Increasing Information in Socio-Technical MAS Considered Contentious

Vivek Nallur*, Julien Monteil*, Tyler Sammons†, Mélanie Bouroche*, Siobhán Clarke*
*School of Computer Science and Statistics, Trinity College Dublin, Ireland
{vivek.nallur|julien.monteil}@scss.tcd.ie
†University of Notre Dame, Indiana, USA
Tyler.S.Sammons.6@nd.edu

***Abstract*—Socio-technical systems differ from typical MAS formulations in that efficiency of the system is not the only concern of the participating agents. Human attributes such as concern for social equity, lying and irrationality are also present, alongside the normal computation being undertaken by the agents. Typically, non-rational attributes are considered as noise and therefore not considered as important attributes of the system. In this paper, we consider noise in socio-technical systems and show that the typical reaction of increasing information to counter noise is ineffective. We show that endowing agents with increased memory and/or computational power is not necessarily beneficial to achieving the goal of the system.**

I. Introduction

Self-adaptive socio-technical systems are all around us, though not explicitly recognized as such. From electronic marketplaces where participants are inherently self-interested (e.g., ebay, amazon, etc.), to P2P networks where participants share bandwidth (e.g., torrents), to volunteer computing networks where participants donate time and computing resources(e.g., Folding@Home, Seti@Home, Milkyway@Home, etc.) these systems combine algorithmic mechanisms to optimize some parameter, with a variety of human-targetted features to encourage participation. The success of these systems depends on multiple factors, like how efficiently the system reaches the desired goal, or how motivated the participating humans are. Modelling such systems formally is usually difficult due to the complexity of behaviour of the participating agents. Game theoretic approaches which offer analytic solutions, assume complete rationality on behalf of the agents, which are often hard to reconcile with real/observed autonomic behaviour. Multi-Agent Systems (MAS) are frequently used as a mechanism to get around such assumptions, and also to support computation at large scales (e.g., city-scale modelling).

Creating agents that reasonably approximate the actual behaviour of the individual in the system is a major challenge, in the design of MAS solutions. Agents, while autonomous, must possess the right amount of information such that they are able to compute solutions to their problems. Adding more information, in the form of increased memory or sophisticated computational strategies, is the intuitive approach to overcoming this challenge, in competitive situations. For example, in the area of agent-based auctions, Gode and Sunder had shown that agents with *Zero Intelligence* could achieve equilibrium prices [1]. However, subsequent works almost universally add information, either in the form of increased memory or sophisticated strategy [2]–[4]. In the context of MAS, we define a socio-technical system as *a MAS system which contains a mix of humans and computational devices, jointly making decisions*. The computational devices provide fast, calculating ability along with rational choices, while the humans influence the actual decisions. Due to this, socio-technical systems may exhibit non-rational behaviour in the form of spontaneous cooperation, lying, etc. These behaviours are not uniformly present amongst all agents in a socio-technical system. Some individuals may decide to cooperate at a particular point in time, and then shift to non-cooperation. Furthermore, social equity amongst the rewards achieved by the agents, is a particular concern in socio-technical systems. Social equity is roughly defined as the opportunity for each participating agent to grab some rewards, as opposed to some MAS where it does not matter if the optimal solution involves a majority of agents losing out, at the individual level. When both rational as well as social behaviour are present, the intuitive approach of adding information to agents to enable better performance, is actually counter-productive. In this paper, we show through simulations of two very different systems, that increasing the amount of information available to an agent is not directly proportional to increase in system performance. We use the *Minority Game* [5], a game involving self-coordination, and Vehicular Networks (VN) as our exemplars.

II. Minority Game

The Minority Game (MG), introduced by Challet and Zhang, consists of an odd number (N) of agents, playing a game in an iterated fashion. At each timestep, each agent chooses from one of two actions (+1 or 0), and the group that is in the minority, wins. Since N is an odd integer, there is guaranteed to be a minority. Winners receive a reward, while losers receive nothing. After each round, all agents are told which action was in the minority. This feedback loop induces the agents to re-evaluate their action for the next iteration. While simple in its conception, this game has been used in many fields like econophysics [6]–[10], multi-agent resource allocation [11]–[13], emergence of cooperation [14], and heterogeneity [12], [15]. The Minority Game (MG) is intuitively applicable to many domains, such as traffic (the driver that chooses the 'less-travelled' road experiences less congestion), packet traffic in networks, ecologies of foraging animals, etc. The Minority Game has been used, primarily, to investigate the *efficiency* of strategies in the system when individuals have bounded rationality (for an extensive review, see [16]). We conceive of a Minority Game as a special (simplified) case of a socio-technical system.

Variable	Value
Memory	3
Strategies-per-Agent	8
Population Size	501
Simulation Period (rounds)	2000

TABLE I: Experimental constants

A. Description of Experiments

An agent in MG is characterized by two variables, m and s, where m refers to the memory of the agent, and s refers to the set of strategies used by the agent. The agent keeps score of which strategy performs best, and uses that strategy in conjunction with the history (the last m winning actions) to decide its next action. The agent is free to change its strategy, and associated action, after every iteration. In their paper [5], Challet and Zhang show that with a small m and s, the agents self-organize into an efficient state, where the number of winners (the minority) after some iterations, is greater than would be expected by a random throw of a coin. In our experiments, we measure the following two variables:

- **Mean Dispersion**: the dispersion of payoff is the absolute difference between the number that were in the majority and the number that were in minority. The ideal dispersion value per iteration is 1, when the size of the majority is one more than the size of the minority. In reality, the system tends to start at some high value and continues fluctuating around a mean value without ever settling down. The *mean value of dispersion* over a large number of iterations is therefore, used to measure the efficiency of the system.
- **Social Payoff**: the social payoff models the dispersion of rewards amongst the agents. This is so because we do not want a situation where the same group of agents continuously choose one or the other action. This could lead to a very efficient system, but would also be unfair to the individual agent. In an equitable system, we would like each agent to have an equal chance of being in the minority. We measure social payoff, as the rewards achieved by each agent. In the ideal case, all agents would achieve a high mean social payoff, with a standard deviation of zero. In reality, the mean social payoff varies highly, with a high value of standard deviation.

Together, they represent the efficiency and the equity of a self-coordinating MAS. Agents in the MG vary over m and s. Over all the experiments, we keep the variables constant, as shown in Table I. Each datapoint on each graph is a mean of data from 10 simulations.

B. Addition of Cooperation

We model the introduction of cooperation into MG, by introducing the concept of a logical neighbourhood. An agent is randomly assigned a set of neighbouring agents, and after every iteration, the agent is able to query its neighbours about their next action. Depending on the responses of the neighbours, the agent is able to choose the minority action. Clearly, the greater the size of the neighbourhood, the greater should be the ability of the agent to predict correctly.

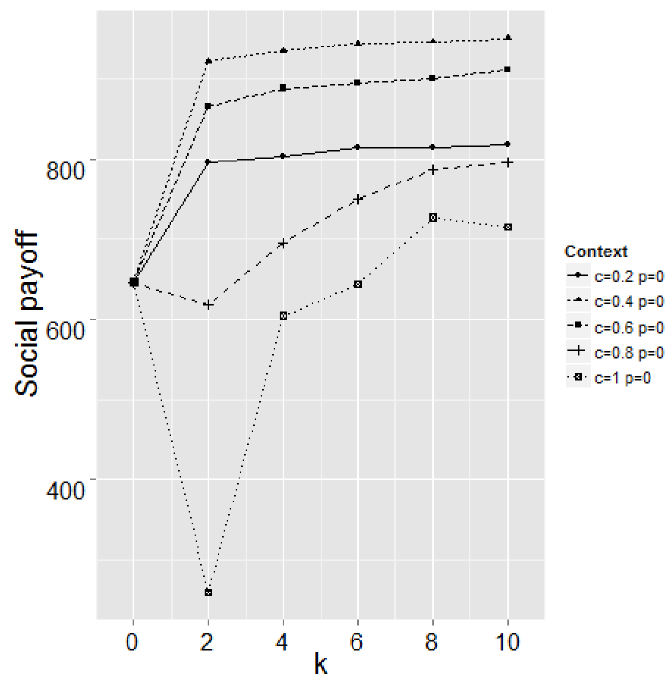

Fig. 1: Effect on social payoff as cooperation increases

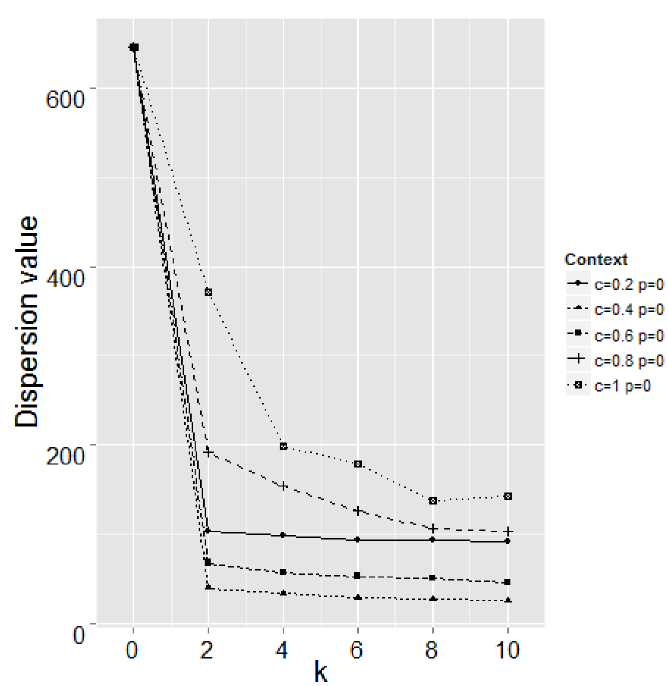

Fig. 2: Effect on mean dispersion as cooperation increases

In Figure 1, each line represents the percentage of the population that is able to cooperate, *i.e.*, `c = 0.2` means that `20%` of the population is able to cooperate and query other agents, while the rest are not. `c` varies from `0.2 to 1`. `k` represents the number of neighbours that each agent is able to query. `k` varies from `2 to 10`. We see that the social payoff initially (at `k = 2`) increases for three values (`c = 0.2, 0.4, 0.6`), but drops for two values (`c = 0.8, 1`). However, surprisingly, as `k` increases to `4,6,8,10`, social payoff decreases for the first three values (`c = 0.2, 0.4, 0.6`), but increases for the latter two (`c = 0.8, 1`). Figure 2 shows the efficiency of the system, represented by the mean dispersion. As expected, the addition of cooperation increases the efficiency of the system with all levels of cooperation. Each level also exhibits a plateau in terms of the value of k, where adding more cooperation does not increase the efficiency any more. Therefore, for the rest of this paper, we only show social payoff, that fluctuates with addition of more information into the system.

C. Addition of Lying

The addition of cooperation introduces more information into the system. However, in a socio-technical system it is very conceivable that an agent could lie about its intentions to gain a competitive advantage. In MG, if an agent is able to successfully mislead its neighbours into choosing the wrong group, it increases the probability of increasing its own payoff. Thus, lying in a social context is a very rational act for an MG-agent. We model what happens when cooperating agents have a tendency to lie, *i.e.*, if according to its best strategy, the next chosen action would be `1`, upon querying the agent would lie (and reply with a `0`), with a certain probability `p`.

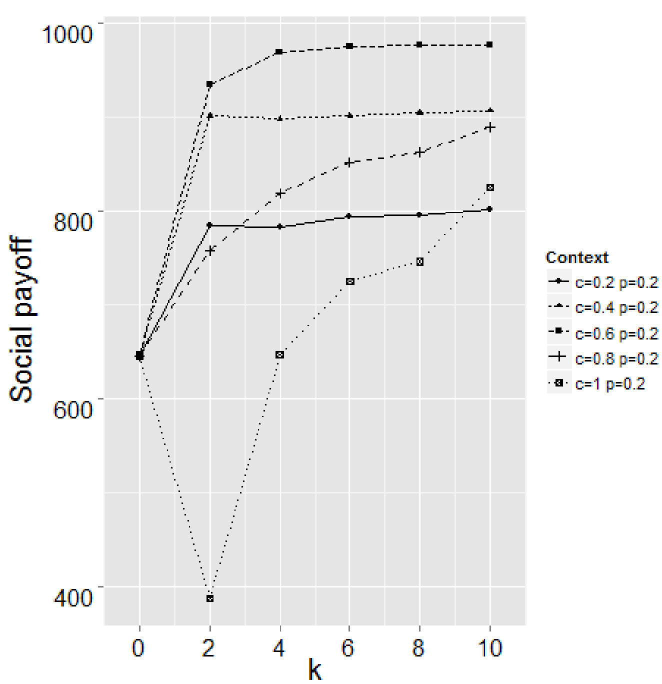

Fig. 3: Social payoff with a small probability to lie (0.2)

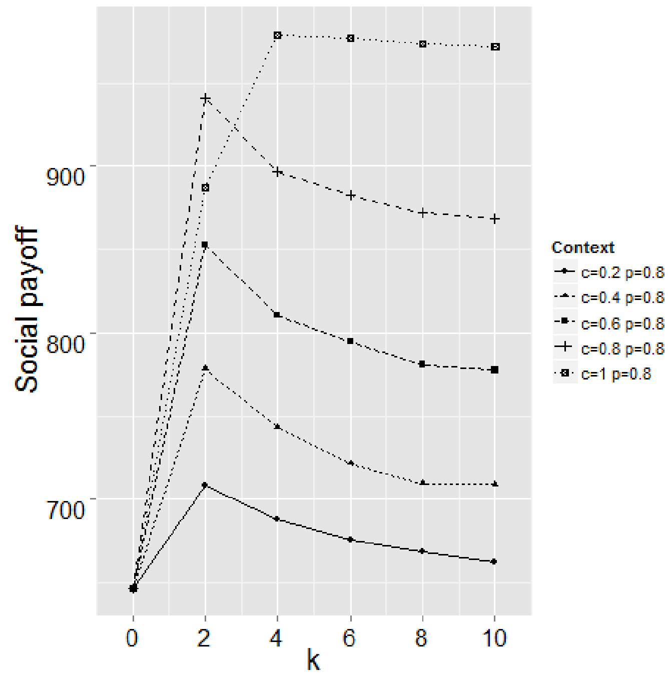

Fig. 4: Social payoff with a higher probability to lie (0.8)

Figures 3 and 4 show the change in social payoff, with each agent now having the ability to lie. Surprisingly, in Figure 3, almost all values of k result in the social payoff being above 800, whereas without lying (in Figure 1), only three manage to reach the 800 mark. In Figure 4 shows an initial rise, with a subsequent fall. The fall is explained by the fact that when each agent is lying with a probability of 0.8, the amount of correct information in the system is actually falling quite low.

D. Heterogeneous Population

It is slightly unrealistic to have a population of agents that all lie with the same probability. Now, we model a more heterogeneous population of agents, *i.e.*, a situation where the presence of lying is not uniform, rather each cooperating agent has a `p` drawn from a gaussian distribution. This adds more realism to the multi-agent population.

Figure 5 shows the fluctuation in social payoff when agents lie with a probability drawn from a Gaussian distribution centred on 0.5 in order to cover the full range of values from 0 to 1. One very interesting interesting feature in all the figures is that the optimal value of `k` changes, not only with increase in neighbourhood size, but also within percentage of population. That is, for a given population, change in cooperation (c) and lying (p) causes the optimal value of `k` to change.

III. Vehicular Networks

The study of the minority game can be related to another socio-technical system, Vehicular Networks (VN). In VN, some vehicles, said to be cooperative, are equipped with embedded sensors and are able to exchange information with other vehicles, and can be controlled in order to increase the global traffic safety, and efficiency of the transportation

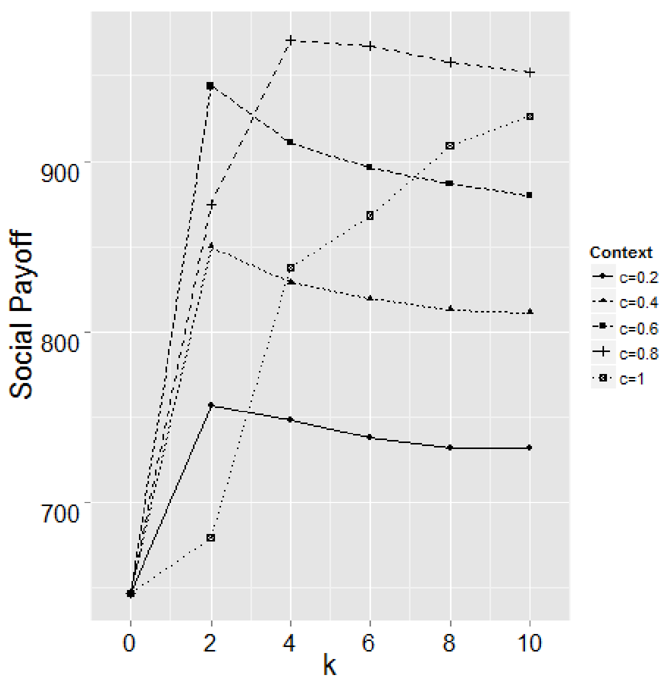

Fig. 5: Effect on social payoff when agents are heterogeneous

network. The rest of the vehicles are simply operated by drivers and adopt a very specific driving behaviour. The resulting system is a complex, mixed-traffic environment where vehicle-sensors, communication, and driver-behaviour, all contribute to the dynamics of the system. Such Vehicular Networks are currently the subject of a lot of research, as the economic benefit of autonomous and connected vehicles is assessed to be enormous (worth £51 billion in the UK, by 2030 [17]).

In such Vehicular Networks, an agent can be represented as an entity encompassing a driver and the vehicle, providing the agent with behavioural capabilities (the way the driver drives), as well as with technical capabilities (engine power, embedded communication devices and sensors). There are two types of agents: non-cooperative agents and cooperative ones. The classical approach used to model vehicle dynamics of non-cooperative agents at a microscopic level, is to use car-following models (see [18]). The acceleration of a vehicle is written as a non-linear function of its speed, relative distance and relative speed to the leading vehicle, as well as behavioural parameters Θ specific to the drivers (maximum tolerated acceleration, desired speeds, etc.). This is comparable with the Minority Game, except that the agent does not have a fixed number of strategies in mind, but will link its perception to the decision in a continuous way. Cooperative agents act based on enhanced perception available via sensors and communication exchanges. There are different ways of incorporating cooperation at microscopic levels. In this paper we chose the multi-anticipation [19], [20] approach. Multi-anticipation refers to the situation where the acceleration of a driver is computed as a function of multiple leaders and followers, as opposed to only one leader such as in the traditional approach [21]. Some results showed that multi-anticipation can help reduce traffic congestion and remove traffic instabilities [22].

A. *Description of Experiments*

In our experiments, an agent is distinguished by the strategy it follows. If cooperative, it follows the multi-anticipative model, whereas if non-cooperative, it follows the usual (and well-known) Intelligent Driver car-following Model (IDM) [21]. Additionally, the MOBIL lane-changing model overrides the car-following behaviour when a lane change is judged useful by the agent. Note that any other realistic car-following or lane-changing model could have been chosen for this analysis. The agent is characterized by its behavioural parameters $\Theta = (V_{\max}, \mathtt{T}, \mathtt{a}, \mathtt{b}, s_0)$, where $V_{\max}$ is the desired speed, $\mathtt{T}$ is the desired safety time headway, $\mathtt{a}$ is the maximum acceleration, $\mathtt{b}$ is the comfortable deceleration, s_0 is the minimum net stopped distance from the leader. Those parameters account for the behaviour of the driver and the technical capabilities of the vehicle. In our experiments, by analogy with the MG, we measure the following variable:

Social payoff: Social payoff is defined as the inverse of the standard deviation of the speeds. It can be seen as the equivalent of the defined social payoff in the MG, as a low standard deviation corresponds to homogeneous speeds of agents along the network, which relates to equity in terms of experienced travel times and speed trajectories. High values of standard deviation correspond to high disturbances in the traffic stream, meaning that agents can have a very different experience depending on whether they are caught in congestion or not. This indicator both represent the efficiency and the equity in the VN Multi-Agent System.

Variable	**Value**
Population Size	500
Number of Lanes	2
Simulation Period (rounds)	900

TABLE II: Experimental constants

Over all the experiments, the variables shown in Table II were kept constant. Each datapoint on each graph is a mean of data from 30 simulations. Agents in the VN vary over their behavioural parameters. This can be related to the probability to lie in the Minority Game, which can be seen as a behavioural parameter. The neighbourhood size is assumed to only take even numbers: $\mathtt{k} = 0$ corresponds to the classic car-following law (one leader no follower), $\mathtt{k} = 2$ corresponds to two leaders and one follower, and so on).

B. *Results*

As mentioned previously, the multi-anticipative strategy is a strategy of cooperation that is known to help increase traffic flow efficiency. Positive outcomes are expected when implementing this strategy. However, we will show that depending on the percentage of cooperative agents in the fleet, the optimal number of the neighbourhood size $\mathtt{k}$ changes, and that when introducing variability, an increase of cooperation does not necessarily mean a system with better social payoff.

1) Homogeneous Population: The behavioural parameters of the agents were chosen to be identical, and defined as in a realistic traffic, see [23]. In Figure 6, we see that the optimal payoff is obtained for $k = 4$. After that threshold, adding information does not lead to a better social payoff.

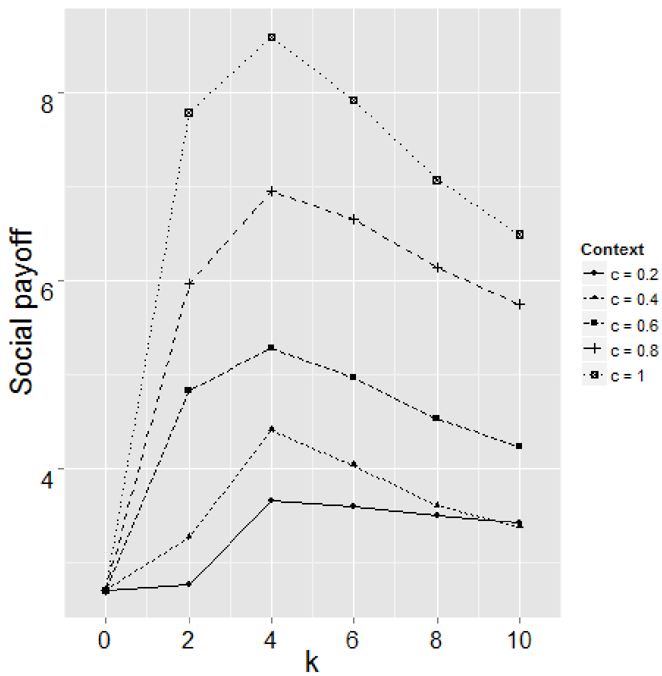

Fig. 6: Vehicular Network: agents are homogeneous

2) Heterogeneous Population: in order to consider some variability in Vehicular Networks, the distribution of the behavioural parameters of agents were drawn, like in the MG, from a Gaussian distribution centred on the realistic parameters values defined earlier.

The first consideration resulting from figure 7 is that the multi-anticipative strategy does not work well when there is a small proportion of cooperative agents. The strategy with 20% of cooperative agents has a worse performance than the strategy without multi-anticipation. The strategy with 100% of cooperative agents shows a relatively stable social payoff for k varying from 4 to 8. This leads to the following discussion: adding information can be detrimental when the variability between the agents is high, or simply when this variability is realistic. *The variability in agent behaviour causes good strategies, that were analytically proven to perform well, to perform badly.*

The second consideration is the one that was observed in the Minority Game. There is an optimal k for each given percentage of population participating in cooperation$c$, given the variability of behaviours and given the implemented strategy to incorporate cooperation (multi-anticipation in this study). This can be explained: in a highly physics-dependent system such as VN, considering too much information can blur the knowledge of the agents and an agent will end up giving too much weight to unimportant neighbours, e.g. neighbours which are too far away from its immediate surroundings.

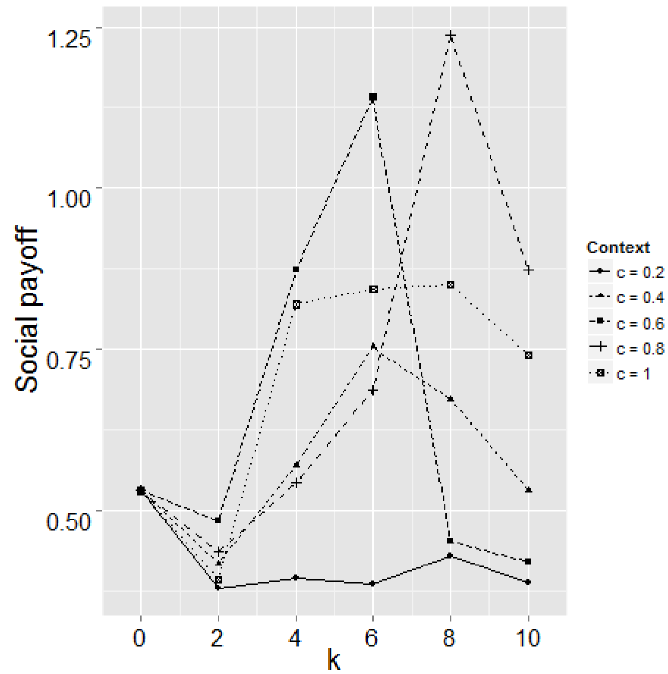

Fig. 7: Vehicular Network: agents are heterogeneous

As vehicular networks are likely to face low levels of cooperation at first, those results are of critical importance. Cooperation strategies that are analytically designed, must be tested in a non-deterministic framework which takes into account the variability of agent behaviours. An ideal cooperation strategy must not display such variations in the global social payoff.

IV. Observations

Both case-studies showcase the existence of two different kinds of social variability, that naturally exists amongst agents:

- *Variability in Sociability / Cooperation*: In a social MAS, agents may or may not cooperate with other agents. In the case of cooperation, agents may also be distributed unevenly (either logically or spatially), which could affect the number of neighbours they have access to. A uniform policy/protocol for deriving information could lead to severely different conclusions being reached by different agents. The more the amount of information is being considered, the worse the possible divergence is between the agents.
- *Variability in Individual Behaviours*: Agents in a social MAS could also comprise a heterogeneous population, specially in domains such as vehicular networks, where agents have multiple behavioural and technical parameters. Variability could range from the unintentional (computational limitation) to accidental (malfunctioning sensor) to the deliberate (lying). This means that information acquired from neighbouring agents must be treated with caution, and cross-verified wherever possible. In this scenario, increasing the

amount of information derived from neighbouring agents, acts in a counter-productive manner.

V. Conclusions

This paper presents the results of adding socio-aware behaviour to two multi-agent systems, the Minority Game and Vehicular Networks, that were previously known to achieve efficient outcomes. In previous work on both kinds of multi-agent systems, work has focused on making the system reach an efficient outcome, using more computational power or sophisticated strategies. We show that the presence of socio-aware behaviour acts as a confounding factor, *i.e.,* adding more information or computational power does not necessarily result in a better outcome, for a given strategy of cooperation. In our experiments, there exists a sweet spot of adding information where both performance and social outcomes are better. But adding more information can be detrimental, and this is related to the human variability in a socio-technical system. Socio-technical systems of the future must consider social aspects of sharing information, and the probabilistic nature of trustworthiness. For future work, we will investigate multiple distributions of socio-aware agents in a population and attempt to quantify and correlate the impact of these distributions on outcomes achieved.

Acknowledgement

This work was partially supported by the EU FET-Project Diversify FP7-ICT-2011-9, and by Science Foundation Ireland grant 10/CE/11855 to Lero - the Irish Software Engineering Research Centre (www.lero.ie).

References

[1] D. K. Gode and S. Sunder, "Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality," *The Journal of Political Economy*, vol. 101, no. 1, pp. 119–137, 1993. [Online]. Available: http://www.jstor.org/stable/2138676

[2] D. Cliff, "Simple bargaining agents for decentralized market-based control," *HP Laboratories Technical Report*, 1998. [Online]. Available: http://www.hpl.hp.com/techreports/98/HPL-98-17.pdf

[3] S. Gjerstad and J. Dickhaut, "Price formation in double auctions," *E-Commerce Agents*, pp. 106–134, 2001. [Online]. Available: http://www.springerlink.com/index/3m4fnyp3hcl2yglq.pdf

[4] A. E. Roth and I. Erev, "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term," *Games and Economic Behavior*, vol. 8, no. 1, pp. 164–212, 1995. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S089982560580020X

[5] D. Challet and Y.-C. Zhang, "Emergence of cooperation and organization in an evolutionary game," *Physica A: Statistical Mechanics and its Applications*, vol. 246, no. 34, pp. 407 – 418, 1997. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378437197004196

[6] G. Bianconi, T. Galla, M. Marsili, and P. Pin, "Effects of Tobin taxes in minority game markets," *Journal of Economic Behavior & Organization*, vol. 70, no. 1-2, pp. 231–240, May 2009. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167268108002084

[7] D. Challet, "Inter-pattern speculation: Beyond minority, majority and $-games," *Journal of Economic Dynamics and Control*, vol. 32, no. 1, pp. 85–100, Jan. 2008. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0165188907000413

[8] T. Patkowski and M. Ramsza, "Multi-market minority game," *Advances in Complex Systems (ACS)*, vol. 08, no. 01, pp. 65–74, 2005. [Online]. Available: http://EconPapers.repec.org/RePEc:wsi:acsxxx:v:08:y:2005:i:01:p:65-74

[9] C. A. Zapart, "On entropy, financial markets and minority games," *Physica A: Statistical Mechanics and its Applications*, vol. 388, no. 7, pp. 1157–1172, Apr. 2009. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378437108009801

[10] D. H. Wolpert and J. W. Lawson, "Designing agent collectives for systems with markovian dynamics," *Proceedings of the first international joint conference on Autonomous agents and multiagent systems part 3 - AAMAS '02*, p. 1066, 2002. [Online]. Available: http://portal.acm.org/citation.cfm?doid=545056.545074

[11] Y. Li and R. Savit, "Toward a theory of local resource competition: the minority game with private information," *Physica A: Statistical Mechanics and its Applications*, vol. 335, no. 1-2, pp. 217–239, Apr. 2004. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378437103011348

[12] B. A. Mello and D. O. Cajueiro, "Minority games, diversity, cooperativity and the concept of intelligence," *Physica A: Statistical Mechanics and its Applications*, vol. 387, no. 2-3, pp. 557–566, Jan. 2008. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378437107009995

[13] R. Metzler and C. Horn, "Evolutionary minority games: the benefits of imitation," *Physica A: Statistical Mechanics and its Applications*, vol. 329, no. 3-4, pp. 484–498, Nov. 2003. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378437103006265

[14] D. Dhar, V. Sasidevan, and B. K. Chakrabarti, "Emergent cooperation amongst competing agents in minority games," *Physica A: Statistical Mechanics and its Applications*, vol. 390, no. 20, pp. 3477–3485, Oct. 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378437111003876

[15] G. W. Greenwood, "Deceptive strategies for the evolutionary minority game," in *Proceedings of the 5th International Conference on Computational Intelligence and Games*, ser. CIG'09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 25–31. [Online]. Available: http://dl.acm.org/citation.cfm?id=1719293.1719308

[16] E. Moro, "The minority game: an introductory guide," in *Advances in Condensed Matter and Statistical Physics*, E. Korutcheva and R. Cuerno, Eds. Nova Science Publishers, 2003, pp. 263–286. [Online]. Available: http://arxiv.org/abs/cond-mat/0402651

[17] KPMG, "Connected and autonomous vehicles - the UK economic opportunity -," March 2015.

[18] M. Brackstone and M. McDonald, "Car-following: a historical review," *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 2, no. 4, pp. 181 – 196, 1999. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S136984780000005X

[19] S. P. Hoogendoon, S. Ossen, and M. Schreuder, "Emprirics of multi-anticipative car-following behavior," *Transportation Research Record: Journal of the Transportation Research Board*, no. 1965, pp. 112–120, 2006. [Online]. Available: http://pubsindex.trb.org/view.aspx?id=776914

[20] H. Lenz, C. Wagner, and R. Sollacher, "Multi-anticipative car-following model," *The European Physical Journal B - Condensed Matter and Complex Systems*, vol. 7, no. 2, pp. 331–335, 1999. [Online]. Available: http://dx.doi.org/10.1007/s100510050618

[21] T. M. Kesting, A. and D. Helbing, "Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity," *Philosophical Transactions of the Royal Society*, vol. 368, no. 1928, pp. 4585–4605, 2010. [Online]. Available: http://dx.doi.org/10.1007/s100510050618

[22] J. Monteil, R. Billot, J. Sau, and N.-E. El Faouzi, "Linear and weakly nonlinear stability analyses of cooperative car-following models," *Intelligent Transportation Systems, IEEE Transactions on*, vol. 15, no. 5, pp. 2001–2013, Oct 2014.

[23] V. Punzo, M. Montanino, and B. Ciuffo, "Do we really need to calibrate all the parameters? variance-based sensitivity analysis to simplify microscopic traffic flow models," *Intelligent Transportation Systems, IEEE Transactions on*, vol. 16, no. 1, pp. 184–193, Feb 2015.