.} \end{cases}$$

This setup thus far is identical to the “basic” reinforcement learning model in Erev and Roth (1998), which has only one parameter, namely the value of $q_{ij}(1)$. In our experiment subjects were not informed about network payoff structures. At the beginning of a session 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 there are no hints about which strategies might be best, we restrict initial propensities to be identical across strategies: $q_{ij}(1) = q_i(1)$ for all j . However, note that the initial propensities may still vary across individual agents. The parameter $q_i(1)$ 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.

Our model differs from Erev and Roth's basic reinforcement learning model with respect to the stochastic response rule: In our model, player i picks strategy k in round t with probability

$$p_{ik}(t) = \frac{e^{\beta_i q_{ik}(t)}}{\sum_j e^{\beta_i q_{ij}(t)}}$$

where the sensitivity parameter β_i determines the degree to which differences in the propensities $q_{ij}(t)$ across strategies lead an individual to discriminate between strategies with high and low propensities. Such a probabilistic response would arise if each individual chose the strategy with the maximum propensity score in a given period, but where the propensity scores are subject to a random component drawn from an extreme value distribution (see McFadden, 1974).

In order to find suitable values for the two parameters $q_i(1)$ and β_i , which must be specified individually for each network user i , we estimated them from our data using the time series we have collected from each subject. To do this we searched (using a numeric grid search algorithm) for combinations of $q_i(1)$ and β_i that minimize the sum of the squared deviations between a subject's actual choices and the probabilities of making these choices under the learning model, or

$$\sum_t \sum_j (d_{ij}(t) - p_{ij}(t))^2,$$

where $d_{ij}(t)$ is a dummy variable taking the value 1 when subject i chose strategy j in round t and 0 if she did not.

The parameter estimates vary greatly across individuals. The $q_i(1)$ parameter ranges from 1 to 100, with an average of 81.7 and a standard deviation of 23.0. The β_i parameter ranges from 0 to 0.4, with an average of 0.038 and a standard deviation of 0.110.

For each of the twelve groups of eight players in our experiment, we then ran computer simulations based on the learning algorithm described above and using the parameters we estimated from the eight human subjects in that group. Each simulation ran for 60 periods on the

relevant network, and we conducted 1000 runs of simulations for each eight-player group.²⁰ Table 10 summarizes our findings (averaging over the 1000 runs) using the same metrics as before.

Table 10: Comparison of observed and simulated traffic flows
- 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>Simulation</i>	3.62	4.62	3.82	3.39	4.29
	<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>Simulation</i>	1.16	1.17	1.20	1.30	1.17
	<i>Data</i>	1.02	1.02	1.08	0.93	1.07
Switching Propensity	<i>sMSE</i>	49%	41%	50%	25%	49%
	<i>Simulation</i>	24%	20%	30%	23%	23%
	<i>Data</i>	18%	21%	27%	20%	25%
Avg. network latency	<i>sMSE</i>	57.0	57.0	59.3	66.0	
	<i>Simulation</i>	57.1	51.4	58.7	60.6	
	<i>Data</i>	57.5	55.0	58.6	63.4	

The heterogeneity across the simulated players' learning model parameters produces, relative to the sMSE, a more balanced distribution of travelers on the high and the low road *and* less variability in route choices, both in terms of the standard deviation of the number of road users²¹ and in terms of individuals' switching propensities. At the same time, the variability in traffic flows is not close to zero, as would be the case in Wardrop equilibrium. Thus, the simulation results are characterized by the same types of deviations from equilibrium that we observe in the data.

²⁰ Note that even with a very large number of rounds, the learning model could produce data that deviates substantially from equilibrium, since the process is strongly path dependent. That is players may 'lock-in' on strategies, even if the resulting pure strategy profile does not constitute an equilibrium.

²¹ Braess Improved 1st leg is the only exception. 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.

Similarly, when we turn to the network latency metric we find that the direction of the discrepancies between equilibrium and simulation data is the same as between equilibrium and experimental data. For instance, the simulation produces travel times in the PKD Baseline treatment that are slightly greater than in equilibrium and travel times in the PKD Improved treatment that are lower than in equilibrium. The reason for this—in both the experiment and the simulation—lies in the combined effect of lower variability in road usage and more balanced traffic flows. While the lower variability reduces latency in both networks relative to the mixed strategy equilibrium, 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, the “lock in” effect is even stronger than in the experimental data: 48% of our simulated agents display a tendency towards the low road in the third quarter, and 98% of these (and also 98% 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.

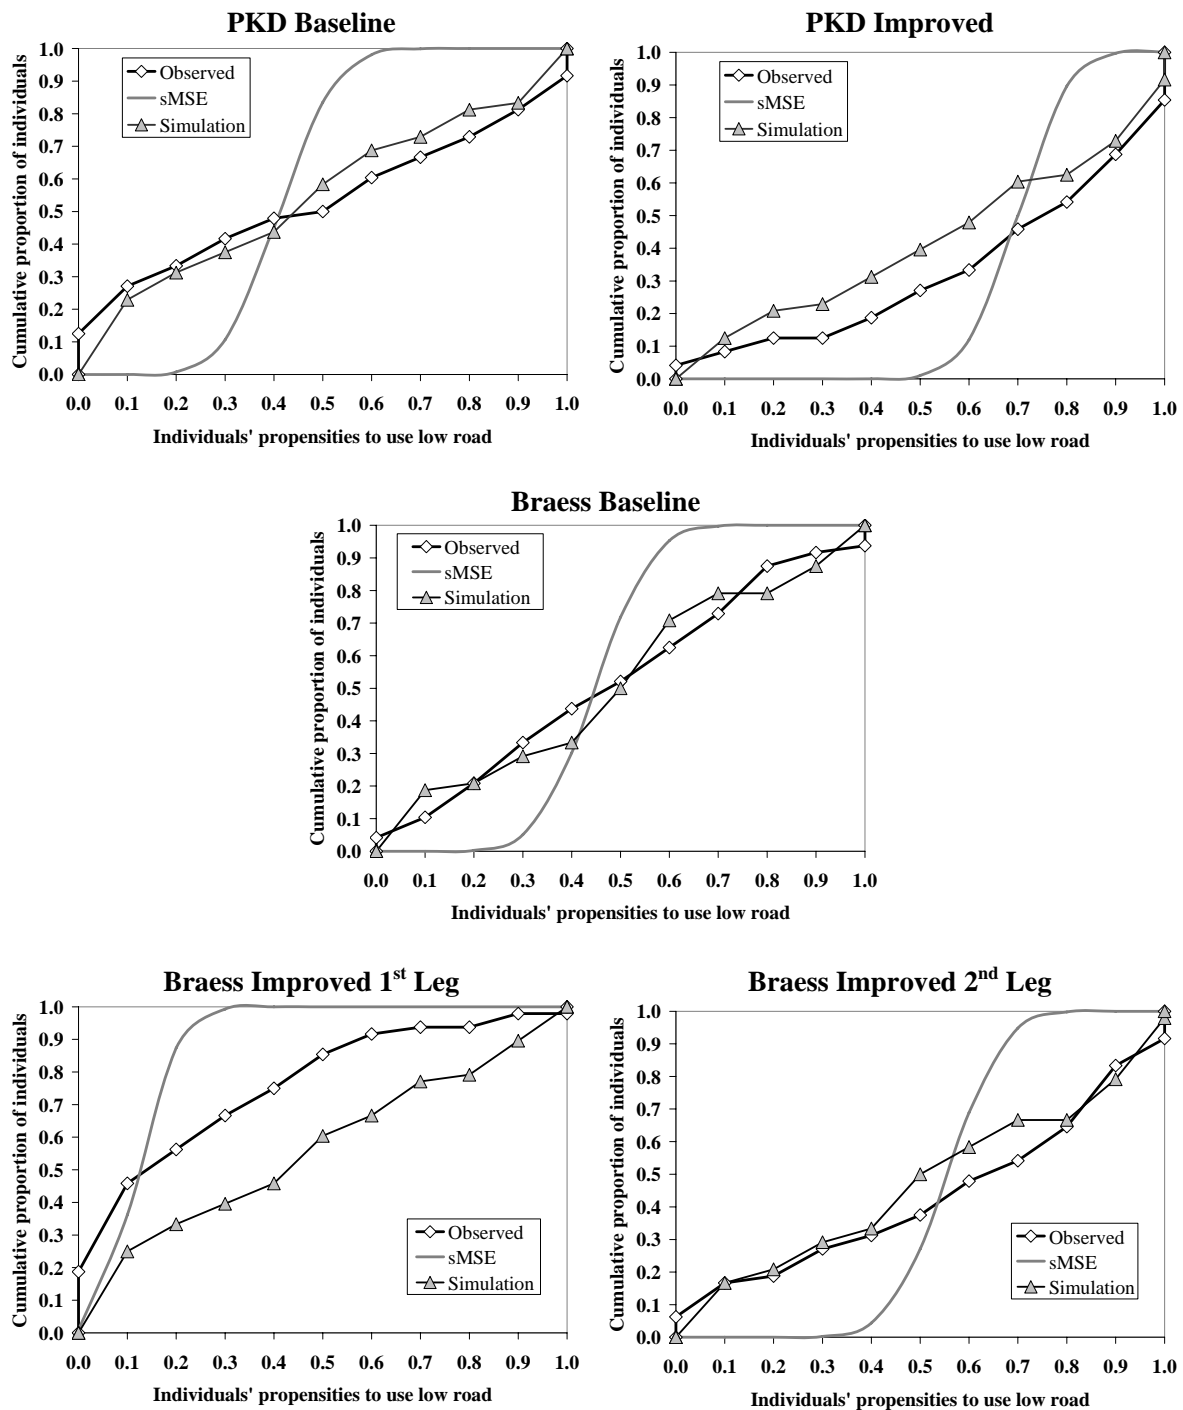
However, there are also discrepancies between the outcomes of the learning model and the experimental data and, as seen in Table 10, these are not systematically smaller than the discrepancies between equilibrium and experimental data. An interesting question is whether this is because of shortcomings in the detailed assumptions underlying the learning model (i.e. the specifications of the stochastic response and propensity updating rules) or is because this class of learning models *in general* cannot fully capture the dynamic route-choice patterns in our experimental data. We do not have a conclusive answer to this question—however, we considered several alternative formulations of the stochastic response rule and the propensity updating rule and found that the simulation results reported above are quite robust. Of course, one could potentially improve the 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, we briefly compare our data with the implications of asymmetric mixed strategy equilibria. In such an equilibrium, some players take one of the routes with certainty while others mix over route choices, which gives scope for some heterogeneity amongst players. As a result, switching levels are lower than in the symmetric mixed strategy equilibrium and in some cases similar to those we find in the data. But note that although several of these asymmetric 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.²² 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 6. (For comparison we have also plotted the corresponding distributions that emerge in the symmetric equilibrium and in the simulations.) Almost all subjects switch

²² 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.

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.

Figure 6: Cumulative distributions of players' propensities to use low road
- Data from last 30 rounds of a phase -



6. Conclusion

The notion of Wardrop equilibrium 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 when 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 equilibrium traffic flows imply the Pigou-Knight-Downs Paradox: Equilibrium travel times are the same with and without the improvement. In our experiment 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 equilibrium implies the Braess Paradox: 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 of 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 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, 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.
- Avineri, E. and J.N. Prashker (2006). The Impact of Travel Time Information on Travelers' Learning under Uncertainty. *Transportation* **33**, 393-408.
- Bazzan, A.L.C. and F. Klügl (2003). Learning to Behave Socially and Avoid the Braess Paradox in a Commuting Scenario. *Mimeo*, Universidade Federal do Rio Grande do Sul.
- Beckmann, M.J., C.B. McGuire and C.B. Winsten (1956). *Studies in the Economics of Transportation*. Yale University Press: New Haven, Connecticut.
- Braess, D. (1968). Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* **12**, 258-268.
- Camerer (2003). *Behavioral Game Theory*. Princeton University Press: Princeton, New Jersey.
- Downs, A. (1962). The Law of Peak-Hour Expressway Congestion. *Traffic Quarterly* **16**, 393-409.
- 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.
- Frank, M. (1981). The Braess Paradox. *Mathematical Programming* **20**, 283-302.
- Haurie, A. and P. Marcotte (1985). On the Relationship between Nash-Cournot and Wardrop Equilibria. *Networks* **15**, 295-308.
- Helbing, D. (2004). Dynamic Decision Behavior and Optimal Guidance through Information Services: Models and Experiments. In: M. Schreckenberg and R. Selten (eds.), *Human Behaviour and Traffic Networks*, Springer: Berlin, 47-95.
- Horowitz, J.L. (1984). The Stability of Stochastic Equilibrium in a Two-Link Transportation Network. *Transportation Research B* **18(1)**, 13-28.
- Iida, Y., T. Akiyama and T. Uchida (1992). Experimental Analysis of Dynamic Route Choice Behavior. *Transportation Research B* **26(1)**, 17-32.

- Kahneman, D. (1988). Experimental Economics: A Psychological Perspective. In: R. Tietz, W. Albers and R. Selten (eds.), *Bounded Rational Behavior in Experimental Games and Markets*, Springer: Berlin, 11-18.
- Kameda, H., E. Altman, T. Kozawa and Y. Hosokawa (2000). Braess-Like Paradoxes in Distributed Computer Systems, *IEEE Transactions on Automatic Control* **45(9)**, 1687-1691.
- Mahmassani, H.S. and K.K. Srinivasan (2004). Experiments with Route and Departure Time Choices of Commuters Under Real-Time Information: Heuristics and Adjustment Processes, in: M. Schreckenberg and R. Selten (eds.), *Human Behaviour and Traffic Networks*, Springer: Berlin, 97-131.
- McFadden, D. (1974). Conditional Logit Analysis of Qualitative Choice Behavior. In: P. Zarembka (ed.), *Frontiers in Econometrics*, Academic Press: New York, 105-142.
- Pas, E.I. and S.L. Principio (1997). Braess' Paradox: Some New Insights, *Transportation Research B* **31(3)**, 265-276.
- Penchina, C.M. (1997). Braess Paradox: Maximum Penalty in a Minimal Critical Network, *Transportation Research A* **31(5)**, 379-388.
- Rapoport, A., T. Kugler, S. Dugar and E. Gisches (2005). Choice of Routes in Congested Traffic Networks: Experimental Tests of the Braess Paradox. *Mimeo*, University of Arizona.
- Rapoport, A., T. Kugler, S. Dugar and E. Gisches (forthcoming). Braess Paradox in the Laboratory: An Experimental Study of Route Choice in Traffic Networks with Asymmetric Costs. In: T. Kugler, J.C. Smith, T. Connolly and Y.J. Son (eds.), *Decision Modeling and Behavior in Uncertain and Complex Environments*. Springer: Berlin.
- Rapoport, A., V. Mak and R. Zwick (2006). Navigating Congested Networks with Variable Demand: Experimental Evidence. *Journal of Economic Psychology* **27**, 648-666.
- Schneider, K. and J. Weimann (2004). Against all Odds: Nash Equilibria in a Road Pricing Experiment. In: M. Schreckenberg and R. Selten (eds.), *Human Behaviour and Traffic Networks*, Springer: Berlin, 133-153.
- Selten, R., T. Chmura, T. Pitz, S. Kube and M. Schreckenberg (2007). Commuters Route Choice Behaviour. *Games and Economic Behavior* **58**, 394-406.
- Wardrop, J.G. (1952). Some theoretical aspects of road traffic research, *Proceedings of the Institute of Civil Engineers Part II* **1**, 325-378.

Zieglmeyer, A., F. Koessler, K. Boun My and L. Denant-Boèmont (forthcoming). Road Traffic Congestion and Public Information: An Experimental Investigation, *Journal of Transport Economics and Policy*.

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.

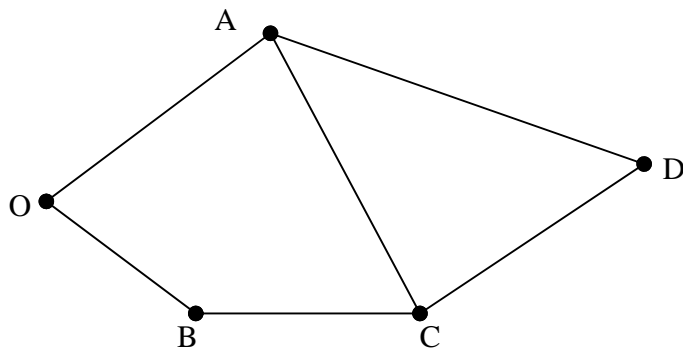
It is important that you do not talk, or in any way try to communicate, with the other participants during the experiment. If you have a question, raise your hand and a monitor will come over to where you are sitting and answer your question in private.

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.