

# Implicit Coordination in a Network of Social Drivers: The Role of Information in a Commuting Scenario

Ana L.C. Bazzan\*, Manuel Fehler, and Franziska Klügl

<sup>1</sup> Instituto de Informática, UFRGS, Caixa Postal 15064,  
Porto Alegre 91.501-970, RS, Brazil  
bazzan@inf.ufrgs.br

<sup>2</sup> Dep. of Artificial Intelligence, University of Würzburg,  
Am Hubland, Würzburg 97074, Germany  
{fehler, kluegl}@informatik.uni-wuerzburg.de

**Abstract.** One of the major research directions in multi-agent systems is dedicated to learning how to coordinate and whether individual agents' decisions can lead to globally optimal or at least acceptable solutions. Our long term goal is to study the effect of several types of information to guide the decision process of the individual agents. This present paper addresses simulation of agents' decision-making regarding route choice, and the role of an information component. This information can be provided by group colleagues, by acquaintances from other groups (small-world), or by route guidance. Besides, we study the role of agents lying about their choices. We compare these scenarios, concluding that information (from some kind of source) is beneficial in general: lying helps only to a certain extent, and route guidance is the best type of information.

## 1 Introduction

In multi-agent systems it is almost impossible to oversee the issue of learning how to coordinate. In particular, we focus here on a route choice scenario, in which social drivers have to select a route based on information sharing. The behavior of these agents is influenced by information, be it route recommendation, be it driving experiences exchanged within a group. Thus, we address not only the individual agent behavior, but also how information sharing and recommendations could influence this behavior so that a coordinated situation emerges.

A route choice scenario is normally characterized by an agent facing repeated action selection. In a previous paper [7], we concluded that the more reliable the information that an agent gets about the current and future state of the environment is, the more his actions depend on his beliefs about the decisions of the other agents.

The rest of this paper is organized as follows. The next section briefly reviews some background ideas on decision-making regarding binary choice and organization of agents in groups, with and without information sharing. In Section 3 we present the scenarios used to simulate route choice under several conditions, whereas the results of the corresponding simulations are presented in Section 4. The last section summarizes the conclusions and outlines the possible extensions.

---

\* Author partially supported by CNPq.

## 2 Background

### 2.1 Coordination Games and Route Choice

The *El Farol* Bar Problem (EFBP) [1], and, more generally, the Minority Game (MG) [5] are particular instances of coordination models. Basically, these deal with the situation in which  $N$  players or agents have the choice between two alternatives (e.g. buy/sell stocks, go to a bar, use route A or B). Variants of these games have been proposed in economics, computer science, and physics. Two relate to multi-agent systems and traffic problems. In [2], personalities are introduced in the MG and populations of agents with these personalities interact in a commuting scenario.

In the second one [6] under the focus of dynamic adaptation, in a commuting scenario, in which drivers have to daily select a route to drive from home to office and back. One route, namely  $M$  (main), provides more capacity. The other alternative is a secondary one (thus  $S$ ). At the end of the trip, every agent gets a reward that is computed based on the number of agents who selected the same alternative. This mimics the actual travel time experienced by the driver himself. The agents know nothing about the reward function or about other agents, but their decisions do influence the reward each receives. Rewards for each agent  $i$  were computed based on the the number of agents selecting each route; the less agents select one alternative, the more reward each of them receive for their choice. The overall goal is to reach the best distribution of agents. This is the case when the system reaches the user equilibrium and no agent can change to another alternative at a strictly lower cost. Also, especially for traffic scenarios, the Wardrop Second Principle (optimum) states that the average travel time is minimal [9]. This corresponds to the best social optimum in economic terms.

To investigate this, a simple model for adaptive choice was initially developed [6]. Each agent decides which route to select based on the probability according to which it selects the main route. In the adaptive scenario the agent updates this heuristic with a certain periodicity according to the rewards he has obtained selecting that alternative up to that point. An important factor is how often and in which intervals the heuristic is updated. Changes in this and other parameters are detailed in [6]. Without any information from outside, this yields a configuration where, on average, the agents learn the optimal heuristic.

### 2.2 Social Attachments and Networks

Also related to the present paper, we discuss some previous works on simulating the Iterated Prisoner's Dilemma (IPD) under different conditions, with agents having various kinds of social attachment.

Watts and Strogatz [10] studied networks of coupled elements through an analogy with the *small-world phenomenon*. The small-world concept is based on the fact that, in large societies, there is normally a shortcut between any two persons via a path of acquaintances [8].

In [3], the performance of a society composed of agents playing the IPD in the presence of agents with attachment to others was analyzed. These agents may have an altruistic behavior towards its acquaintances. These so-called altruistic agents are interested in the good performance of their group as a whole, as well as on their own,

since the social group provides also a base for support in case the agent itself is not performing well.

This contributes to the further understanding of how coordination mechanisms can be developed. It is not enough to consider pure rationality when agents are autonomous but also interact in a social group.

### 3 Description of the Information Sharing Scenarios

The main motivation is to check whether egoistic agents who seek only to maximize a utility function or are overconcerned with self-interest can miss good opportunities for themselves. We use a scenario similar to the one described in Section 2.1 adding the issue of people interacting in groups, using the iterated route choice and the two options (main and secondary routes). The goal remains to have agents distributed between the two routes so that no agent is better off by deviating from its selection. The learning scheme is based on information sharing among the group and, in some scenarios, among people in a network of acquaintances, bringing in the idea of small world explained in Section 2.2. The reward functions are slightly changed here, as explained next.

The  $N$  agents acting in these scenarios have to select one of the two available routes as explained in Section 2.1. After all agents have made their decisions and have driven, they receive their rewards, which are inversely proportional to their travel times. Since we use an abstract model here, the travel time is actually computed as in Eq. 1.

$$R_i = \begin{cases} \frac{4}{3} - \frac{M}{N} & \text{if } M \text{ selected} \\ 1 - \frac{S}{N} & \text{if } S \text{ selected} \end{cases} \quad (1)$$

where:

- $R_i$  is the reward for driver  $i$
- $M$  and  $S$  are the number of drivers driving main and secondary routes respectively
- $N$  is the total number of drivers ( $M + S$ )

These formulas arise from the fact that we assume that topological constraints of both routes allow twice more vehicles at the main route than at the secondary one (this is a more didactic example than just allowing 50-50% distribution). Thus, at the equilibrium,  $\frac{2}{3}$  of the vehicles should drive on main. The reward on main would then be, for each driver  $i$ ,  $R_i = \frac{2}{3} - \frac{M}{N}$ . Similarly, the reward on the secondary route would be  $R_i = \frac{1}{3} - \frac{S}{N}$ . However, these formulas would yield negative rewards (e.g. in case all drivers go to side, the reward would be  $-\frac{2}{3}$ ), which is not desirable. Therefore, we normalize the rewards adding  $\frac{2}{3}$  in all cases, what puts the distribution of rewards between 0 and  $\frac{4}{3}$ . Table 1 shows rewards for some particular distributions of drivers.

When agents meet their colleagues, they can share information about which route was the best in the group. As said before, from time to time these agents can meet acquaintances from other groups, sharing information about routes as well.

This set of simulations is based on groups composed by either “social” people (in the sense that they share their information within their group and with acquaintances

**Table 1.** Rewards for particular distribution of drivers between main and secondary routes, for N=150

#nb. main	#nb. sec.	rew. main	rew. sec.
149	1	$\approx 1/3$	$\approx 1$
1	149	$\approx 4/3$	$\approx 0$
<b>100</b>	<b>50</b>	<b>2/3</b>	<b>2/3</b>
50	100	1	1/3
75	75	5/6	1/2

about which route is good), “nasty” people (always lie giving the opposite information), and “noisy” drivers (do not belong to any group since they do not commute frequently).

We can see these groups as being formed by colleagues (e.g. people working together). Agents in the group can ask for rewards of colleagues (or, more directly, which route was good) and eventually (if the reward is higher) do what other(s) have done. This is important because even if agent  $i$  has a good driving history (e.g. it is always on time because it has a good strategy for selecting routes), it does not bring much to be the only one on time in the office. Ideally, assuming that they work together or meet every morning for important decisions, it is desirable that *every one* in the group is on time! (thus the motivation for route information sharing).

Some of these agents also know people belonging to other groups (e.g. partners in leisure activities, etc.) in a small-world-like configuration, so that they have opportunity to ask for route tips from time to time. Finally, there are agents who like to experiment a different route with a given probability.

Noisy drivers select routes randomly and do not ask for information. However, they might meet other non-noisy agents (e.g. at the gas station) and give information.

Nasty agents have an operational behavior similar to the social drivers except that they do not care about the group or acquaintances; they just want to drive in the route with as few drivers as possible. Thus, in a quite “naïve” way, they give the opposite information (e.g. if one used the main route and got a good reward, it advises others to use the secondary one, if asked). In fact, in some scenarios, given that the information propagates, this can be a good strategy, as it will be shown. At this stage, we do not have groups of mixed people because we want to compare performance at the group level.

Finally, in some scenarios, we introduced a private message system (PMS) which recommends drives what to do. The aim of the PMS is to direct people to an equilibrium situation, based on the optimal distribution of drivers. Again, we assume that noisy drivers do not get this kind of message.

Several scenarios were simulated, changing the composition of groups, the characteristics of agents, and the presence or not of PMS. These scenarios are summarized in Table 2.

All scenarios of type 1 are without PMS, and without noisy and nasty drivers. In scenario **1A**, infos are exchanged among group colleagues only; everyone select the route of the best colleague in the group. Scenario **1B** is pretty much the same, except that agents have a probability  $p_1$  **of trying a different route**. **1C** is also similar to 1A, except that agents have a probability  $p_2$  **of asking for infos outside their groups**.

**Table 2.** Description of scenarios (\* means that the parameter assumed different values, as described in the experiment’s specific section)

scenario	PMS	nb. of nasty	nb. of noisy	$p_1$	$p_2$	$p_3$
<b>1A</b>	no	0	0	0	0	0
<b>1B</b>	no	0	0	*	0	0
<b>1C</b>	no	0	0	0	*	0
<b>2A</b>	no	0	*	0	0	0
<b>2B</b>	no	0	*	0.01	0	0
<b>2C</b>	no	0	*	0	0.2	0
<b>3A</b>	no	*	0	0	0	0
<b>3B</b>	no	*	0	0.01	0	0
<b>3C</b>	no	*	0	0	0.2	0
<b>4A</b>	yes	0	0	0	0	0
<b>4B</b>	yes	0	0	0.01	0	*
<b>4C</b>	yes	0	0	0	0.2	*
<b>5A</b>	yes	*	0	0	0	0
<b>5B</b>	yes	*	0	0.01	0	0
<b>5C</b>	yes	*	0	0	0.2	0
<b>5D</b>	yes	*	0	0	0	*
<b>6A</b>	yes	0	*	0	0	0
<b>6B</b>	yes	0	50	0	0	*
<b>6C</b>	yes	0	50	0	0.2	0

Scenarios of type 2 are similar to type 1 but include noisy drivers. In **2A** we changed only their number. In **2B** we also have a non zero probability  $p_1$ . **2C** is similar to 2B but  $p_2$  is used instead of  $p_1$ .

Scenarios of type 3 are similar to type 2 but include nasty instead of noisy drivers. Their number was changed in scenario **3A** while **3B** and **3C** are similar to 2B and 2C respectively.

Scenarios of type 4, 5, and 6 all have agents receiving recommendations from the PMS. In type 4 there are no noisy or nasty drivers. In **4A** all agents do follow the recommendation. In **4B** they **deviate from the recommendation with probability  $p_3$**  (this means doing the opposite as recommended), plus they try different routes with probability  $p_1$ . **4C** is similar to 4B but we use  $p_2$  instead of  $p_1$ .

In type 5, there are nasty drivers. In **5A** we change only the number of them. In **5B** and **5C**  $p_1$  and  $p_2$  are non zero respectively. In **5D** we also vary  $p_3$ .

Type 6 is similar to type 5 but includes noisy instead of nasty drivers. In **6A** we vary the number of noisy drivers. The other drivers all follow the recommendations. In **6B** the number of noisy drivers is constant and we vary the probability to deviate from the recommendation ( $p_3$ ), and in **6C** we vary the probability  $p_2$ .

Simulations take 500 time steps (in the graphics we show up to 1000 *simulation steps* because each time step takes two simulation steps due to the nature of the simulation: decide and drive). The parameters used in the simulations are:  $N = 150$  drivers, and these are divided in 15 groups. Since agents are created with a probability of belonging to a group, groups may not have exactly  $N/15$  members. Depending on the scenario,

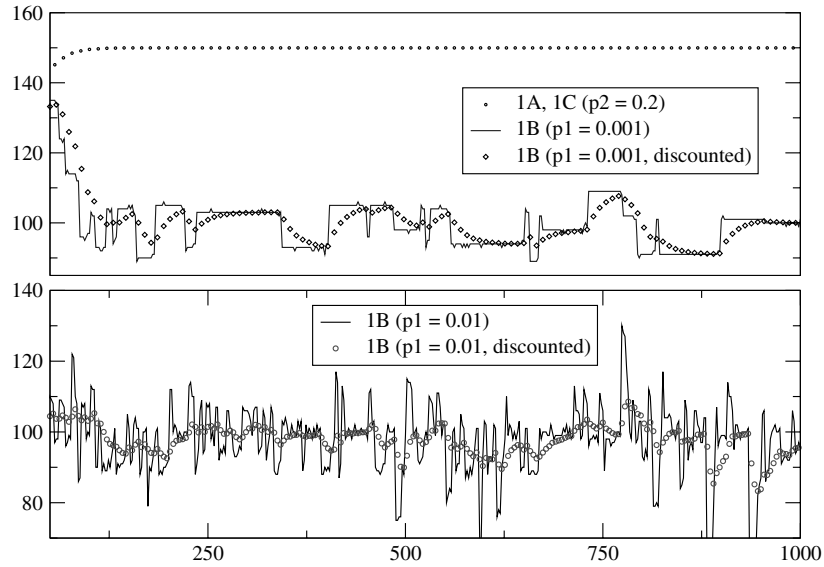
two of those groups are fixed: one group with only noisy and one group with only nasty drivers.

## 4 Results

We now present and discuss the main results regarding the scenarios just described. In the graphics we depict only the number  $M$  of drivers on main (the secondary is given by  $N - M$ ). In fact, we measure both the actual number at each time step and a discounted average (by a factor  $\delta$ ) as follow:  $\overline{M}_\delta = \delta \times \overline{M}_{old} + (1 - \delta) \times M_{current}$ . In the simulations presented here  $\delta = 0.9$  so that we put much less weight on the current time step. In most graphics we plot only the discounted average which is more smooth than the actual number  $M$  (see for instance how it looks when both are plotted at Fig. 2). Remember that the optimal situation is when two-thirds of the drivers (100) are on main and one-third (50) on the secondary route.

### 4.1 Scenarios Without Recommendation

In Fig. 1 we depict all simulations of type 1: case 1A is the topmost curve in the upper box. Here we see that after a short time the  $N$  agents select the main route! This happens because they share the information and all copy the best strategy in the group. Since there is no probability of trying something new, they are all stuck in a very poor situation in terms of performance (see Table 1). Similar performance is noticed for case 1C. This happens because, although agents do ask acquaintances in other groups, since they all behave the same regarding copying strategies and there are no noisy



**Fig. 1.** Number of drivers on main for scenarios 1A, 1B, and 1C, with  $p_1 = 0.001$  (upper box) and  $p_1 = 0.01$  (lower box)

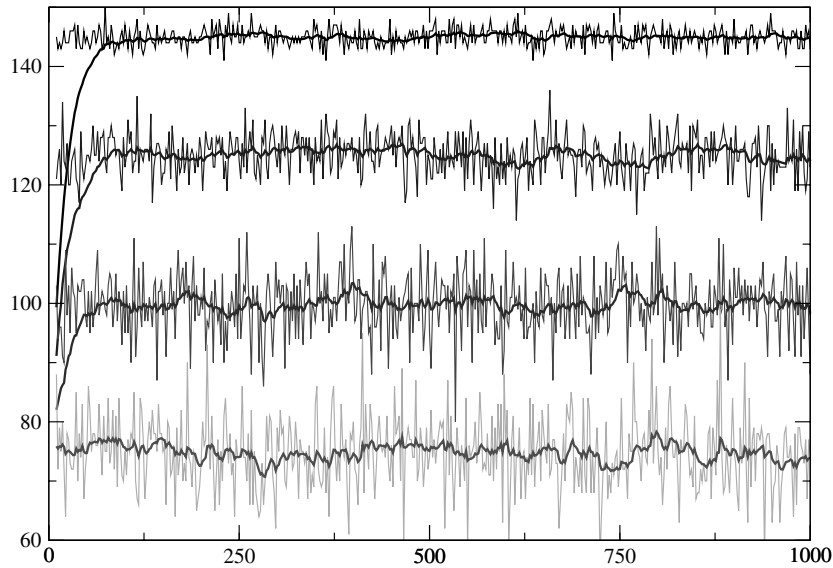
drivers, this behavior brings nothing. Thus, the result is that again all  $N$  drivers end up selecting main.

Of course these are deterministic and thus very unrealistic situations. In case 1B we used probability  $p_1$  of people trying different routes. This is depicted in Fig. 1 as well. The curves in the upper box are for  $p_1 = 0.001$  and the curves in the lower box are for  $p_1 = 0.01$ . We can see that in both cases the average number of drivers on main fluctuates around 100 drivers, with a much higher deviation in the latter case, which is explained by the fact that the experimentation happens more frequently.

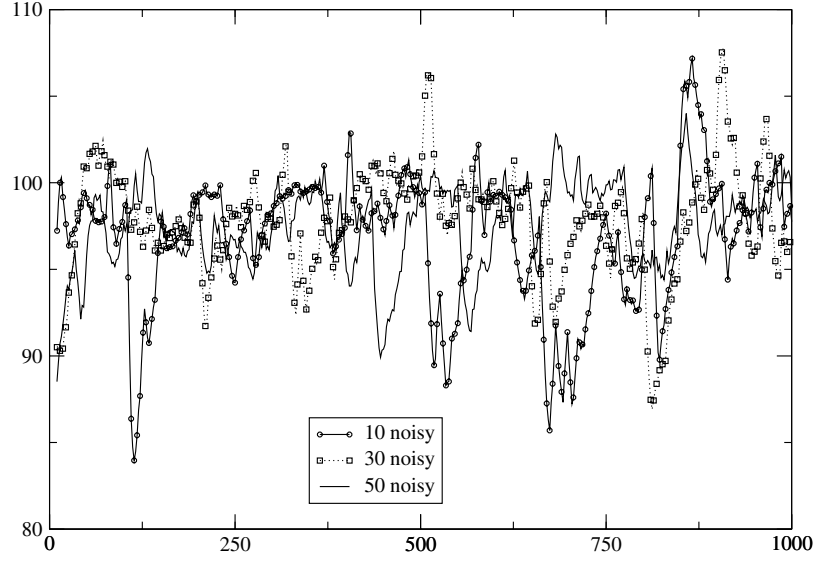
Fig. 2 shows the simulations for scenario 2A for various numbers of noisy drivers, with  $p_1 = p_2 = 0$ . From top to bottom, their numbers are 10, 50, 100, and 150. When there are only 10, on average 5 go to main and 5 to secondary, what puts the number of people at main to 145 on average (remember that the non-noisy just copy the best route and thus go to main as in 1A). Similar situation happens with 50 (125 on main), 100 (100 on main), and when all are noisy (of course 75 go to each route on average). Notice that the deviation from the expected average increases with increasing number of noisy drivers.

In the simulations depicted on Fig. 3 (scenario 2B), a probability  $p_1 = 0.01$  of trying another route was used, and the number of noisy drivers varied: 10, 30, and 50. From now on, in most of the figures, we show only the discounted average in order to plot several curves in a single graphic. The overall average is below 100 due to the noisy drivers who select randomly thus bringing the average as close to 75 the higher their number is.

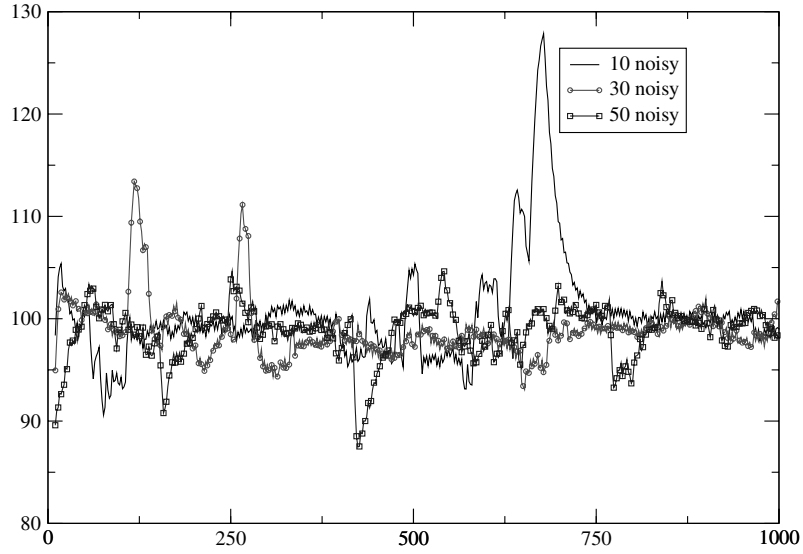
In scenario 2C, the probability  $p_2$  was set to 0.2 (keeping  $p_1 = 0$ ), changing the number of noisy drivers (10, 30, and 50). Fig. 4 shows the results: the more noisy drivers,



**Fig. 2.** Number of drivers on main for scenario 2A, for 10, 50, 100, and 150 (from top to bottom) noisy drivers



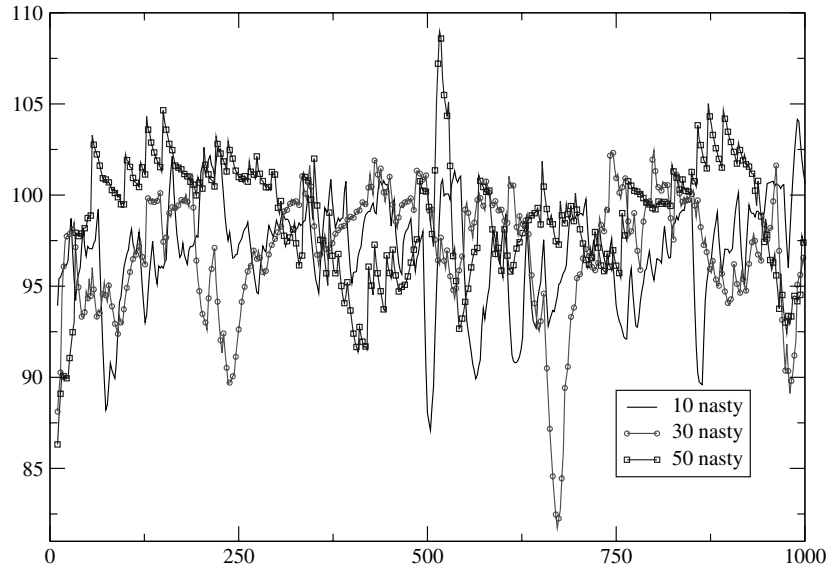
**Fig. 3.** Number of drivers on main for scenario 2B



**Fig. 4.** Number of drivers on main for scenario 2C

the more the small-world metaphor is effective. The high/low peaks in the curves appear because when agent  $i$  asks agent  $j$  in another group and learns a better route, then all agents of the group of agent  $i$  will copy and change, if agent  $i$  got a good reward.

Cases with nasty drivers are discussed next. In 3A,  $p_1$  and  $p_2$  are zero and we vary only the number of nasty drivers: 10, 30, 50, 75, 100, and 140. With 10 nasty drivers, those who select the side route at the beginning (5 on average) perform very well because



**Fig. 5.** Number of drivers on main for scenario 3B

all social drivers go to main (again, because they copy the best route in their groups). The remaining nasty drivers get wrong information from the others in the group and go to main as well. Thus, a few nasties are better off driving the secondary route. Similar situations happen for 30 and 50. However, with 75, 100, and 140 nasty drivers their performance decreases (in this order) because around 50% of them go to each route and they are simply too many in each. Compared to situations 1A and 1C (no nasty drivers), here, with the presence of false information, at least some of the nasty drivers have a good performance (especially when they are few).

To escape from this unrealistic situation, simulations were performed in which both social and nasty drives have a probability  $p_1$  (here equal to 0.01) to try another route. Results are depicted in Fig. 5. There is not so much difference comparing cases 3B and 1B: nasty drivers do not cause more fluctuation because every driver can find a better route.

The scenario with nasties was also combined with the small world (3C) using  $p_2 = 0.2$ . Only 10 nasty drivers are not enough to influence the other 140 (in a situation similar to 1C – see Fig. 1, topmost box), so that all social drivers get stuck selecting main, while the nasty drivers go to both and those on the side route have a much better reward. However, when we increase the number of nasty drivers, the wrong information they give to people outside their group plays the role of diversity and makes other agents select a different route, avoiding everyone being stuck. Because nasty drivers also get information from outside their group, they find out the best route (from someone who does not lie) and eventually all go to either main (2/3) or side (1/3).

An interesting conclusion here is that nasty drivers, although not intentionally, provide the diversity necessary for the whole society to converge to the equilibrium, which is not in the best interest of those nasty drivers who give the false information (since this was actually intended to free their own routes!).

## 4.2 Scenarios with Recommendation

Now we discuss similar scenarios in which all drivers (except for noisy) receive a recommendation from the PMS in order to direct  $2/3$  of them to the main route. In case of 4A, there are no nasty or noisy drivers and all probabilities ( $p_1$ ,  $p_2$ , and  $p_3$ ) are set to zero. One simulation is shown in Fig. 6. The top most curve shows the expected: on average,  $2/3$  of the drivers use the main route.

More interesting are the cases where those probabilities are non-zero. We start with variations in the probability of deviating from the recommendation and change route (4B), i.e.  $p_3$  assuming values of 0.01, 0.1, 0.5, and 1.0, while  $p_1$  is set to 0.01. These results are also depicted in Fig. 6. With  $p_3 = 0.01$ , the number of drivers on main is not different from the above case. For  $p_3 = 0.1$ , the number of drivers on main goes down to around 95. Similar for  $p_3 = 0.5$  (75) and  $p_3 = 1.0$  (50). The case 4C is depicted in the same figure, only for  $p_2 = 0.2$  and  $p_3 = 0.01$  (others are similar). Asking people from other groups introduces more noise in the scenario (compared to 4A).

Summarizing case 4, the best situation happens of course when everybody follows the recommendation; however this is not realistic. The more noise there is (introduced by the probabilities  $p_3$ ), the less drivers deviate from equilibrium.

However, case 4 does not account for noisy or nasty drivers. The latter is studied in case 5. We have simulated case 5A with different number of nasty drivers, keeping  $p_1$ ,  $p_2$ , and  $p_3$  at zero. Nasty drivers also get recommendations but, when asked, give the wrong selection. In all cases the average fluctuates around 100 drivers on main because nasty drivers have low influence due to the recommendation. This is pretty much the case also regarding cases 5B and 5C. Neither probability  $p_1$  nor probability  $p_2$  affects the distribution which is in equilibrium.

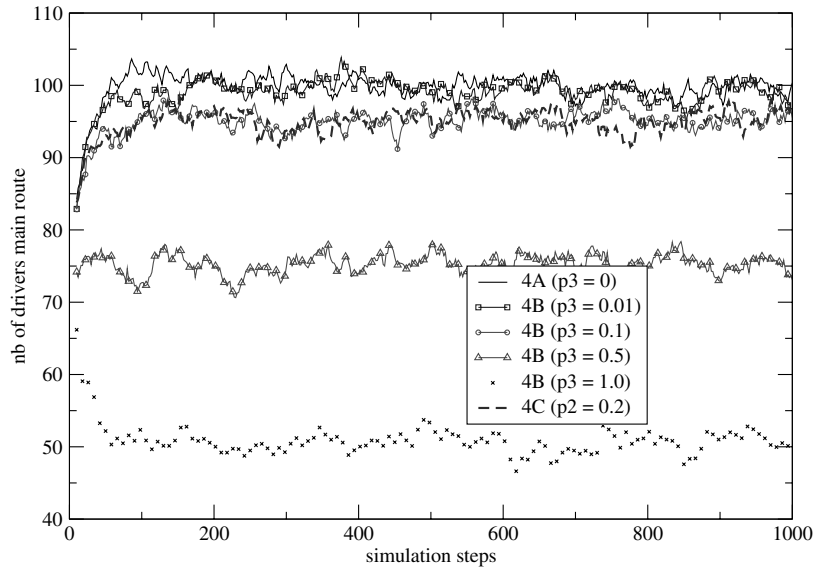
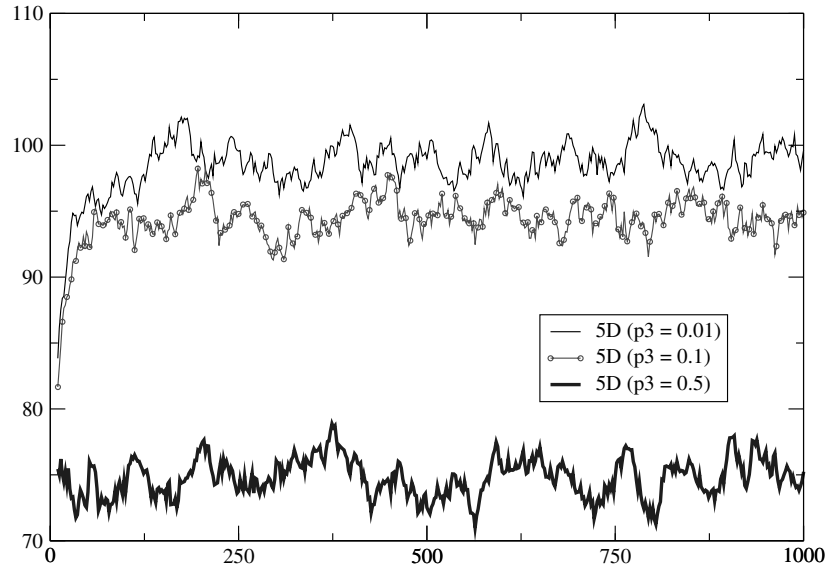
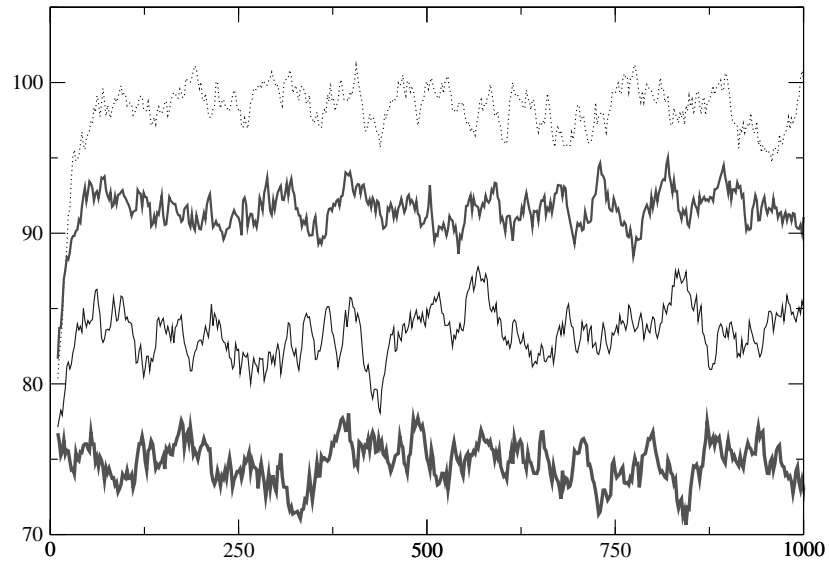


Fig. 6. Number of drivers on main for scenario 4

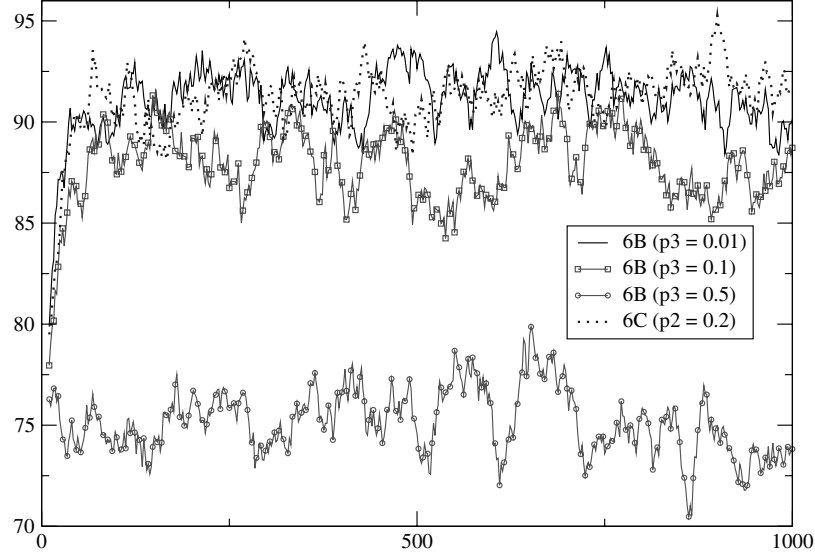


**Fig. 7.** Number of drivers on main for scenario 5D



**Fig. 8.** Number of drivers on main for scenario 6A, for 10, 50, 100, and 150 (top to bottom) noisy drivers

However, this changes when we introduce the probability  $p_3$  of drivers deviating from the recommendation. Fig. 7 shows this situation for different values of  $p_3$  (0.01, 0.1, and 0.5), while keeping  $p_1$  and  $p_2$  at zero. As expected, the number of drivers on main decreases and this happens the more drivers deviate from the recommendation.



**Fig. 9.** Number of drivers on main for scenarios 6B and 6C

With  $p_3 = 0.5$ , the overall behavior is similar to random selection, even if no noisy driver is present.

Scenario 6A includes noisy drivers, and  $p_1$ ,  $p_2$ , and  $p_3$  are zero. Fig. 8 shows the results of the simulations for number of noisy drivers equal to 10, 50, 100, and 150 (from top to down). As expected, in each case we have  $2/3$  of the social drivers who follow the recommendations and half of the noisy drivers in the main route, which destroys the equilibrium.

Now, when we add non-zero probability for  $p_3$  (6B), the number of drivers on main decreases further because some drivers deviate (Fig. 9). In this figure we also plot for  $p_2 = 0.2$  (6C).

### 4.3 Final Remarks About All Scenarios

Regarding the *performances of groups*, we discuss only cases which are far from the equilibrium because these situations indicate the performance of a given group.

Regarding the analysis of the simulations *without recommendation*, generally, the performance of *noisy drivers* is just boring: they tend to perform around the equilibrium when they are too many (see last entry on Table 1 whose average reward for main and side is  $\frac{5}{6} * \frac{1}{2} + \frac{1}{2} * \frac{1}{2} = 2/3$ ) or when they are too few (they have little influence). In other cases, they tend to destroy the equilibrium pattern, pushing the number of drivers on main below 100.

In the case of *nasty drivers*, their best performance happens for 3A with 10 nasty drivers since only half of these few go to  $S$  because of the false information they give inside their group. With the increasing number of nasty drivers, their performance as a group drops. In fact, their worst performance is in case 3A with 140 nasty drivers, followed by 3A with 100 nasty.

In the case of *social drivers*, their best performance is case 3A when there are 140 nasty drivers (which, not surprisingly is the worst performance for these). All of the nasty drivers go to secondary route most of the time, leaving the main route for the few social ones.

The worst performance for social people are 1A and 1C because they all copy the best route in the group and select the same route. In 2A, the performance of social drivers is bad only when there are few noisy drivers; when there are many, these select randomly leaving less people on main, thus increasing the reward of the social drivers. 3A is also bad for social drivers because the few nasty drivers do very well as explained above.

Regarding the performance for the situations *with recommendation*, the observations above do not apply (except for noisy drivers). Regarding nasty drivers, in no case they do very good or very bad. There are high fluctuations over the average for case 5D when  $p_3$  is 0.5 because drivers deviate from the recommendation frequently. Social drivers do well in case 6A, when there is a high number of noisy drivers (again, these select randomly leaving less people on main). They have a low performance in case 4B when  $p_3$  is 1.0 because they all deviate from the recommendation, thus ending up in a situation opposed to the equilibrium.

## 5 Conclusions and Future Work

This paper discusses the effect of information sharing in a scenario of adaptation regarding which route to take in a route choice scenario. The simulations of the situations described in Section 3 show that it is interesting to have a system giving recommendations to the drivers. However, the performance of the groups decreases when too many drivers deviate from the recommendation, which seems to be a current practice regarding driving because people not always trust the recommendation given bad past experiences or they just want to experiment new routes.

Also, when there is no social attachment and the behavior is myopic towards maximization of short time utility, the performance is bad for nasty drivers (at least), except when they are few. Being a nasty driver pays off up to a certain level only. Another point that can be stressed is that in scenario 3C nasty drivers help the social ones, even if this is not intended.

Regarding the noisy drivers, these tend to destroy good patterns of equilibrium (and hence performance). However they also help social drivers when these are stuck in a bad choice in situations, especially when the probability of meeting acquaintances is non-zero. Therefore, the recommendation system must consider the ratio of noisy drivers. When the number of noisy drivers is unknown or too high, the recommendation may fail.

Finally, this work is based on a series of assumptions that may not be bearable in every real world application. For instance, we assume a global control component that is able to compute the exact utility of the agent decisions for producing the recommendation. In [4], we investigate a scenario in which the control system has imperfect information.

In the future, we plan to also investigate the influence of group size. Also, we will add an adaptation component to the social and nasty drivers, as well as to the PMS so that they can adapt to the changing conditions, with an eventual recognition of nasty drivers in order to mark them as untrustable as information source.

## Acknowledgments

The first author is partially supported by CNPq. We also thank the anonymous reviewers for their valuable suggestions.

## References

1. B. Arthur. Inductive reasoning, bounded rationality and the bar problem. Technical Report 94-03-014, Santa Fe Institute, 1994.
2. A. L. C. Bazzan, R. H. Bordini, G. Andriotti, R. Viccari, and J. Wahle. Wayward agents in a commuting scenario (personalities in the minority game). In *Proc. of the Int. Conf. on Multi-Agent Systems (ICMAS)*. IEEE Computer Science, July 2000.
3. A. L. C. Bazzan and A. P. Cavalheiro. Influence of social attachment in a small-world network of agents playing the iterated prisoners dilemma. In S. Parsons and P. Gmytrasiewicz, editors, *5th Workshop of Game Theoretic and Decision Theoretic Agents*, pages 17–24, July 2003. held together with AAMAS 2003.
4. A. L. C. Bazzan and R. Junges. Congestion tolls as utility alignment between agent and system optimum. In *Proceedings of the Fifth Int. Joint Conference on Autonomous Agents and Multiagent Systems*, 2006. submitted to AAMAS 2006.
5. D. Challet and Y. C. Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica A*, 246:407–418, 1997.
6. F. Klügl and A. L. C. Bazzan. Route decision behaviour in a commuting scenario. *Journal of Artificial Societies and Social Simulation*, 7(1), 2004.
7. F. Klügl, A. L. C. Bazzan, and J. Wahle. Selection of information types based on personal utility - a testbed for traffic information markets. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 377–384, Melbourne, Australia, July 2003. ACM Press.
8. S. Milgram. The small world problem. *Psychol. Today*, 2, 1967.
9. J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers*, volume 2, pages 325–378, 1952.
10. D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684):397–498, June 1998.