

# Reducing the Effects of the Braess Paradox with Information Manipulation

Ana L. C. Bazzan and Franziska Klügl

**Abstract.** The Braess Paradox is a well known phenomenon in transportation engineering: adding a new road to a traffic network may not reduce the total travel time in it. In fact, some road users may be better off but they contribute to an increase in travel time for other users. This situation happens because drivers do not face the true social cost of an action. Previous works have shown that in commuting scenarios, where people use the same traffic network routinely, a continuous learning and adaptation process is a realistic scenario: road users can adapt to the traffic conditions and will eventually learn to avoid the situation in which the cost is higher. However, this learning process can take a long time. Moreover, because the process is very sensible to the cost function and to the number of agents using the network, a more efficient approach to distribute agents in the network is to let the traffic control center to acquire and process data regarding the occupancy of the available roads and compute the optimal distribution (from the point of view of the whole system). With this information, manipulated information can be passed to the road users. The interesting point is what happens when drivers simultaneously receive this kind of information and are involved in learning processes. Thus this paper reports results obtained after simulations of several situations related to the Braess scenario: only uninformed agents using the network; with different shares of uninformed agents; drivers adapting to the traffic conditions under different learning probabilities; drivers receiving forecast; and drivers receiving manipulated information.

## 1. Introduction

The Braess Paradox was originally presented by Braess in 1968 [4]. This “paradox” consists of a phenomenon which contradicts the common sense: in a traffic network, when a new link connecting two points (e.g. origin and destination) is constructed,

---

Author partially supported by CNPq.

it is possible that there is no reduction regarding the time necessary to commute from the origin to the destination. Actually, frequently this time increases and so the costs for the commuters.

Several authors have investigate the original problem ([1, 2, 7, 8, 10] among others) as it will be detailed in the next section.

The present work describes a model already partially used by Tumer and Wolpert [10] who have investigated the use of a multi-agent system to control routing of packages in a computer network using the so-called Collective Intelligence (COIN) formalism. The authors conclude that the network is also sensible to the paradox in some cases. It happens because the agents, by trying to reduce their individual routing times in a greedy way, end up increasing the global time.

We depart from Tumer and Wolpert work in what concerns the scenario and the determination of private goals for each agent. We use the classical scenario proposed by Braess, i.e. road traffic, in which a new link added to the network is very attractive to the users since the commuting time there is low. Besides, instead of using the COIN formalism, we introduce a kind of “manipulation” of the information given to agents about the state of the system in order to distribute the agents more evenly.

The discussion on the Braess Paradox and previous approaches to it is presented in the next section. Section 3 describes the approach based on providing information manipulation to the road users and all the situations we simulated. The results of these simulations are reported in Section 4 while Section 5 discusses the conclusions and future work.

## 2. The Braess Paradox and Some Previous Approaches

### 2.1. Basic Braess Scenario

In the scenario proposed by Braess [4], drivers or agents can select between two or three available paths to commute from origin **O** to destination **D**, as depicted in Figure 1, items (a) and (b) respectively.

In the network, commuting times are computed by means of the functions depicted for each link. They are all function of the flow or number of vehicles (e.g.  $T_{OQ} = f * 10$ ), where  $f$  is the number of vehicles in that particular edge.

In the configuration depicted in Figure 1, item (a), it is clear that if we have 6 vehicles (the original number in Braess paper), the equilibrium occurs when routes OQD and OPD carry 3 vehicles each. No one would be better off changing route.

In order to understand the paradox which occurs in the configuration depicted in Figure 1, item (b), let us consider an increasing number of vehicles starting with only one. For a single vehicle, route OQPD is much cheaper: it takes only  $10 + 11 + 10 = 31$  time units to commute from O to D. Taking OPQ or OQD would take 61 time units. It is easy to see that for more than 2.58 drivers, the new route is not advantageous anymore, so drivers would not use OQPD at all.




FIGURE 1. Two configurations of the network in the classical scenario of the Braess Paradox: (a) two possible routes: OPD and OQD (four-link network); (b) with additional route OQPD (five-link network)

Table 1 shows the costs for each of the five links, as well as the cost over all links, for different distribution of drivers among the three routes. The paradox is clear: when the six drivers are distribute in the two original routes, then the total cost is 176 units. However, when they find themselves two in each route, then the total cost increases to 196. This is so because the last case is the user equilibrium: each route takes 92 units so there is no incentive to change route. In the former case, routes OQD and OPD cost 83 each while route OQPD costs 70 thus being an incentive for drivers to change to OQPD. The problem is that when too many change, then the cost can be very high (see for instance the last line in the table, where the total cost is 236 because all six drivers use OQPD).

Nb. of Drivers			Cost of Link					Cost
OP	QD	QP	OP	QD	QP	OQ	PD	Total
3	3	0	53	53	10	30	30	176
2	2	2	52	52	12	40	40	196
1	2	3	51	52	13	50	40	206
1	3	2	51	53	12	50	30	196
2	1	3	52	51	13	40	50	206
2	3	1	52	53	11	40	30	186
3	1	2	53	51	12	30	50	196
3	2	1	53	52	11	30	40	186
6	0	0	56	50	10	0	60	176
0	6	0	50	56	10	60	0	176
0	0	6	50	50	16	60	60	236

TABLE 1. Cost for link and total cost for different distributions of drivers

## 2.2. Related Work

Following the seminal work by Braess [4], many authors have been proposing applications and modifications on the formulation of the problem in order to avoid the paradox. Smith’s paper [9] aims to show how, in a particular simple case, total journey time varies with the travel time along an uncongested link. This result is of particular relevance to towns with a good bypass or, better, a good outer ring road.

From the perspective of the economics of traffic, Arnott and Small [2] analyse three paradoxes in which the usual measure for alleviating traffic congestion, i.e. expanding the road system, is ineffective. The resolution of these paradoxes — among them the Braess Paradox — employs the economic concept of externalities (when a person does not face the true social cost of an action) to identify and account for the difference between personal and social costs of using a particular road. For example, drivers do not pay for the time loss they impose on others, so they make socially-inefficient choices. This is a well-studied phenomenon more generally known as The Tragedy of the Commons [5].

Regarding the Braess Paradox, in the scenario analysed by Arnott and Small the travel time for each of the two original routes is 20 minutes, while after the addition of the new route the travel time for the equilibrium situation rises to 22.5 minutes for each route.

Returning to traffic engineering and physics communities, in Penchina [8] the “simplest anti-symmetric” two-path network is described which exhibits the Braess paradox. A discussion of the good effects (non-paradoxical) of a bridge (especially a two-way bridge) is also included. Their Minimal Critical Network and graphical solution technique gives a clear understanding of the paradox for this network. They are also especially useful for analysis of sensitivity to such extensions as, e.g. changes in parameters, elastic demand, general non-linear (even non-continuous) cost functions, two-way bridges, tolls and other methods to control the paradox, and diverse populations of users. It is shown that the paradox occurs in a simpler network than previously noted, and with a larger Braess penalty than previously noticed.

Yang and Bell [11] work deals with network design via the Braess paradox and show how this capacity paradox can be avoided by introducing the concept of network reserve capacity into network design problems. Pas and Principio [7] examine properties of the paradox and show that whether the paradox does or does not occur depends on the conditions of the problem (link congestion function, parameters and the demand for travel). Akamatsu [1] explores the properties of dynamic flow patterns on two symmetrical networks.

Although the literature from the transportation branch also proposes some ways of avoiding the paradox, all of them concentrate on the parameters of the network, not on the driver itself, thus relegating the human component which plays an important role. One can argue that even if the “right” network is constructed (i.e. one which avoids the paradox from the point of view of the mathematics of the

problem as it was proposed in the literature), drivers will always seek to maximize their individual payoffs which frequently leads to sub-optimal global distribution of traffic in the roads of the network. Besides, even if the “right” parameters are taken into account in the design of the network, the increasing demand for mobility in our society leads to a rapid obsolescence of the network. Thus we argue that it is not correct to talk about “right” parameters.

### 2.3. Multi-agent Approaches

Multi-agent systems approaches are interesting in the Braess scenario since one can focus on the issues related to the driver and its decision-making process. Tumer and Wolpert [10] have investigated the use of a multi-agent system to control routing of packages in a computer network, as well as the sensitivity of the network to the Braess paradox using the Collective Intelligence (COIN) formalism. Performance using COIN is compared to the case in which each agent in a set of agents estimates the “shortest path” to its destination. Since each agent decision (which is based on this estimation) ignores the effects of the decision of other agents on the overall traffic, performances based on estimation of the shortest path are badly sub-optimal. The authors conclude that the network is also sensible to the paradox in some cases. It happens because the agents, by trying to reduce their individual routing times in a greedy way, end up increasing the global time.

In the COIN approach, the new goals are tailored so that if they are collectively met the system maximizes throughput. The world utility,  $G(\zeta)$ , is an arbitrary function of the state of all agents across all time. The utility for an agent is given by the difference between the total cost accrued by all agents in the network and the cost accrued by agents when all agents sharing the same destination are “erased”.

Bazzan and Klügl [3] present a learning heuristic and depart from Tumer and Wolpert work in what concerns the scenario and the determination of private goals for each agent. The classical scenario proposed by Braess, i.e. road traffic, is used, in which a new link added to the network is very attractive to the users since the commuting time there is low. However, since it overlaps (see Figure 1) with the existing paths, the global commuting time increases.

The approach in [3] is based on learning and adaptation by means of an heuristic capable of minimizing both the global and local performance losses. Moreover, the heuristic was developed with the multi-agent paradigm in mind requiring very little processing of each simulation agent making it very scalable. Besides, it is important to notice that agents do not need to explicitly communicate in order to coordinate their choices. Thus, the learning heuristic proposed – called A2B (Adapt to Braess) – improves the performances because it implicitly includes some factors of the global performance in the individual ones. Comparing the situation in which the agents use this heuristic to the situation in which the selection of routes happens at random, the performance greatly improves, i.e. the total commuting time decreases.

In short, A2B is based on reinforcement learning. Agents form tactics for selecting a route according to the reward obtained in that route in the past. Each tactic is associated with a learning rule.

Be  $i$  the index of an agent, and  $j$  the index of an action available to agent  $i$ . In a set of actions  $A_i = (a_1, \dots, a_j)$ , a learning rule is a rule which specifies the probabilities  $P_i = (p_{i,t}(a_1), \dots, p_{i,t}(a_j))$  as a function of the rewards obtained by selecting actions in the past. In the future, each action is selected according to its probability.

In the Braess scenario, this means that each agent remembers how many times s/he used each route and the total amount of time spent traveling on each of them. With this information in mind, the agent then calculates the average travel time taken at each route and selects the one with the shortest time. The average travel time in each route is actually computed by means of a discount factor  $\delta$ , in order to allow us to play with the fact that it is not desirable that agents remember *all the past* with the same weight they remember the more recent outcomes. Thus, the last time measured at a given route is multiplied by  $\delta$  while the averaged past time is multiplied by  $(1 - \delta)$ .

Hence,  $p_r$  – the probability of selecting each route  $r$  – is updated according to Equation 1, where  $r$  is the index of each available route,  $\sum_r$  is the sum over all routes, and  $\sum_t T_{OD,r}$  the sum of travel time for commuting from O to D using route  $r$  along time  $t$ .

$$p_r = \frac{\frac{1}{\sum_t T_{OD,r}}}{\sum_r \left( \frac{1}{\sum_t T_{OD,r}} \right)} \quad (1)$$

### 3. Approaches and Scenarios

Although this scenario is not one of binary decision, the basic conditions from the *iterated route choice* (IRC) scenario [6] are valid. This is a model for adaptive choice in which each agent has no information about other agents. They decide which alternative to select based on a local inference about the costs or rewards of each available action from the action set. This inference is based on the update of the probability according to which an agent selects each alternative action.

In the adaptive scenario the agent  $i$  updates these probabilities with a certain periodicity according to the rewards it has obtained selecting alternative  $r$  up to that point. The update of the heuristic is done via reinforcement according to the following formula:

$$heuristic(i) = \frac{\sum_r reward_r(i)}{\sum_t reward(i)} \quad (2)$$

This basic scenario can be extended to give agents forecast information. Now, the decision-making process was performed in two phases. First, agents make their initial selection (first decision) based on the adaptation process introduced above. Based on this information collected from agents, a control center computes the reward for every agent and sends this information back: the agents receive a forecast information regarding the potential reward they would have if they would keep their first decision. Then, they have a second chance to actually take their first choice or change it (second decision). Finally the actual selections are made yielding the actual rewards.

The third situation occurs when “manipulated information” is given to agents to try to force an equilibrium distribution between the alternatives. For example: if the current distribution is  $number_{OQPD} = 500$  and  $number_{OPD} = number_{OQD} = 500$ , it is clear that the system could do better in terms of *global sum* of travel times. Also, the rewards of agents are sub-optimal (not necessarily for each individual). When the control system perceives such a situation, it tries to induce agents to come closer to the global optimum by given manipulated forecast to them. In the particular case above, the control system can give a bad forecast for route  $OQPD$  to try to divert drivers to other routes.

In our scenario, drivers or agents can select between the three available paths to commute from origin **O** to destination **D**, as depicted in Figure 1, item (b). In this network, commuting times are computed by the means of the cost functions depicted for each link. They are all function of the flow or number of vehicles as already explained in Section 2. Moreover, since this is a commuting scenario, agents perform this selection repeatedly. We then measure the number of drivers selecting each of the available routes, as well as the total time of commuting.

Thus, in some situations we simulate, agents just select a route at random, while in others they consider their past experience regarding performance measured by commuting time <sup>1</sup>.

The implementation was done using SeSAm, a shell for developing agent-based simulation<sup>2</sup>. In the implementation, agents and vehicles can be consider a single unit. The situations we have simulated are:

- I Drivers are uninformed i.e. they all randomly select one of the 3 routes
- II Drivers select one of the 3 routes according to the algorithm proposed in [3] (A2B) which includes adaptation based on the past performance
- III Different proportions of uninformed drivers select randomly while the rest use the A2B learning algorithm
- IV All drivers receive forecast information about the occupancy of the road it has selected and either keep their decision or reselect a route
- V Drivers receive manipulated information regarding the forecast

<sup>1</sup>Up to now we use only this measure which is also standard in the field. However we do not completely agree that *commuting time* is the only issue drivers consider, although it is certainly the most important. We are investigating how to include driver comfort, knowledge of the network, willingness to change route, and other issues to the performance function.

<sup>2</sup>available for download at [www.simsesam.de](http://www.simsesam.de)

Simulations of situation I is clear:  $N$  drivers randomly select between routes OPD, OQD, and OQPD. If  $\vec{P}$  is the array of probabilities for selecting routes, then the initial distribution is  $\vec{P}_0 = (1/3, 1/3, 1/3)$ . One can argue that this initial distribution does not reflect the commuting times. However, since the cost of commuting is influenced by the number of drivers at any of the routes (which is unknown), each agent actually has no preference at the beginning. Therefore, it puts equal probability to each of the routes.

Situation II was simulated taking into account a behaviour by agents which follows these ideas: drivers use the selection strategy A2B. This keeps track of choices in the past, sums the performances regarding each of the three possible route, and updates an array of probabilities for selecting each route. The performances depend on the choices of all other agents, what configures a typical problem of coordination.

While in situation II all drivers are informed, in situation III we play with the parameter *rate of random drivers*, i.e. the rate of drivers who are new in the network so they do not have enough information to use the learn-and-adapt algorithm. Thus they just select a route randomly. We use rates of 0 (which is situation II), 25, 50, 75, and 100% (situation I).

In situation IV drivers receive forecast information about the selected route, as explained above when we introduced the IRC model plus its extensions. Also here we played with the number of random drivers so when this rate is 0% everyone gets the information and adapt. When it is 100% everyone ignores the forecast. The difference in this case is that the forecast information may influence the decisions of the agents since not all will keep their first intention. It is expected that the noise increases and that the third route is used more frequently.

Finally, in situation V, the traffic control center gives “false” forecast to prevent drivers from using a route which is expected to be increase the global costs. An interesting question here is that agents can learn that the information cannot be trusted, as they not only receive this kind of information but rely on their learn-and-adapt heuristic. The situation regarding false or manipulated forecast happens frequently with radio broadcast of traffic messages. When it brings no real payoff, drivers tend to have a critical view of it.

## 4. Description of the Experiments and Results

In this section we discuss the results regarding the five situations presented above. In all cases, the  $N$  agents interact and have to implicitly coordinate since they are using shared resources. For the experiments, we use  $N = 1500$  agents,  $\delta = 0.8$ , and the learning probability is 0.2. Cases with less agents and in particular with  $N = 6$  are reported in [3].

As for the cost functions to compute the commuting time, they are shown in Table 2. For  $N = 1500$ , the paradox happens with parameters shown in Table 3.



Links	Function
OQ	$T_{OQ} = \beta_1 * f_{OQ} + \alpha_1$
PD	$T_{PD} = \beta_1 * f_{OQ} + \alpha_1$
OP	$T_{OP} = \beta_2 * f_{OP} + \alpha_2$
QD	$T_{QD} = \beta_2 * f_{QD} + \alpha_2$
QP	$T_{QP} = \beta_4 * f_{QP} + \alpha_4$

TABLE 2. Functions to compute the time to traverse each route in the network.

Parameter	Values
$\alpha_1$	70
$\alpha_2$	500
$\alpha_4$	5
$\beta_1$	0.2
$\beta_2$	0.05
$\beta_4$	0.5

TABLE 3. Values for parameters when  $N = 1500$ .

With these parameters, the global costs can be computed for the particular situations of interest:

1. when the  $N = 1500$  drivers are distributed equally among the 3 routes, the global cost is 795 for  $OQPD$ ,  $OQD$ , and  $OPD$ ; this is also the average cost per driver, while the global cost is 1845 (over all sub paths).
2. when all drivers avoid the  $OQPD$  route and 750 use each remaining route, the total cost is 445 for  $OQPD$  and 757.5 for the other two, thus adding up 1520.

In the simulations, the dimensions measured were mainly distribution of drivers in each route, and the probability of selecting each route (averaged over the  $N$  agents). The cost is a direct function of the distribution of drivers so we do not plot costs here.

In situation I the paradox can be observed: since the selection of the 3 routes is random, some drivers do select route  $OQPD$  and this increases the global commuting time.

The case of situation II for  $N = 1500$  agents was discussed in [3]. The majority of the agents tended to avoid the third route, although this does not happen completely due to the inertia in the learning process. The distribution of agents between the two remaining routes tended to be equal, as it maximizes the average performance of the agents. The learning process is slow and sensible to the cost parameters.




FIGURE 2. Number of drivers selecting route *OQPD* along time, under different conditions (only random selection, with learning, with forecast, and under information manipulation ( $N = 1500$ )).

The simulations of situation III show that the increase in the number of random drivers increases the usage of the route *OQPD* thus increasing the global cost.

As for situation IV, the forecast given to the drivers does not help when it is the correct one, that means, when it is the information actually computed from the first intention of each driver. This happens because, in the majority of the cases, the forecast is worse than the expected value a driver computes for the selected route. Thus, the driver rejects his first selection and perform a second selection randomly. As above, situations in which too many drivers choose randomly brings to an usage of *OQPD* and therefore, to an increase in costs. Even if drivers have a tolerance to the forecast (i.e. they do not simply reject their first selection but tolerate a worse forecast up to some threshold), the situation does not improve because in this case the system has a performance similar to that in situation II.

Therefore, the best results were achieved with situation V, i.e. with information manipulation. The gain in performance is summarized in Figure 2. We plot the number of drivers who select route *OQPD* along time, for the situations above: with only random selection, with the learn-and-adapt process, with forecast (and no random driver), and under manipulation of information (and no random driver). One sees that only the last case brings some of the number of drivers to avoid the route *OQPD*. In the other cases, around  $1/3$  of the drivers still select this route (curves overlap).




FIGURE 3. Number of drivers selecting routes *OQPD* along time, under manipulation and with different rates of random selection ( $N = 1500$  agents).

Now, it is interesting to investigate what happens if not all drivers are informed, i.e. when random drivers perform the route selection together with drivers who learn and get manipulated forecast. We have simulated this situation with different shares of uninformed drivers: 0, 25, 50, 75, and 100% and these results are depicted in Figure 3).

The case with the lowest number of drivers selecting *OQPD* is the one in which everybody is informed (thus, the same curve which appeared in Figure 2). This shows that the role of information is significant and, moreover, that the role of manipulated information is even more important than the correct forecast since the latter can introduce noisy in the learning process of the drivers, especially when too many drivers are unaware of this information. For higher shares of uninformed drivers, the trend is that the route *OQPD* is selected more often. Of course, in case 100% of drivers do not get information, the manipulation has no effect and thus, the selection is random and inefficient both for each driver and for the whole system.

## 5. Conclusion and Future Work

In the classical scenarios on the Braess Paradox reported in the literature, if drivers or agents of any kind, act to maximize their own profits, the global performance of the system may decrease. This happens because the global goal opposes the individual goals in most cases.

While studying the Braess Paradox, we noticed that this could be an interesting scenario for investigating issues related to learning and adaptation as well as information manipulation in order to minimize both the global and local performance losses. Moreover, the heuristic was developed with the multi-agent paradigm in mind requiring very little processing of each simulation agent making it very scalable. Besides, it is important to notice that agents do not need to explicitly communicate in order to coordinate their choices.

In this paper we have studied the effect of information manipulation in the Braess scenario. The simulations show that it is useful to manipulate the forecast information given to the agents. Doing so, the control system is able to divert them to the more convenient alternative, from the point of view of both the overall system and the individual agents (as this is the situation in which individual rewards are the highest). This holds in higher or lower degrees for different shares of uninformed drivers.

The main conclusion is that having agents provided with the most accurate information (information about the *actual* state of the system) is not necessarily good.

This work is based on a series of assumptions that may not be bearable in every real world application. First, we assume a global control component that is able to compute the *exact* utility of the agent decisions for producing the forecast information. Although this degree of exactness in forecast might not be necessary, under which circumstances the existence of such a forecast system is realistic in general?

Second, and more interesting from the perspective of mechanism design, is the assumption that the central component acts in the interest of the highest system performance. However, when we have other interests involved (for instance when there are several competing “central” controller components), some of them may act primarily to disturb the others.

Therefore, the next directions of this research are twofold: the system may be able to learn the share of uninformed and account for this share when giving manipulated information; and designing of more complex agent architectures to accommodate competition and self-interest.

## References

- [1] AKAMATSU, T. dynamic traffic equilibrium assignment paradox. *Transportation Research B* 6 (2000), 515–531.
- [2] ARNOTT, R., AND SMALL, K. The economics of traffic congestion. *American Scientist* 82 (1994), 446–455.
- [3] BAZZAN, A. L. C., AND KLÜGL, F. Learning to behave socially and avoid the braess paradox in a commuting scenario. In *5th Workshop on Decision Theoretic and Game Theoretic Agents* (July 2003), S. Parsons and P. Gmytrasiewicz, Eds., pp. 107–112. held together with AAMAS 2003.

- [4] BRAESS, D. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* 12 (1968), 258.
- [5] HARDIN, G. The tragedy of the commons. *Science* 162 (1968), 1243–1248.
- [6] KLÜGL, F., AND BAZZAN, A. L. C. Simulation of adaptive agents: Learning heuristics for route choice in a commuter scenario. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)* (Bologna, Italy, July 2002), vol. 1, ACM Press, pp. 217–218. extended abstract.
- [7] PAS, E. I., AND PRINCIPIO, S. L. Braess’ paradox: Some new insights. *Transportation Research B* 31, 3 (1997), 265–276.
- [8] PENCHINA, C. M. Braess paradox: maximum penalty in a minimal critical network. *Transportation Research A* 31, 5 (1997), 379–388.
- [9] SMITH, M. J. In a road network, increasing delay locally can reduce delay globally. *Transportation Research* 12, 6 (1978), 419–422.
- [10] TUMER, K., AND WOLPERT, D. Collective intelligence and braess’ paradox. In *Proceedings of the AAAI* (2000), AAAI Press.
- [11] YANG, H., AND BELL, M. G. H. A capacity paradox in network design and how to avoid it. *Transportation Research A* 32, 7 (1998), 539–545.

Ana L. C. Bazzan  
Instituto de Informática  
UFRGS  
Caixa Postal 15064  
91.501-970 Porto Alegre, RS, Brazil  
e-mail: [bazzan@inf.ufrgs.br](mailto:bazzan@inf.ufrgs.br)

Franziska Klügl  
Dep. of Artificial Intelligence  
University of Würzburg  
Am Hubland  
97074 Würzburg, Germany  
e-mail: [kluegl@informatik.uni-wuerzburg.de](mailto:kluegl@informatik.uni-wuerzburg.de)