An agent-based model of the emergence of cooperation and a fair and stable system optimum using ATIS on a simple road network☆

Ido Klein, Nadav Levy, Eran Ben-Elia*

GAMES Lab, Department of Geography and Environmental Development, Ben-Gurion University of the Negev, Israel

ARTICLE INFO

ABSTRACT

Traffic congestion threats the growth and vitality of cities. Policy measures like punishments or rewards often fail to create a long term remedy. The rise of Information and Communication Technologies (ICT) enable provision of travel information through advanced traveler information systems (ATIS). Current ATIS based on shortest path routing might expedite traffic to converge towards the suboptimal User Equilibrium (UE) state. We consider that ATIS can persuade drivers to cooperate, pushing the road network in the long run towards the System Optimum (SO) instead. We develop an agent based model that simulates day-to-day evolution of road traffic on a simple binary road network, where the behavior of agents is reinforced by their previous experiences. Scenarios are generated based on various network designs, information recommendation allocations and incentive mechanisms and tested regarding efficiency, stability and equity criteria. Results show that agents learn to cooperate without incentives, but this is highly sensitive to the type of recommendation allocation and network-specific design. Punishment or rewards are useful incentives, especially when cooperation between agents requires them to change behavior against their natural tendencies. The resulting system optimal states are to most parts efficient, stable and not least equitable. The implications for future ATIS design and operations are further discussed.

1. Introduction

More than 53% of humanity (74% in Europe) lives in cities, and this figure will only increase, especially in developing and transition economies (World Bank, 2016). People are attracted to cities because of the benefits of agglomeration – access to a wide range of services and employment compared to rural areas. However, it is becoming harder to maintain road traffic in smooth working order. Traffic congestion is a sign of a city's vitality, but it also incurs negative externalities, such as time losses and delays, air pollution, noise and decreasing safety (Mayeres et al., 1996). Non-recurring conditions, like accidents or bad weather, further contribute to exacerbating congestion on the roads (Schrank et al., 2012). As more people are attracted to cities, future traffic congestion levels are unlikely to decrease.

Traffic and transport engineers traditionally believed that increasing road capacity would solve congestion problems. However, road-space expansion is fiscally costly and strengthens car-dependency while weakening public transport attractiveness and in the long run increasing urban sprawl (Goodwin, 1996; Guiliano, 1995). Public transport, though more environmentally sustainable than car travel is not able to provide services to anyone to anywhere and anytime within acceptable operational and budget thresholds. Evidence suggests that a notable share of driving is by choice (Handy et al., 2005), while car image retains strong symbolic and

---

affective values (Steg, 2005). Thus, cars continue to be the most popular mode of transport.

While private cars maximize personal mobility, they remain costly from a social viewpoint. Thus, similar to Hardin's 'Tragedy of the Commons' (Hardin, 1968), we can claim a *"social mobility dilemma"* - too many cars on too little road space, ending up in clogged traffic. Different strategies have attempted to discourage car travel to solve this social mobility dilemma. Bamberg and Schmidt (2003) distinguish between soft strategies that focus primarily on influencing travelers' attitudes, perceptions and motivations in order to change behavior and hard ones aiming at influencing behavior by changing the physical, organizational or financial conditions of car travel.

Hard strategies like road pricing are based on microeconomic foundations involving negative economic incentives that are designed to internalize the marginal costs of congestion that car users naturally overlook (Rouwendal and Verhoef, 2006; Vickrey, 1969). However, the lessons learned from the implementation of congestion charges, in places like Singapore, Stockholm and London confirm that to be effective pricing interventions have to be very strong. For example, being initially high, the congestion charge in London has further spiked considerably since its inception a decade ago (Timms, 2013). Unsurprisingly, congestion charging remains controversial regarding social equity and fairness, as well as public and political acceptability in liberal democratic societies (Eriksson et al., 2006; Viegas, 2001). Research based on cognitive psychology theories (Geller, 1989; Kahneman and Tversky, 1984) suggests that positive incentives – or rewards – can efficiently substitute charging (Ettema et al., 2010). However, Ben-Elia and Ettema (2011) asserted that rewards encourage mainly temporal shifts in car use and can be regarded as an unfair and costly subsidy of car users. Unsurprisingly, rewards have mainly been implemented as temporary relief strategies (e.g. for planned road closures in the Netherlands), rather than to be maintained in perpetuity like congestion charges.

With the proliferation of information and communications technologies (ICT), soft strategies aimed at voluntary behavior changes based on the provision of travel information, have become increasingly popular (Chorus et al., 2006; Davies, 2012). Travelers can thus make better mobility choice, that will likely improve the level of service on the entire transport system (Ettema and Timmermans, 2006; Levinson, 2003). In reality information is likely providing individual travelers with an illusionary sense of self-control over ambiguous travel conditions (Kemel and Paraschiv, 2013), while their perceived trust in travel information remains relatively high (Ben-Elia et al., 2013, 2008; Ben-Elia and Shiftan, 2010). Consequently, many travelers are developing a large degree of dependence on information for their everyday mobility decisions.

More importantly, the collective outcome of information dependency seems to expedite traffic to converge towards the sub-optimal state of User Equilibrium (UE). UE occurs when the trip time distribution is steady and travelers cannot improve their travel time by switching routes (Arnott et al., 1996, 1993, 1991; Emmerink et al., 1998, 1995a, 1995b). Traffic theory asserts that UE is a collective outcome of the non-coordinated choices of fully informed agents competing for shortest paths on a decentralized network of limited capacity (Arnott et al., 1993). With the proliferation of real-time crowdsourced routing apps (like Waze[©]), car traffic could be approaching a state of User Equilibrium more than ever before. Several experimental studies Lu et al. (2011, 2014), Selten et al., 2007) attest that fuller information brings about a social outcome that is similar if not exactly replicating UE. Moreover, some researchers claim that real-time navigation applications are leading to a state possibly worse than UE (Varga, 2014a). Depending on the network design, traffic congestion in a UE state might be worse off for travelers, in particular when all routes are adversely affected by the same global factor, such as bad weather (Lindsey et al., 2014).

Traffic theory relates the first principle of Wardrop (1952) as a working assumption that corresponds analytically to UE. Various discrete choice models for explaining route-choice were developed in order to extend Wardrop's first principle (for a review see Prashker and Bekhor (2004) as well as in multiple multiplayer economic experiments (e.g. Rapoport et al., 2014, 2009; Selten et al., 2007). However, it is actually Wardrop's second principle that considers a secondary state, where the aggregate travel times (or costs) are minimized. This system-optimum (SO) remains elusive - merely a theoretical concept in the eyes of many transportation researchers and practitioners. Achieving this SO has been regarded mostly as impractical as apart for certain corner solutions, a road network hardly ever arrives naturally at a global optimum without resorting back to hard policy strategies such as fines and tolls (Rothengatter, 1982). However, simulations and experiments on SO-based traffic assignment in various road networks and regarding the detrimental effects of pre-trip information on network performance illustrate the importance of bridging the gap between the aggregate travel time in UE and in SO (Jahn et al., 2005; Rapoport et al., 2014).

Nowadays, the sharing economy is allowing cooperation to emerge between large number of people to collaboratively use resources in a way that minimizes the cost of both consumption and coordination, using ICT as a medium (Kaplan and Haenlein, 2010). Sharing economy implementations from the field of transportation include ridesharing and carsharing that revolve around sharing rivalrous private goods i.e. the consumption of one prevents or disrupts the consumption of the other, and excludable – the owners of the good can prevent others from consuming it. In contrast, the reality of a road networks is of common-pool goods. While a road network is rivalrous, as the addition of more travelers influences its performance, it is not excludable – as it is impossible to prevent travelers from using it (unless physically barred). To bring the sharing economy principles to road networks and improve their performance, a behavioral approach is needed whereby the road network capacity is cooperatively shared in an optimal way.

As further described in Section 2, cooperation in large groups of mostly anonymous and rivalrous members (like car drivers) hardly ever emerges spontaneously, we therefore suggest an advanced traveler information systems (ATIS) concept for supplying fairly system-optimal routing information that could persuade and enable drivers to cooperate on the network's usage (e.g. by learning to taking turns on different possible routes). The ability to allow cooperation to emerge, the quality of the possible day-to-day traffic states and the potential dynamics such an ATIS concept might generate are the main aims of our paper. As this idea has no previous comparative result, we decided to focus and therefore tested our approach on a simplified road network that includes only four nodes and four arcs, while generalizations for more elaborate, directed and decentralized networks are left for future research. The rest of the paper is organized as follows: Section 2 presents a theoretical and methodological literature review on cooperation and

Table 1
The Route Choice Game. Assuming T > P > S > R, Nash equilibrium is where both players choose the defection strategy (Route 2), and the aggregate score is 2*p, which is also the solution. The system optimum (S + T)/2 occurs when players take turns, choosing route 1 and 2 alternately.

		Player 2	
		Route 1	Route 2
Player 1	Route 1	R,R	S,T
	Route 2	T,S	P,P

its relation to the system optimum; Section 3 describes the methods applied in our study including the day-to-day agent-based model and ATIS design, Section 4 reports the results of the agent-based simulation runs to test the performance of the ATIS for efficiency, stability and equity; and Section 5 presents a discussion and conclusions.

2. Literature review

We divide our overview to theoretical aspects and methodological ones related to the emergence of cooperation in decentralized networks.

2.1. Theoretical aspects

Traffic theory and Game Theory have several parallel concepts: UE is similar to a Nash equilibrium (Nash, 1951), a state where no player has any incentive to deviate from her chosen strategy. SO is similar to Kaldor-Hicks efficiency – a measure where the aggregate utility of all players is maximized. Dawes (1980) identifies these two features as the generators of social dilemmas i.e. the utility of all players is better when they all cooperate than if all defect, and the individual utility from defecting is always higher than cooperating. This logic explains why the natural behavior of travelers brings the network to a stable state of UE but rarely without hard interventions to SO.

Stark (2010) refines the concept of partial cooperation, which enables the system to reach an optimum by inducing some of the players to act cooperatively, while other players choose the rational action – to defect. This concept is best explained using the route-choice game, shown in Table 1. This 2 × 2 game considers a fast route and a slow route that are sensitive to congestion. While the fast route is always better than the slow route for the individual player, and hence, UE occurs when both players choose it, the social optimum occurs when one player is on the fast route and the other is on the slow one. Hence, in a single game the chance that rational players will reach is hypothesized to be very small. Conversely, partial cooperation can become the game's equilibrium if it is repeated, whereby the players learn to cooperate by taking turns and alternately switching between the faster and slower routes, and thus sharing fairly the time savings between themselves.

Helbing et al. (2005) tested partial cooperation in an economic experiment that included a route-choice game carried out first between two and then between four players. The players were informed that they would be able to achieve maximal time savings by coordinating their actions (i.e. a context awareness signal). The results showed that players learn to coordinate their actions after a certain adaptation period. These results are further backed by the studies by Browning and Colman (2004) and Colman and Browning (2009), who use a genetic algorithm simulation to investigate the formation of strategies in four different games. Helbing et al. (2005) realized that in games with asymmetric equilibrium points, coordinated alternating reciprocity can evolve without communication between the players. However, even with participating four players and following the findings of Boyd and Richerson (1988) that it is far more difficult for cooperation to emerge as group size increases, it seems that day-to-day learning, though necessary, is not sufficient to encourage players to cooperate in large numbers over a long period of time.

Generalization of the previous studies to N-players is intractable, as demonstrated by Zhao et al. (2008) in the battle of the sexes game. Nonetheless, we argue that the possibility to coordinate players' actions, and bring about the emergence of cooperation in large groups, could be implemented within the scope of ATIS technology. Several conceptual approaches are possible here: A one shot approach based on the correlated equilibrium principles (Aumann, 1974). A second approach is the online routing approach suggested by Varga (2014a, 2014b) to utilize the available road resources in real-time. A third approach considers a day-to-day dynamic traffic framework (Cantarella and Cascetta, 1995; Cascetta, 1989). The one shot setting ignores learning and long run formation of behavior, while the online routing approach, overlooks agents' freedom of choice and their consequent adaptive learning over time. Therefore, in this study we adopt the day-to-day approach.

The parallelization of the concept of SO in the route choice game and in complex road networks is not intuitive as the computation process of SO in complex road networks is computationally expensive, contrary to the simplicity of the route-choice game, where the road network has only two routes and a predefined number of road users. The unit of flow in question is affected by this difference – while in the simple route choice game one can relate to each player as a single unit of flow, in complex road network we do not relate to the players individually, but as continuous streams of flow. This has an effect over the optimization method – while in simple road network linear optimization is feasible, in complex road networks this optimization method is unavailable. It is also important to note that each agent is in fact a unit of flow therefore, and different from classical traffic assignment optimization approaches, integer programming is also a possible optimization approach for more elaborate networks. Moreover, multiple SOs can possibly co-exist in a

complex road network, a fact that raises the relevance of the fairness concept – some SO states will be fairer than others, in the sense that the inequality in travel time will be larger in some SO states than in others. Balancing between efficiency and fairness is one way of solving this issue (Jahn et al., 2005; Zhu and Ukkusuri, 2017). Nonetheless, adopting the day-to-day approach is a different solution to achieving fairness – the allocation of resources may not necessary be fair in a certain one-shot incident, but across a longer period of time resources will be allocated in a fair and stable way.

Our idea is based on prescriptive information delivered to drivers in the form of daily route recommendations that are essentially individually tailored recommendations aimed at persuading them on each day to choose a route based on an overall policy whereby if all drivers cooperate, the network will likely reach a stable and more importantly fair system optimum. This is similar to the work of Agogino and Tumer (2008), who explore the role of rewards as a part of agent's learning in dynamic and stochastic domains. However, we also exploit two key technological behavioral and concepts: Persuasive Technology and Choice Architecture. Persuasive Technology (Fogg, 2002) harnesses technological mediums to bring a change to human habitual behavior. Choice architecture (Thaler and Sunstein (2008), considers the use of nudges that are small features in the choice environment that highlight the better (e.g. healthier, safer) alternatives without restricting the consumer's freedom of choice and without changing the physical environment or the choice set. Avineri and Waygood (2013) demonstrate an application of nudges to motivate carbon reducing travel options like bicycling and public transit. Thus by implementing route recommendations as a type of nudge, drivers might be persuaded and learn to cooperate via an SO-ATIS application.

2.2. Methodological aspects

Dixit et al. (2015, pp. 2-3) suggest that multiplayer economic experiments seem suitable for the examination of social mobility dilemmas, such as traffic congestion. Experiments could well overcome the pitfall of hypothetical bias – participants actively deal with scenarios which have implication over their direct utility, and forced to reveal their preferences, as opposed to participants who are just stating their preferences with no strings attached. However, as human decision making behavior is often based on different sets of heuristics (for example: affect; Finucane et al., 2000, availability; Gilovich et al., 2002, anchoring; Tversky and Kahneman, 1974, to name a few), to understand complexity driven by the actions of bounded-rational human players, it may be better to initially simulate the plausible system dynamics using programmable agents, thus focusing future research on factors which likely will have a substantial effects over network performance while limiting the number of costly treatments required to be conducted with human participants.

Simulation techniques are often used to represent the resulting dynamics of the collective outcomes of individual human decision makers (Meijer and Hofstede, 2003). Agent-based models (ABM) are well suited to simulating such dynamics in transportation systems, e.g., the Multi-Agent Transportation Simulation (MATSim) environment for simulating transportation networks (Balmer and Rieser, 2009). The ABM consists of a collection of heterogeneous agents, the rules governing them, and the rules of agents' interactions with the environment they live in (Shalizi, 2006). As such, ABM is a bottom-up method for capturing possible macroscopic regularities (Epstein, 1999). In addition, ABM is suited for detailed hypothesis testing as discussed by Helbing and Balietti (2013).

ABM simulations are quite commonly used in transportation studies (see Bazzan and Klügl (2014) for a review). Shang et al. (2017) explore the effect of information percolation (the spread of information) using a day-to-day ABM, and find that percolation and the convergence to UE are positively correlated. Wei et al. (2016) also use day-to-day ABM to explore the effect of social interaction on network dynamics and find that while the number of interactions a traveler encounters influences her route choice strategy, the evolution of the aggregate network flow is not much affected. He et al. (2013) test the performance of route guidance strategies on an asymmetric two-route traffic network using a time step ABM, and find that the mean velocity feedback strategy (choosing the route with the lower mean velocity) is the only one that manages to equalize travel time on both routes. Wahle et al. (2002) uses time step simulation to study the impact of different types of travel information over two-route traffic network scenario, and discover that information about route density enhances network performance.

To the best of our knowledge, no ABM for simulating driver cooperation using ATIS has yet been developed neither for day-to-day nor real-time traffic evolution. The exception is Levy and Ben-Elia (2016), who consider a simple binary network where agents are provided with feedback of the daily network travel times and respond by switching routes when their last action increased it. Their surprising result is a stable and equitable system optimum (i.e. all agents contribute similarly to the common pool). However, given that their approach requires a strong assumption of altruism (or at least partially), it is not inconceivable to consider that this would be quite difficult to maintain with human drivers. Therefore, it remains an uncharted territory if voluntary route recommendations (as a soft policy) could change driver behavior from competition to cooperation without the need to work against human nature to behave in a self-interested manner and without the need to resort to hard policies like punishments and rewards.

3. Methods

We develop an agent-based model for studying the emergent day-to-day traffic states on a simple road network with partial supply of information. The model is comprised of five layers: temporal dimension, network design, agent behavior, routing policies allocations and incentive mechanisms. We elaborate on each one of these components below.

3.1. Temporal dimension

As the concept of partial cooperation involves a repeated-like game, to examine whether the system actually converges in the long

run to an optimum, the model in question should be analyzed according to the framework of day-to-day traffic dynamics (e.g. Cantarella and Cascetta, 1995; Cascetta, 1989). To verify the robustness of agent behavior, the simulation first entails a phase where no route recommendations are given, so that agents do not pre-adapt to the recommendations. We hypothesize the system will converge to UE at this phase. Subsequently, a second phase starts where recommendations are given to the agents, and their behavior is then hypothesized to adapt somehow to them.

3.2. Network design

The road network is formally represented by a directed graph G = (V, E) with vertex (i.e. node) set V (V ∈ {"origin", "throughA", "throughB", "destination"}) and edge (i.e. arc or link) set E (E ∈ {"origin" → "throughA", "origin" → "throughB", "throughA" → "destination", "throughB" → "destination"}). G, therefore has a path (i.e. route) set P E ∈ {"origin" → "throughA" → "destination", "origin" → "throughB" → "destination"}). While it is possible to expand the network design into a more elaborate directional graph with multiple nodes and edges that include independent link costs, here we focus on a simplified road network to investigate core factors that influence the sensitivity of cooperation and emergence of SO. Congestion on path set P is derived directly from the agents' choices, as they compete over road capacity. Competition means dependency of the travel time (T_i) on route i (i = A, B) on the number of travelers taking it (N_i) using a monotonously increasing volume-delay-function (VDF) of the form $T_i = \alpha_i + \beta_i N_i^{\delta}; \alpha, \beta, \delta \geqslant 0$, where: α is free-flow travel time, β is the congestion effect, and δ is the power effect. We consider mainly the effects of recurring congestion and do not explicitly account for the influence of non-recurring incidents on traffic buildup. While this assumption can be later relaxed, it is also not entirely imaginary in a future where road networks will be populated by intelligent autonomous vehicles that are dramatically expected to reduce both the number and severity of road accidents (Fagnant and Kockelman, 2015). Let N_A^{SO} be the share of drivers on route A under the system optimum (SO) found by optimizing the aggregate travel time as a function of N_A i.e. $min\{T_A * N_A + T_B * N_B\}; s. t: N_A, N_B \geqslant 0$ and $N_A + N_B = D$, where D is the total number of agents departing from the origin vertex. Let N_A^{UE} be the share of drivers on route A under the unique user equilibrium (UE) found for the value of N_A we obtain $T_A = T_B$. Naturally, multiple UE might exist in more elaborated networks (e.g., Schulz and Stier-Moses, 2003, 2006).

It is important to note that route recommendations that work in one network design might not necessarily work in another, due to the diverse features each network has. In order to examine the efficiency of route recommendations in different topologies of road networks, two attributes of road networks can be used as a scale: the "Price of Anarchy" and the "Share of Agents on the Faster Route under the System-Optimum" (AFRSO). The Price of Anarchy, as defined by Roughgarden (2005) and formulated in the road networks context by Mak and Rapoport (2013), is the ratio between the aggregate cost of the system under the worst UE (when travel costs on all routes are equal) and the aggregate cost of the system under SO (when the overall travel costs are minimized). In our case, the costs are the agents' experiential travel times. As that price of anarchy is influenced by the total flow of agents populating the system (O'Hare et al., 2016) which for consistency purposes we wish to keep constant, we chose to make use of the AFRSO instead for the comparisons between networks. Similarly, we use the share of agents on the faster route under UE – AFRUE for the same purpose. Under SO, one route almost always has a shorter travel time than the other depending on the specific network design parameters, and AFRSO will also change accordingly. As AFRSO increases, more agents are required to use the faster route in SO. Therefore the agents will likely experience lower travel times more often and we hypothesize that their propensity to cooperate by complying with route recommendations will also likely increase.

While the investigated network is simple in a sense that it emphasizes the steady flow of the network, as opposed to intermittent and transient flow more reminiscent of microsimulation models that complicate the vehicle routing problem, we consider this necessary in order to reach some basic comprehension as to the dynamic quality of SO-targeted route recommendations to bring about the emergence of cooperation. Naturally, many modifications to this model are possible in future research with more elaborate car following models that account for acceleration and deceleration of vehicle dynamics. We return to this issue when we elaborate on the simulation procedures.

3.3. Agent behavior

We assume that the agents' decision making process is based on two strategies: exploration and exploitation. These strategies have been well-established in the psychological literature regarding reinforced learning (Hills et al., 2015). Exploration suggests that once in several trials, an alternative from the choice set is randomly chosen, based on a predefined probability function. Exploitation, in contrast, is based on a maximization strategy where the agents assess the utility of the different alternatives based on a reinforced learning process that is derived from their past experiences. This is similar to the studies by Horowitz (1984) and Wahle et al. (2002) on learning in a simple binary network. Two cognitive models for exploitation are worth mentioning: Total Recall and Sampling and Weighting.

Total Recall (TR) can be considered as a simplistic machine learning model. Agents behave similarly to robots capable of perfectly recording all the outcomes (gains and losses) of their previous choices. Accordingly, the agents evaluate the mean utility of each alternative, and choose the one with the highest utility. In essence *TR* emulates a fully rational choice process. In contrast, *The Sampling and Weighting (SAW)* model (Erev et al., 2010) is a generalization of the Total Recall model. We chose it to model agent behavior due to its successful predictions in market entry games. Reward shaping (e.g. Babes et al., 2008) is an additional interesting model for formulating agent behavior. Regarding the different alternatives, the SAW model assumes agents consider both their short-term memory that reflects the outcomes of the most recent experiences, and long-term memory that takes into account the entire range of the agent's experiences. These memory components are weighted using a weighting parameter. Compared to *TR*, SAW can be

attributed to a bounded-rational choice process. The SAW utility function is described in Equation 1.

$$U_j(t) = W * \frac{\sum_{k=0}^{k_j} -C_j(k)}{k_j} + (1-W) * \frac{\sum_{i=k_j-\delta}^{k_j} -C_j(k)}{\delta}$$

(1)

where U_j is the average utility when choosing alternative j, C_j is the travel cost on alternative j, k_j indexes the trips the agent travelled on alternative j out of t trips, W is the long/short term memory weighting parameter, and δ is the number of recent trips drawn from the agent's long-term memory.

3.4. Route recommendations allocation policies

A key problem of traffic dynamics is that optimality conditions require the agents often choose in contradiction to their natural heuristics mentioned above. We consider the provision of route recommendations to the agents as a possible solution to the aforementioned problem by providing each agent with a daily personalized route recommendation (i.e. prescriptive information), designed in such a way that if all agents were to comply with such guidance, cooperation would emerge and the network will be at a stable SO state. Helbing et al. (2005) first conceived this idea, but it was never tested analytically or empirically.

We assume that agents are capable of gradually learning to cooperate by complying to the recommendations using two additional memory slots – travel time experienced when complying with the recommendations, and travel time experienced when not complying. When we translate these memory slots into ATIS concepts, we give the agents an indicator to the utility of the recommendation. This mechanism is a prerequisite in order that the agents will learn to comply with the ATIS while human memory is different due to other traits such as forgetfulness, recency, etc. Regret minimizing behavior – the notion of making decisions that reduce future regret can also be considered here as shown by Blum et al. (2006) whereby in routing games regret minimizing behavior converges to UE. However, our model differs from regret minimizing behavior because it includes two separate decision planes – compliance to the recommendation and the actual route-choice. Compliance to the recommendation is not necessarily regret minimizing. Box 1 depicts the algorithm we applied for agents' decision making with recommendation adaptation.

While route recommendations are always provided to all agents, their allocations between the agents can vary. The efficiency of the recommendation allocation process (i.e. allocation policy) is critical, as the perception of trust the agents have regarding them is directly dependent on their own experience. If the allocation process is perceived as unreliable, agents would learn not to trust the recommendations making the ATIS essentially ineffective. Furthermore, in SO, one route is more attractive than the other, and if the allocations are distributed in a way that discriminates agents, the trust in them would further decrease.

Although many allocations are possible, we focus here on five specific ones. These allocations are sorted into two approaches: Predefined-static and Adaptive-Dynamic. *Predefined-static* – The recommendation allocation is independent of the agent's actions or characteristics, that is, the recommendations are allocated for all days in a predefined way. Two types of predefined allocations can be considered: Random and Queue. Their operationalization is depicted in Fig. 1.

Under a random allocation, the number of agents that receive a recommendation to take a specific route remains constant. However, the specific identity of those agents can change stochastically from one day to another. Although some agents end up receiving a recommendation to take the slower route while others should take the faster route, the average chances of the better/worse recommendation remains essentially equal across all agents, i.e. their expected value is kept identical.

Under the "Queue" allocation, agents receive the better/worse recommendations in sequential turns, based on the number of times they received the better recommendation – if they received recommendations to take the faster route more times than others,

Box 1
Algorithm of agent's decision making process with recommendation adaptation.

```
If random number < exploration rate:

 • Choose random route[A or B]

   else:

 • If U_complied > U_Not Complied:
   • Choose route according to recommendations system
 • Else:
   • If U_a > U_b:
     • Choose route A
   • Else if U_a < U_b:
     • Choose route B
   • Else:
     • Choose random route[A or B]
```
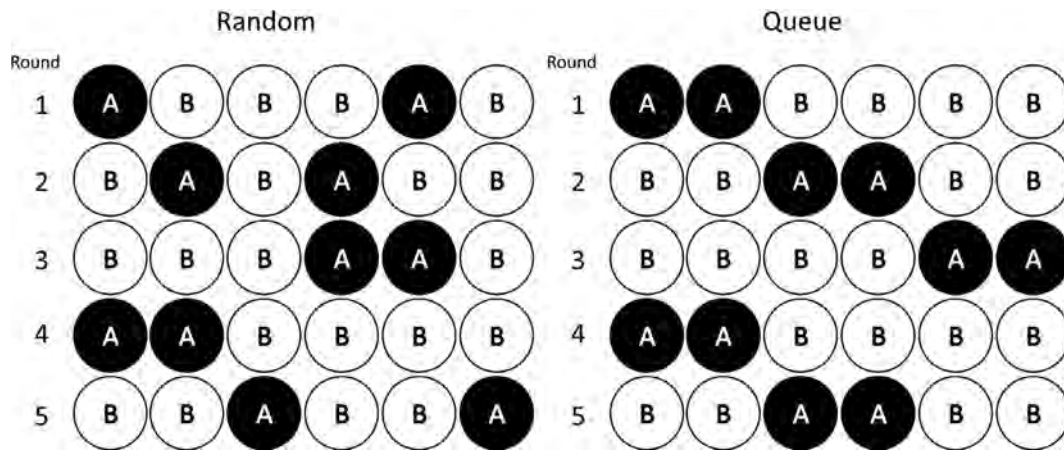
Fig. 1. Hypothetical representation of "Random" and "Queue" allocations. The system optimum occurs when $N_a = 2$ and $N_b = 4$. In random, recommendations are allocated to agents stochastically. In queue, recommendations are allocated to agents by their turn.

they will receive the worse recommendation, and vice versa – if they received it more times than others, they will receive the better recommendation. Over time, all agents will have certainly received an equal number of better/worse recommendations. This process is designed to mimic the turn-taking strategies in the route-choice game theoretical framework.

As opposed to predefined allocations, dynamic ones are not independent of the agents' actions or their characteristics. That is, whether an agent receives a better or worse recommendation depends on its previous actions and attributes, as illustrated in Fig. 2. In this sense, the recommendation system is recording the behavior of the agents and adapting to different traits the agents reveal. We consider three approaches how this can take place: Justice, Anti-Merit, and the Reformer allocations.

Under the "Justice" allocation better recommendations are given each day to those agents whose average travel time is higher. The rationale behind this allocation is the presumption that if the equality between the agents will increase, their trusts in the recommendations will also increase, and consequently their propensity to comply with the recommendations. This allocation also serves the agenda of Rawls (1999), by striving to create a more equitable set of travel times between the agents. Under the "Anti-Merit" allocation, better recommendations are provided to those agents who least cooperate. At first glance, this mechanism seems unfair as it is rewarding agents who do not comply with the recommendations. However, the rationale here is to increase those agents' trust in the guidance system, encouraging them to learn to eventually cooperate. The "Reformer" allocation is similar to Anti-Merit in its rationale, but differs in its implementation. Better recommendations are given to those agents whose difference between the experienced travel time when not cooperating and the experienced travel time when complying is the largest. The connotation "Reformer" implies this mechanism is designed to persuade agents in concordance with what they actually remember, thus in a sense "reforming" them to maintain cooperating. Box 2 describes the algorithm applied to the recommendation allocation policy computation.

3.5. Incentive mechanisms

Cooperation depends on whether the agents remember their previous actions. Rather than relying solely on the agents' own learning, hard policy measures involving punishments and rewards – can be also incorporated if needed. Helbing et al. (2005) postulated that the agents who do not cooperate and instead behave selfishly should pay society for the proportional time they added to all other agents in the system. Conversely, cooperating agents should be compensated for their proportional increase in travel time.
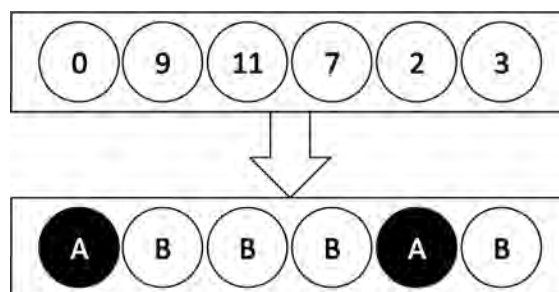

Fig. 2. Illustration of the operationalization of Dynamic allocations (SO occurs when $N_a = 2$, $N_b = 4$, and route A is faster). The parameter according to which the recommendations are allocated serve as an input layer, route recommendation recommendations are the output allocated according to the minimal values in the input layer.

Box 2
Algorithm for allocation computation.

Set *faster route in SO, slower route in SO, AFRSO* according to network design.
   Set *agentlist* **as** the list of agents and their attributes.
   If *allocation* = queue:

• **Foreach** *Agent* in *agentlist* **set** *attribute* **as** number of times agent was sent to *faster route in SO*.

  If *allocation* = random:

• **Foreach** *Agent* in *agentlist* **set** *attribute* **as** random number.

  If *allocation* = justice:

• **Foreach** *Agent* in *agentlist* **set** *attribute* **as** average utility (U)

  If *allocation* = anti-merit:

• **Foreach** *Agent* in *agentlist* **set** *attribute* **as** number of times agent complied with recommendations.

  If *allocation* = reformer:

• **Foreach** Agent in agentlist **set** attribute **as** $U_{complied} - U_{not\ complied}$.

  Order *agentlist* (ascending) according to *attribute* and give agents indices.
  Foreach agent in ordered *agentlist*:

• **If** agent's index $\leq$ *AFRSO*:
  • **Send** agent *faster route in SO* as recommendation.
• **Else**:
  • **Send** agent *slower route in SO* as recommendation.

Over time if agents cooperate and learn to trust the ATIS, the size and number of these transactions should decrease. We adopt this idea into the operation of the ATIS, with some modifications: every noncompliance event is met with a penalty (punishment), whereas compliance is met with a reward. This modification becomes crucial when dealing with more than two players, as there is the possibility the faster route is underused and the aggregate travel-time could be even higher than UE.

Punishments are computed as the average time wasted by defectors, and rewards are computed as the sum of the time wasted by cooperators divided equally between them. In real life, the value of time is different across drivers, but in the model we simplify reality by applying the same value of time to all agents. Agents experience the effect of a punishment/reward the same way they would experience changes in travel time. Thus we consider the punishment/reward to be additive to the experienced travel time, the equivalence of which is "translating" the value of time into monetary terms, and imposing monetary fines on defectors, while awarding a monetary incentive to the cooperators. The individual treatment of each road user could be done by tracking them using their mobile phones. The aforementioned "translation" is done because of the added complexity of value of time and its variance across the population. We explore three different combinations of these incentive mechanisms: reward only, punishment only, and no incentive mechanism (i.e. no punishment or rewards). The suggestion by Helbing et al. (2005) considering both punishments and rewards simultaneously was tested initially. However, the simulated results of this mechanism were practically identical to the results of the punishment mechanism. One reason is that the behavioral response to punishments masks every other kind of intervention in its presence. Consequently we modeled separate scenarios for rewards and punishments.

3.6. Simulation procedures

Each of the five recommendation allocation designs and 3 incentive mechanisms previously discussed was simulated – in total $5 \times 3 = 15$ treatment combinations were investigated. In addition, the simulation runs were performed across 49 different road networks, consisting of 49 different free-flow-time parameters for Route B as portrayed in Fig. 3. In total, 36,750 simulations (50 runs of each setting, $3 \times 5 \times 49 \times 50 = 36,750$) were run in order to test the model and compare between different recommendation and incentive mechanisms. Agents were not provided with route recommendations in the first 5000 rounds of the simulation, (out of 10,000 rounds in total) and learned from experience which route is better off, allowing the road network to stabilize and reach a UE-like state. Only in the subsequent 5000 rounds, were the agents provided with route recommendations.
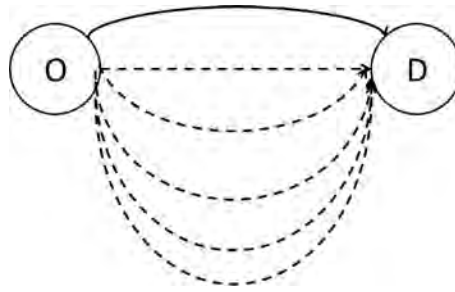
Fig. 3. Illustration of different network designs. Route length expresses free flow time. Route A (full line) is fixed; the free flow time of Route B (dashed line) is allowed to change.

Table 2
Specification of simulation parameters.

	Parameter	Value
Time perspective	Rounds without recommendations	5000
	Rounds with recommendations	5000
Agent behavior	Exploration rate	3%
	Sampling and weighting W	0.5
	Sampling and weighting δ	3
Recommendations	Allocations	Random, Queue, Justice, Anti-merit, Reformer
	Incentive mechanism	Punishment, Reward, None
Network design	N	100
	α_A	50
	α_B	51, 52…99
	β_A, β_B	0.0016666
	Δ	2

In the simulations, Route A is designed with fixed parameters while the share of agents on route B under SO is allowed to wander reflecting its increasing free-flow travel time relative to route A. Table 2 describes the different parameters which were used in the simulation runs. The behavioral parameters were chosen after a process of trial and error of multiple parameter values. Although there were differences between various sets of behavioral parameters, these were not significant.

The simulation runs were written in JavaScript, based on a car-following microsimulation in a similar binary road network built previously in Netlogo (see Levy and Ben-Elia, 2016). As the running time was significantly higher with Netlogo and given the large number of simulations required we preferred to use JavaScript. However, to verify parameter robustness, we regressed the Netlogo volume to delay results and estimated the parameters of the VDF for each network design that was later applied in the JavaScript source code.[1]

It is important noting that by allowing the free-flow-time on Route B (α_B) to wander two decision spaces for the agents were essentially produced. As depicted in Fig. 4, by modifying α_B, the share of agents on route A under UE and under SO also has to change. The first decision space occurs when Route A is only in some cases faster than Route B i.e. Route A has relative dominance over Route B. The second space occurs when Route A is always faster than route B i.e. Route A has absolute dominance over route B. We hypothesize different behaviors will emerge in these two spaces.

3.7. Performance indices

Five performance indices were applied to compare between the different simulated treatments, and are presented as their mean across the last 5000 rounds:

(1) Efficiency – the mean difference between the aggregate utility in SO and the actual aggregate utility, divided by the difference between the aggregate utility in SO and the aggregate utility in UE: $\frac{\sum_{k=1}^{N} U_{SO} - \sum_{k=1}^{N} U_{ACT}}{\sum_{k=1}^{N} U_{SO} - \sum_{k=1}^{N} U_{UE}}$;
(2) Stability – the coefficient of variation of Efficiency (mean to variance ratio), which encapsulates the size of deviation from the mean. The higher the stability index, the less stable the system is; naturally other formulations are possible, e.g. the study by Como et al. (2013).
(3) Cooperation – The mean percent of agents complying with the recommendations (i.e. concordance with the recommendations).

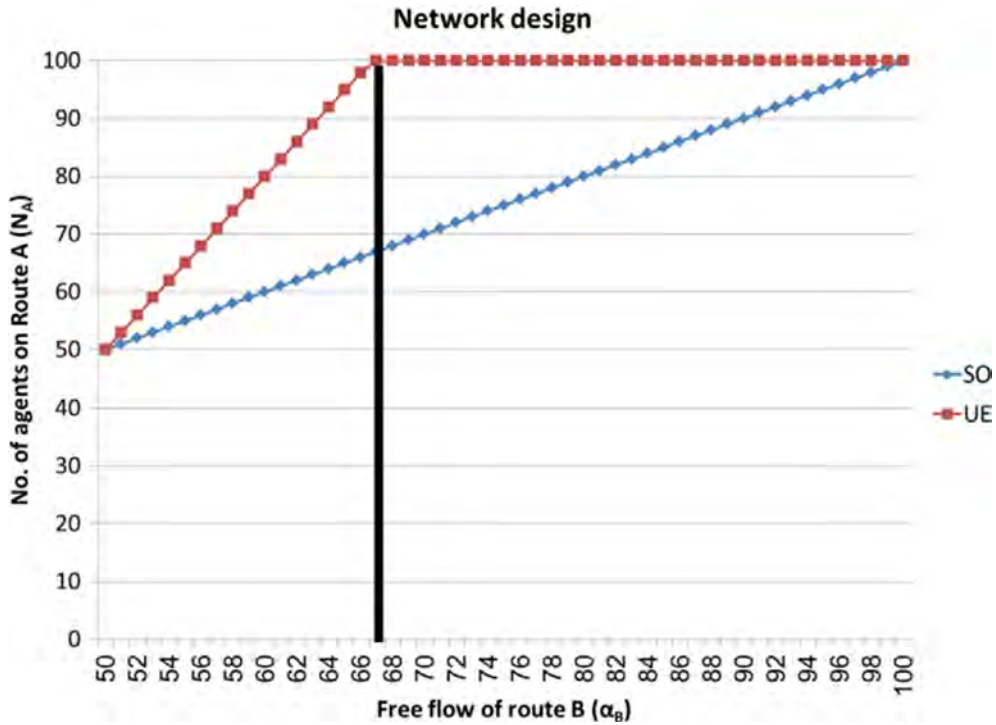[1] Source code is available in https://github.com/idshklein/system-optimal-ATIS.

Fig. 4. Number of drivers on route A in SO (AFRESO) and UE (AFREUE) as a function of α_B.

Where M_k is 1 if the agent complied with the recommendation, and 0 otherwise. Then, compliance is computed in the following manner: $\frac{1}{N}\sum_{k=1}^{N} M_k$;

(4) Willingness to cooperate – The mean percent of agents who are willing to comply with the recommendations. Let W_k be 1 if $U_k^{complied} < U_k^{notcomplied}$, and 0 otherwise. Then, willingness to cooperate is computed in the following manner: $\frac{1}{N}\sum_{k=1}^{N} W_k$;

(5) Equity: The mean Gini index (Gini, 1912) of the agents' cumulative utility. The Gini index measures income inequality in individuals. We understand that it can be used not only for measuring income inequality in monetary terms, but also for measuring time costs of road users, as done regarding accessibility (Jang et al., 2017). Gini index is computed in the following manner: $\frac{\sum_{i=1}^{N}\sum_{j=1}^{N} |U_i - U_j|}{2N\sum_{i=1}^{N} U_i}$

4. Results

Before reporting the results, we note that as hypothesized, all network designs managed to converge to a stable-state UE in the first 5000 rounds with no route recommendations provided. This implies that the agents' behavior appears to be well calibrated to the networks in question.

4.1. Efficiency

The results of the efficiency index (as defined in Section 3.7) are shown in Fig. 5. It is apparent that across all the different allocations and incentive mechanisms, an increase in AFRSO brings an increase in efficiency. It is also clear that networks in the absolute dominance space have less variability than networks in the relative dominance space. However there are differences between the treatments. Table 3 shows the index means in the last 1000 rounds (9000–10,000) for each factor, and results of the Duncan test for between-group differences regarding allocation and incentive mechanism. Table 4 shows the efficiency indices means for those rounds regarding AFRSO values.

In relation to the incentive mechanisms (Table 3), all three treatments are significantly different from each other with the highest gain in efficiency attributed to punishment. The punishment also appears to smooth out the opposing effects of network design over efficiency. The reward is more efficient than having no incentives but the effect of network design still remains quite large.

Regarding the allocations, apart for the difference between random and anti-merit all are significantly different from each other. Reformer eclipsed all other allocations. Queue is also quite efficient, random and anti-merit come third with a similar effect (but with changes across different road networks, as will be later discussed), whereas justice performs rather poorly. Given the poor efficiency performance of justice we will omit it from the rest of the analysis.

Traffic states having no incentive mechanism are shown in Fig. 6. Remarkable patterns of interaction between network design and allocations can be observed. It can be seen that anti-merit and random compete for supremacy in different network designs (random
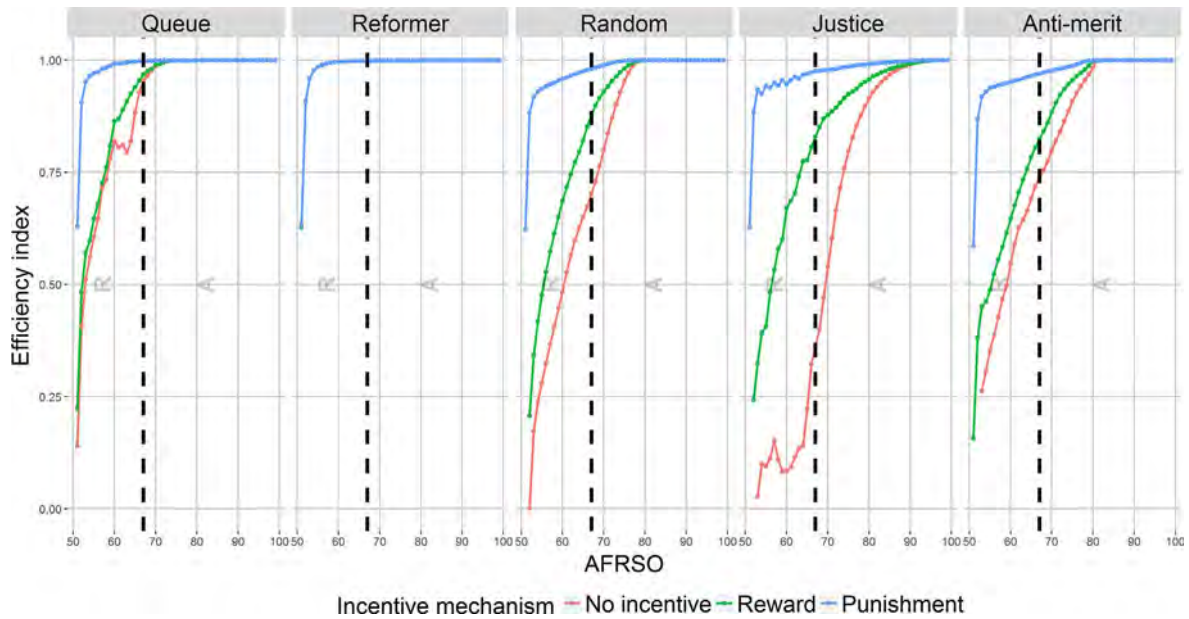
Fig. 5. Efficiency index mean in rounds 9000–10,000 by allocation and incentives. The area to the right of the dashed line represents the absolute dominance area (A), and to area the left of the dashed line represent the relative dominance area (R).

Table 3
Performance indices means in rounds 9000–10,000 for incentive mechanism and allocation design. In the Sig. group column, factors that have the same letter are not significantly different for each other (p < .01), according to the Duncan test.

		Efficiency		Stability		Compliance		Willingness to comply		Equity	
		mean	Sig. group	mean	Sig. group	mean	Sig. group	Mean	Sig. group	mean	Sig. group
Incentive mechanism	Punishment	0.97	a	0.099	a	97.34	a	96.23	a	0.0013	a
	Reward	0.87	b	0.07	b	94.99	b	88.32	b	0.0013	a
	No incentives	0.79	c	0.028	c	94.098	c	85.22	c	0.0013	a
Allocation design	Reformer	0.98	a	0.0207	d	98.5	a	99.5	a	0.0015	b
	Queue	0.93	b	0.063	c	96.51	b	93.021	b	0.00074	d
	Random	0.85	c	0.083	b	93.86	c	85.66	c	0.00084	c
	Anti-merit	0.85	c	0.096	a	93.035	d	81.52	d	0.00230	a
	Justice	0.78	d								

when AFRSO is larger than 70 – most of the absolute dominance space, anti-merit if it is smaller – the whole relative dominance space). The adverse effect of the different network designs prompts the question regarding the interaction of allocations with more complex network designs.

4.2. Stability

Regarding stability (as defined in Section 3.7), Fig. 7 shows that as AFRSO increases, the system becomes more stable. Duncan test reveals that networks with AFRSO > 72 are not significantly different, and networks with AFRSO ≤ 72 are usually significantly different. The instability in the networks in the relative dominance space can be explained by the emergence of oscillations in each network between UE and SO as depicted in Fig. 8, that presents the time sequence of rounds 9500–1000 for one single simulation run. Theses oscillations occur when the network approaches closer to a state of SO, it seems agents respond less to the recommendations and reduce the rate of compliance, until the system returns closer to UE where the recommendations seem to become more effective again bringing the system back to SO and these cycles continue over and over again. Nonetheless, these oscillations disappear when averaging out the results of the entire set of 50 simulations for the same treatment combination. In contrast, when the network is in the absolute dominance space, these oscillations occur less frequently, since the gaps between the recommendations and agents natural behavior is smaller – that is, agents will receive recommendations to take the faster route more often and are less inclined to develop distrust in the recommendations compared to the relative dominance space.

Stability is also influenced strongly by incentives. Punishment as expected provides the largest influence on stability compared to rewards and no incentive. In respect of the difference between allocations, this replicates the same picture observed for efficiency.

Table 4
Efficiency index means in rounds 9000–10,000 by AFRSO. In the right column, factors that have the same letter are not significantly different (p < .01), according to the Duncan test.

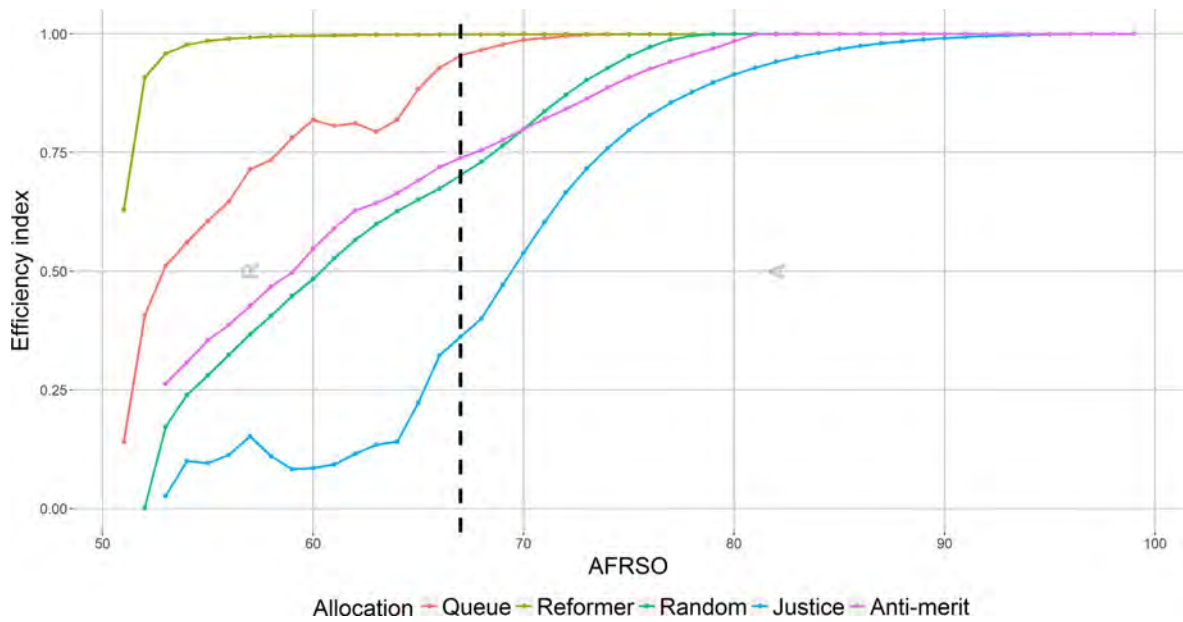| AFRSO	Efficiency mean	Efficiency Sig. group
[80, 81,…99]	[0.997, 0.999]	a
79	0.995	ab
78	0.993	ab
77	0.99	ab
76	0.986	abc
75	0.981	abc
74	0.975	abcd
73	0.969	bcde
72	0.962	cde
71	0.954	def
70	0.945	efg
69	0.935	fgh
68	0.926	ghi
67	0.917	hi
66	0.905	ij
65	0.891	jk
64	0.876	kl
63	0.863	lm
62	0.854	lm
61	0.84	mn
60	0.828	no
59	0.806	op
58	0.785	pq
57	0.767	q
56	0.742	r
55	0.721	r
54	0.694	s
53	0.664	t
52	0.561	u
51	0.151	v

Fig. 6. Efficiency index mean in rounds 9000–10,000, under no incentive mechanism.

4.3. Cooperation

Fig. 9 presents the results for cooperation (as defined in Section 3.7). Duncan test reveals that networks with AFRSO > 80 are not significantly different, and networks with AFRSO ≤ 80 are usually significantly different. Surprisingly, networks in the relative dominance area retain very high cooperation rates, despite lower efficiency. Nevertheless, as the difference between AFRSO and
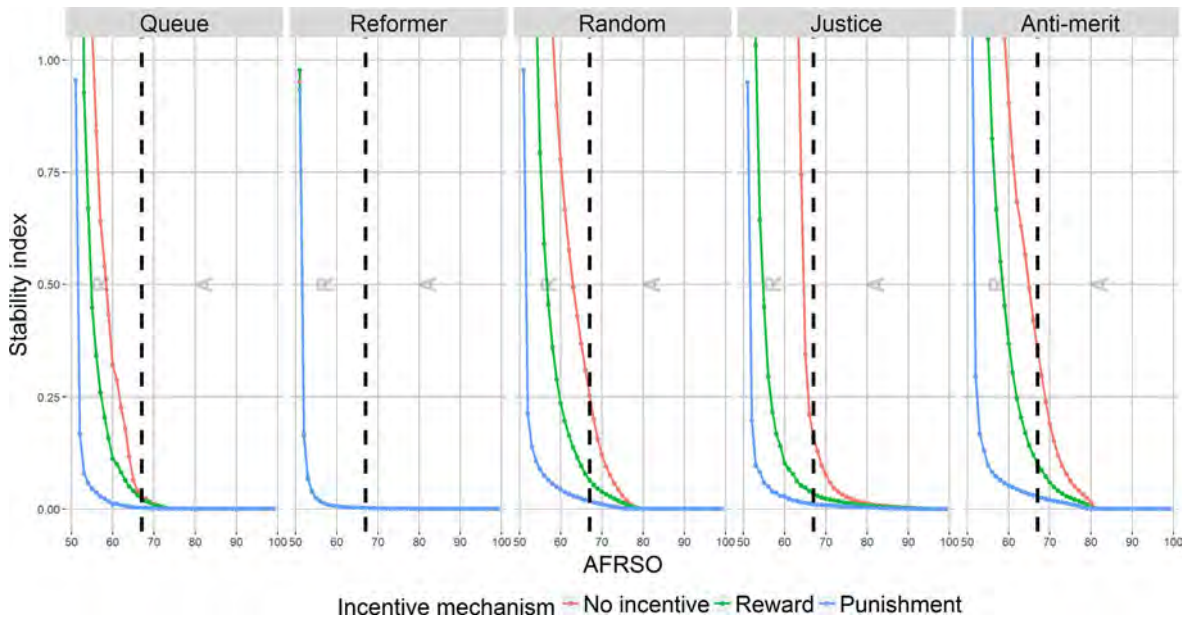
Fig. 7. Stability index mean in rounds 9000–10,000. This plot shows only values between 0 and 1, where zero suggests that the ratio between the standard deviation and the mean is low, and vice versa when the stability index is 1.
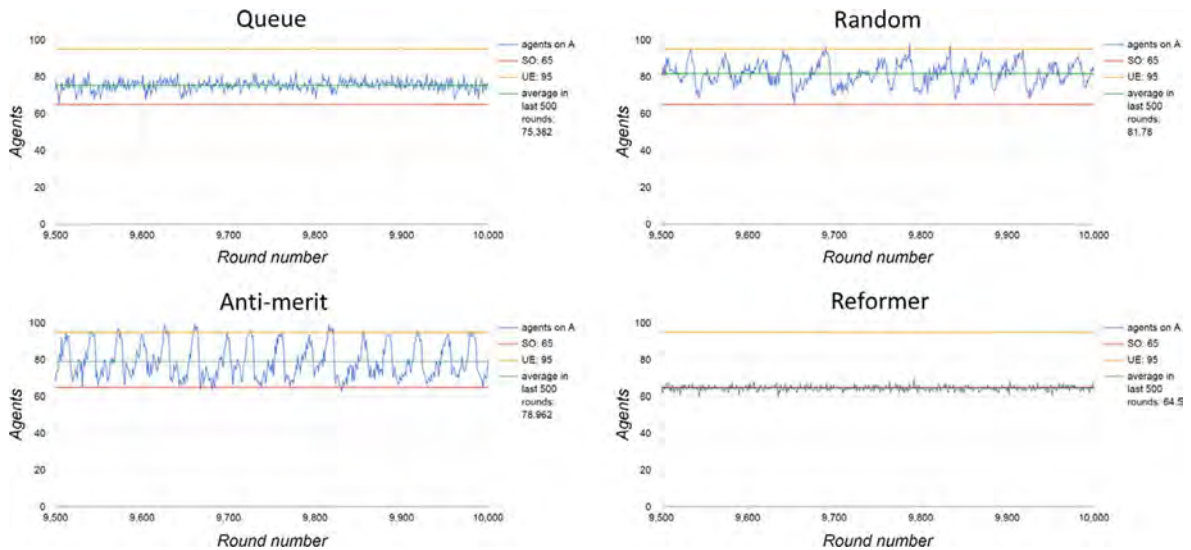

Fig. 8. Rounds 9500–10,000 in a random allocation with no incentive, $\alpha_A = 50$ and $\alpha_b = 65$ of one sample simulation.

AFRUE increases with the rise of AFRSO in the relative dominance space, lower cooperation rates correspond to larger impacts on efficiency. The reason is that while the majority of the agents were cooperating, the few defectors had a large negative influence over the network's efficiency in the relative dominance area. This does not occur in the absolute dominance area for the same reasons described above.

4.4. Willingness to cooperate

Fig. 10 presents the willingness to cooperate rates (as defined in Section 3.7). As described earlier, willingness to cooperate is based on the percent of agents who have $U_{complied} < U_{Not\ Complied}$, i.e. it is subjective, whereas "cooperation" is only based on objective concordance between the choice of the agent and the recommendation it received. This measure divides the agents' population into two classes – those who are willing to cooperate, and those who are not.

In the relative dominance area, there appears to be an anomaly: The rate is much lower than the objective compliance rate. The explanation lies in the fact that some agents may not necessarily be willing to comply to actually choose to comply. For example –
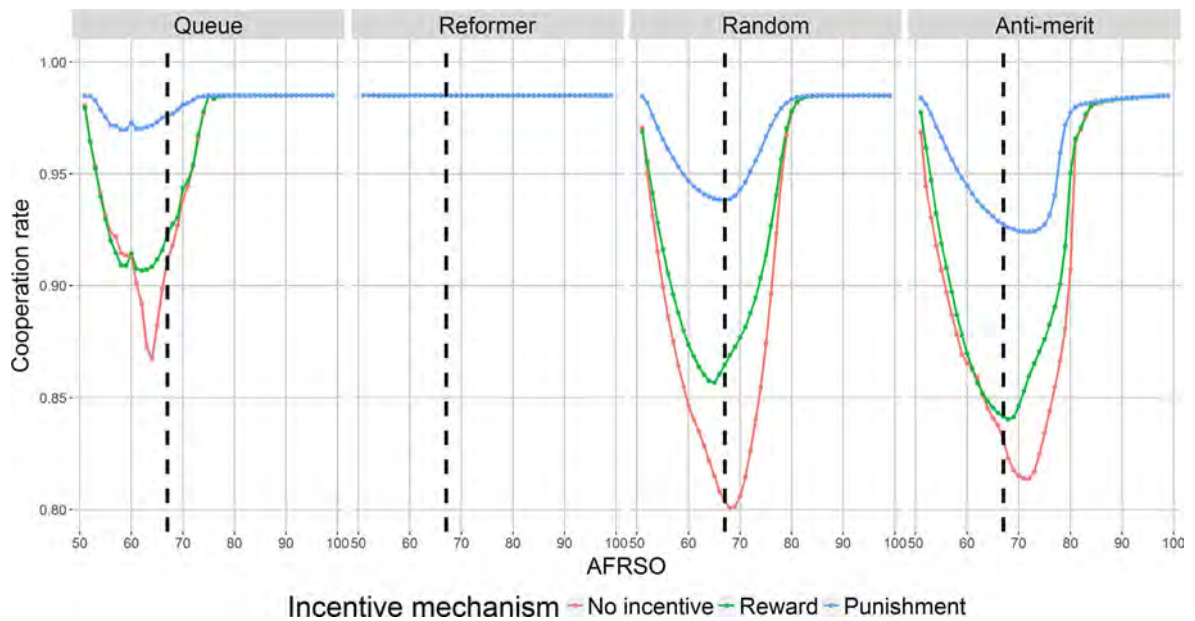
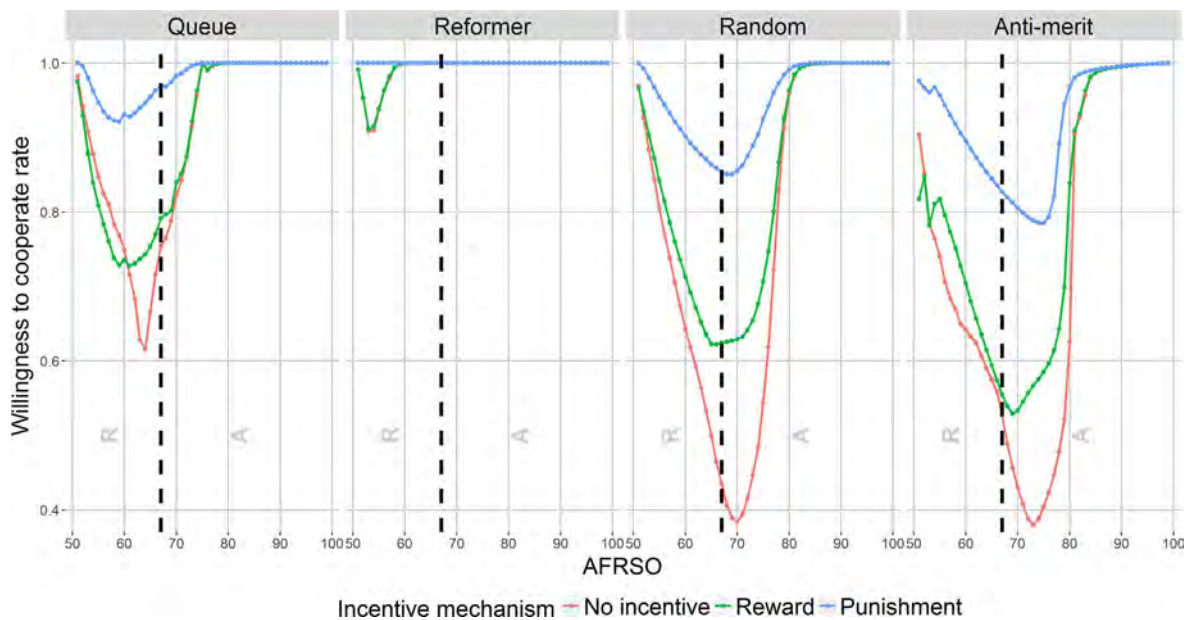**Fig. 9.** Mean cooperation rate in rounds 9000–10,000.

Fig. 10. Mean willingness to cooperate rate in rounds 9000–10,000.

agents who consider that Route A is better than B receive a recommendation to take A but are not willing to comply because compliance works on a different memory slot than route choice but will eventually be complying (i.e. concordance) despite of not willing to comply. In comparison between the allocations, the least anomaly happens under Reformer. Punishments followed by rewards naturally work to increase the willingness to comply for any given network design, while no incentive hampers the willingness to comply especially in Random and Anti-Merit and to some extent in Queue. Duncan test reveals that networks with AFRSO > 81 are not significantly different, and networks with AFRSO ≤ 81 are usually significantly different.

4.5. Equity

Regarding equity (as defined in Section 3.7), the overall picture observed in Fig. 11 is quite positive as all allocations retain very low Gini values. Nevertheless, some allocations and network designs result in less equity than others. Apart for Anti-Merit, punishments have a positive effect on equity, whereas rewards and no incentive have almost no difference. It also appears that the greater
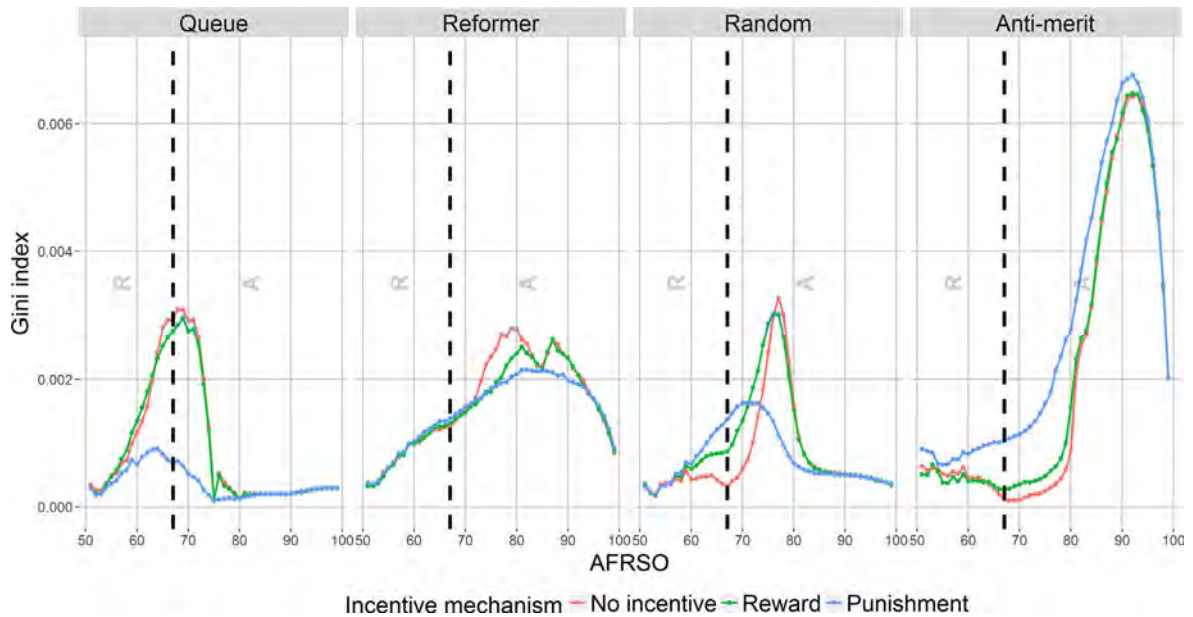
Fig. 11. Mean equity (Gini index) in rounds 9000–10,000.

the gap between AFRUE and AFRSO moving from relative to absolute dominance the larger is the change in equity. Anti-merit is the least equitable allocation but the differences between the other allocations appears negligible. Without incentives, Reformer is most equitable while random is more equitable with punishments. Duncan test reveals some significant differences between networks.

5. Discussion and conclusions

In this study we investigated the potential of a system-optimal ATIS to bring about the emergence of cooperation in a simple road network thus improving its performance. We developed an agent-based model for simulating day-to-day traffic evolution on a simple binary network based on the adaptation of the 'sampling and weighting' algorithm. We included tailored policies (static and dynamic) for allocating route recommendations between agents as well as testing the influence of incentive mechanisms based on punishments and rewards. We allowed the network to converge to UE before introducing the ATIS and measure the efficiency, stability and equity of the resulting system. We only considered direct effects of recurring congestion.

We showed that the network retains efficiency and equity even without hard policy measures. In addition, recommendation allocation policies have different effects and when combined with specific network design parameters, result in different levels of efficiency. The Justice allocation was not found to be efficient which was counterintuitive, whereas the Reformer that works to "convert" deviant agents was found to be very effective. These results provide evidence that cooperation in simplified road networks with a large number of participating agents can emerge without strong assumptions regarding travelers' behavior (e.g. altruism as considered by Levy and Ben-Elia, 2016). Instead, emergence of cooperation can be the outcome of the agents' bounded-rational decision making that does not strongly contradict their self-interest. This illustrates that even without an explicit agenda of 'Tikkun Olam',[2] travelers can learn how to share simplified road networks in a way that maximizes both personal and societal utility.

Self-organization with system-optimal ATIS managed to maintain high levels of equity between the agents in the long run. The Gini index achieved by the simulations runs was always quite low (the highest Gini value was 0.00768), despite of the embedded inequity that the ATIS possesses, i.e. – recommending faster routes to some travelers and slower routes to others. Apparently, as hypothesized in the route-choice game, agents switch routes frequently enough to avoid any one of them being considered a free rider or a saint. This also reflected in the high stability scores for most network design parameters we investigated.

The important conjecture raised by Helbing et al. (2005) regarding the functionality of a system-optimal route guidance system, appears in hindsight to be much more complicated than was initially considered. While the implementation of punishments and rewards simplifies maintaining the efficiency of the network, it seems that the impacts of recommendation allocation policies and network design were not considered profoundly enough. As expected, punishment as an incentive has a strong impact over network efficiency - a kind of hammer effect – quite blunt but effective. For this reason, in earlier runs we could not find differences between the 'punishment + rewards' and 'punishment only' scenarios. This clearly explains the apparent operational success of congestion charging schemes. Rewards also maintain a positive impact over efficiency, albeit to a lesser degree than punishments as they are also on average much smaller per capita. In many of the examined cases, hard policies can be avoided altogether with the proper

---

[2] Tikkun Olam is a Jewish concept that means people bear responsibility not only towards their own material welfare, but also to the welfare of the entire society (Shatz et al., 1997).

consideration of recommendation allocation policies in the ATIS application.

Regarding recommendation allocation policies, although the dynamic 'Reformer' was the best performing policy, its implementation in reality is doubtful. It seems too optimistic to assume that all travelers will respond similarly, and therefore ATIS will not necessarily be able to determine the required signaling sequence in concordance with the cooperation rates and experiential travel times. Reformer is a possible route recommendation allocation policy that could be investigated for fully autonomous vehicles (Automation level 5) under the assumption that V2I telecommunications will become ubiquitous. In the meantime, simpler static allocation designs like 'Queue' can be further investigated on multi-arc directed and decentralized graphs that are more representative of complex urban road networks.

We found that network design parameters (in our simplified context) to have a significant influence over the efficiency of the ATIS. Moreover, it is important to consider in what space the network is operating – relative or absolute dominance regarding the gaps between AFRSO and AFRUE. Absolute dominance almost always promotes high efficiency and stability as drivers adapt easily to the route recommendations without strong resistance to their self-interest. In contrast, relative dominance is much trickier and requires more elaborate manipulations based on recommendation allocation policies and in some cases application of incentives to maintain the network in near-SO states.

Future research directions: First, we only accounted for homogenous bounded-rational agents with equal values of time. Future research should also introduce greater agent heterogeneity to verify if SO can still be maintained if agents respond differently to the same set of recommendations. Moreover, different behavioral response models beyond SAW should also be tested. One interesting aspect is to introduce inertia which was found in various studies to influence route-choice (Cantillo et al., 2007; Cherchi, 2009; Chorus, 2014).

Second, the simplicity of the road network in question limits our ability to generalize our results to more elaborate road networks. More simulations are needed on more network designs including: more complicated network structures (number of routes and ODs), congestion effects, diverse departure times, cost functions and within day dynamics framework, as done by Varga (2014a). Preliminary simulations we carried out in more complex networks show that achieving SO becomes more difficult as the number of routes per OD increases, as asserted by Jahn et al. (2005). However, with multiple ODs system-optimal self-organization is unlikely to emerge spontaneously as cooperation has to include both the routes and OD levels. Another intervention broker to govern traffic between ODs is therefore likely needed. Perhaps, in this case hard policies cannot be avoided and a fair 'traffic light' concept for reconciling between opposing traffic streams should be thought of. A good starting point to this effect would be the Braess paradox network where some degree of overlap exists between the traffic streams.

Third, more work is needed to study the impact of non-recurring congestion on the efficiency and stability of the proposed ATIS design. While it is expected that in the far-future the number of severity of accidents is likely to decrease with the introduction of autonomous vehicles, in the near future road safety will still be a significant stumbling block. Moreover, day-to-day changes globally affecting the network behavior, such as bad weather, are also factors that need to be accounted for in future research (as asserted by Lindsey et al., 2014).

Fourth, additional investigations regarding the route recommendation allocation policies are vital. In this paper, we assume all agents perceive and react to the route recommendations in the same way. However, in reality, this is not necessarily true, and nondiscrimination between the agents, as well as system stability, may be very hard to achieve. More advanced allocations, that take into account other characteristics of the road users, will be able to learn their reactions to system optimal ATIS and treat all road users in a way that their willingness to cooperate will be maximized, creating a state in which stability is maximized and discrimination between the agents is minimized. Another important research direction is the computation time of calculating the system optimal assignment and the allocation of users according to the system optimal ATIS. Currently brute force is used to find a global optimum but online optimization approaches could be a very interesting way forward perhaps using swarm like algorithms (e.g. bee colony optimization; Jovanović et al., 2017).

Lastly, beyond simulations, results need to be tested with real human beings to estimate the parameters of rationally bounded models. Empirical respondent data can be derived from economic multiplayer game-experiments, and this data could, in turn be used in agent-based models. In addition, parameters can be used to build game-based models of mixed populations of human and agents (Klein and Ben-Elia, 2016). Nevertheless, the ABM results can serve as a benchmark to test initial hypotheses and estimate effect sizes (and prior statistical power) for effectively designing experiments with human agents that are always costly. We recently used the outputs of a simulation with N = 10 agents to estimate the necessary sample sizes for a future experiment planned.

Notwithstanding the concurrent limitations, we believe our study shows the importance of further investigations of out-of-the-box approaches for managing congestion. In particular, during a period of transition from human controlled to machine controlled vehicles an SO-ATIS could be tested for possible implementation in existing routing apps operated via drivers' smartphones revealing the degrees that different classes of drivers would be willing to deviate from their shortest paths to more efficient and fair uses of pooled network resources. As the Internet of Things (IoT) proliferates and V2V and V2I telecommunication architectures mature, response to system-optimal routing could be possibly coded into the algorithms controlling the motion of fully autonomous vehicles and given robust recommendation allocation policies that maintain fairness between travelers, can be determined for real urban networks in future research, such an ATIS, could well be embraced by citizens, manufacturers and regulators alike.

Acknowledgments

Space-Time and Mobility, Nanjing. The comments and suggestions of three anonymous reviewers are highly appreciated.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.trc.2017.11.007.

References

Agogino, A.K., Tumer, K., 2008. Analyzing and visualizing multiagent rewards in dynamic and stochastic domains. Auton. Agent. Multi. Agent. Syst. 17, 320–338.

Arnott, R., de Palma, A., Lindsey, R., 1993. A structural model of peak-period congestion: a traffic bottleneck with elastic demand. Am. Econ. Rev. 83, 161–179.

Arnott, R., De Palma, A., Lindsey, R., 1996. Information and usage of free-access congestible facilities with stochastic capacity and demand. Int. Econ. Rev. (Philadelphia) 181–203.

Arnott, R., De Palma, A., Lindsey, R., 1991. Does providing information to drivers reduce traffic congestion? Transp. Res. Part A Gen. 25, 309–318.

Avineri, E., Waygood, O.E., 2013. Applying valence framing to enhance the effect of information on transport-related carbon dioxide emissions. Transp. Res. Part A Policy Pract. 48, 31–38. http://dx.doi.org/10.1016/j.tra.2012.10.003.

Aumann, R.J., 1974. Subjectivity and correlation in randomized strategies. J. Math. Econ. 1, 67–96. http://dx.doi.org/10.1016/0304-4068(74)90037-8.

Babes, M., De Cote, E.M., Littman, M.L., 2008. Social reward shaping in the prisoner's dilemma. In: Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 3. International Foundation for Autonomous Agents and Multiagent Systems, pp. 1389–1392.

Balmer, M., Rieser, M., 2009. MATSim-T: Architecture and simulation times. Multi-Agent Syst. Traffic Transp. Eng. http://dx.doi.org/10.1140/epjb/e2008-00153-6.

Bamberg, S., Schmidt, P., 2003. Incentives, morality, or habit? Predicting students' car use for university routes with the models of Ajzen, Schwartz, and Triandis. Environ. Behav. 35, 264–285.

Bazzan, A.L.C., Klügl, F., 2014. A review on agent-based technology for traffic and transportation. Knowl. Eng. Rev. 29, 375–403.

Ben-Elia, E., Di Pace, R., Bifulco, G.N., Shiftan, Y., 2013. The impact of travel information's accuracy on route-choice. Transp. Res. Part C Emerg. Technol. 26, 146–159. http://dx.doi.org/10.1016/j.trc.2012.07.001.

Ben-Elia, E., Erev, I., Shiftan, Y., 2008. The combined effect of information and experience on drivers' route-choice behavior. Transportation (Amst) 35, 165–177. http://dx.doi.org/10.1007/s11116-007-9143-7.

Ben-Elia, E., Ettema, D., 2011. Rewarding rush-hour avoidance: a study of commuters' travel behavior. Transp. Res. Part A Policy Pract. 45, 567–582. http://dx.doi.org/10.1016/j.tra.2011.03.003.

Ben-Elia, E., Shiftan, Y., 2010. Which road do I take? A learning-based model of route-choice behavior with real-time information. Transp. Res. Part A Policy Pract. 44, 249–264. http://dx.doi.org/10.1016/j.tra.2010.01.007.

Blum, A., Even-Dar, E., Ligett, K., 2006. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. In: Proceedings of the Twenty-Fifth Annual ACM Symposium on Principles of Distributed Computing. ACM, pp. 45–52.

Boyd, R., Richerson, P.J., 1988. The evolution of reciprocity in sizable groups. J. Theor. Biol. 132, 337–356. http://dx.doi.org/10.1016/j.jtbi.2009.01.004.

Browning, L., Colman, A.M., 2004. Evolution of coordinated alternating reciprocity in repeated dyadic games. J. Theor. Biol. 229, 549–557. http://dx.doi.org/10.1016/j.jtbi.2004.04.032.

Cantarella, G., Cascetta, E., 1995. Dynamic process and equilibrium in transportation network: towards a unifying theory. Transp. Sci. A 25, 305–329.

Cantillo, V., de Ortúzar, J.de D., Williams, H.C.W.L., 2007. Modeling discrete choices in the presence of inertia and serial correlation. Transp. Sci. 41, 195–205.

Cascetta, E., 1989. A stochastic process approach to the analysis of temporal dynamics in transportation. Transp. Res. Part B-Methodological 23, 1–17. http://dx.doi.org/10.1016/0191-2615(89)90019-2.

Cherchi, E., 2009. Modelling individual preferences: state of the art, recent advances and future directions.

Chorus, C.G., 2014. Risk aversion, regret aversion and travel choice inertia: an experimental study. Transp. Plan. Technol. 37, 321–332.

Chorus, C.G.C., Molin, E.J.E., van Wee, B., Wee, B.Van, 2006. Travel information as an instrument to change car-drivers' travel choices: a literature review. Eur. J. Transp. Infrastruct. Res. 6, 335–364.

Colman, A.M., Browning, L., 2009. Evolution of cooperative turn-taking. Evol. Ecol. Res. 11, 949–963.

Como, G., Savla, K., Acemoglu, D., Dahleh, M.A., Frazzoli, E., 2013. Stability analysis of transportation networks with multiscale driver decisions. SIAM J. Control Optim. 51, 230–252.

Davies, N., 2012. What are the ingredients of successful travel behavioural change campaigns? Transp. Policy 24, 19–29.

Dawes, R.M., 1980. Social Dilemmas. Annu. Rev. Psychol. 31, 169–193. http://dx.doi.org/10.1146/annurev.ps.31.020180.001125.

Dixit, V.V., Ortmann, A., Rutstr, E., Ukkusuri, S.V., Rutstrom, E., Ukkusuri, S.V., 2015. Understanding transportation systems through the lenses of experimental economics: a review. Available SSRN 2546881, 1–37.

Emmerink, R.H.M., Axhausen, K.W., Nijkamp, P., Rietveld, P., 1995a. Effects of information in road transport networks with recurrent congestion. Transportation (Amst) 22, 21–53.

Emmerink, R.H.M., Axhausen, K.W., Nijkamp, P., Rietveld, P., 1995b. The potential of information provision in a simulated road transport network with non-recurrent congestion. Transp. Res. Part C Emerg. Technol. 3, 293–309.

Emmerink, R.H.M., Verhoef, E.T., Nijkamp, P., Rietveld, P., 1998. Information policy in road transport with elastic demand: some welfare economic considerations. Eur. Econ. Rev. 42, 71–95.

Epstein, J.M., 1999. Agent-based computational models and generative social science. Complexity 4, 41–60.

Erev, I., Ert, E., Roth, A.E., 2010. A choice prediction competition for market entry games: an introduction. Games 1, 117–136. http://dx.doi.org/10.3390/g1020117.

Eriksson, L., Garvill, J., Nordlund, A.M., 2006. Acceptability of travel demand management measures: the importance of problem awareness, personal norm, freedom, and fairness. J. Environ. Psychol. 26, 15–26. http://dx.doi.org/10.1016/j.jenvp.2006.05.003.

Ettema, D., Knockaert, J., Verhoef, E., 2010. Using incentives as traffic management tool: empirical results of the "peak avoidance" experiment. Transp. Lett. 2, 39–51. http://dx.doi.org/10.3328/TL.2010.02.01.39-51.

Ettema, D., Timmermans, H., 2006. Costs of travel time uncertainty and benefits of travel time information: conceptual model and numerical examples. Transp. Res. Part C Emerg. Technol. 14, 335–350.

Fagnant, D.J., Kockelman, K., 2015. Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. Transp. Res. Part A Policy Pract. 77, 167–181.

Finucane, M.L., Alhakami, A., Slovic, P., Johnson, S.M., 2000. The affect heuristic in judgments of risks and benefits. J. Behav. Decis. Mak. 13, 1–17. http://dx.doi.org/10.1002/(SICI)1099-0771(200001/03)13:1<1::AID-BDM333>3.0.CO;2-S.

Fogg, B.J., 2002. Persuasive technology: using computers to change what we think and do. Ubiquity 2002, 80–87. http://dx.doi.org/10.1145/764008.763957.

Geller, E.S., 1989. Applied behavior analysis and social marketing: an integration for environmental preservation. J. Soc. Issues 45, 17–36.

Gilovich, T., Griffin, D., Kahneman, D., 2002. Heuristics and Biases: The Psychology of Intuitive Judgment. Cambridge University Press.

Gini, C., 1912. Variabilità e mutabilità. Repr. Mem. di Metodol. Stat. In: Pizetti, E., Salvemini, T. (Eds.). Rome Libr. Eredi Virgilio Veschi 1.

Goodwin, P.B., 1996. Empirical evidence on induced traffic: a review and synthesis. Transportation (Amst) 23, 35–54.

Guiliano, G., 1995. Land Use Impacts of Transportation Investments-Highway and Transit. In: Hanson, S., Giuliano, G. (Eds.), The Geography of Uraban Transportation. Guilford Press, New York, NY, USA, pp. 305–340.

Handy, S., Weston, L., Mokhtarian, P.L., 2005. Driving by choice or necessity? Transp. Res. Part A Policy Pract. 39, 183–203. http://dx.doi.org/10.1016/j.tra.2004.09.002.

Hardin, G., 1968. The Tradegy of the Commons. Science (80-.) 162, 1243–1248. http://dx.doi.org/10.1126/science.162.3859.1243.

He, F., Yin, Y., Shirmohammadi, N., Nie, Y.M., 2013. Tradable credit schemes on networks with mixed equilibrium behaviors. Transp. Res. Part B Methodol. 57, 47–65. http://dx.doi.org/10.1016/j.trb.2013.08.016.

Helbing, D., Balietti, S., 2013. How to Do Agent-Based Simulations in the Future: From Modeling Social Mechanisms to Emergent Phenomena and Interactive Systems Design. In: Helbing, D. (Ed.), Social Self-Organization: Agent-Based Simulations and Experiments to Study Emergent Social Behavior, Understanding Complex Systems. Springer, Heidelberg, pp. 25–70.

Helbing, D., Schonhof, M., Stark, H.-U., Holyst, J.a., 2005. How individuals learn to take turns: emergence of alternating cooperation in a congestion game and the prisoner's dilemma. Adv. Complex Syst. 8, 87–116. http://dx.doi.org/10.1142/S0219525905000361.

Hills, T.T., Todd, P.M., Lazer, D., Redish, A.D., Couzin, I.D., Group, C.S.R., 2015. Exploration versus exploitation in space, mind, and society. Trends Cogn. Sci. 19, 46–54.

Horowitz, J.L., 1984. The stability of stochastic equilibrium in a two-link transportation network. Transp. Res. Part B 18, 13–28. http://dx.doi.org/10.1016/0191-2615(84)90003-1.

Jahn, O., Mohring, R.H., Schulz, A.S., Stier-Moses, N.E., 2005. System-optimal routing of traffic flows with user constraints in networks with congestion. Oper. Res. 53, 600–616. http://dx.doi.org/10.1287/opre.1040.0197.

Jang, S., An, Y., Yi, C., Lee, S., 2017. Assessing the spatial equity of Seoul's public transportation using the Gini coefficient based on its accessibility. Int. J. Urban Sci. 21, 91–107.

Jovanović, Aleksandar., Nikolić, Miloš., Teodorović, Dušan., 2017. Area-wide urban traffic control: a Bee Colony Optimization approach. Transp. Res. Part C: Emerg. Technol. 77, 329–350.

Kahneman, D., Tversky, A., 1984. Choices, values, and frames. Am. Psychol. 341–350.

Kaplan, A.M., Haenlein, M., 2010. Users of the world, unite! The challenges and opportunities of Social Media. Bus. Horiz. 53, 59–68. http://dx.doi.org/10.1016/j.bushor.2009.09.003.

Kemel, E., Paraschiv, C., 2013. Prospect Theory for joint time and money consequences in risk and ambiguity. Transp. Res. Part B Methodol. 56, 81–95. http://dx.doi.org/10.1016/j.trb.2013.07.007.

Klein, I., Ben-Elia, E., 2016. Emergence of cooperation in congested road networks using ICT and future and emerging technologies: a game-based review. Transp. Res. Part C Emerg. Technol. 72, 10–28. http://dx.doi.org/10.1016/j.trc.2016.09.005.

Levinson, D., 2003. The value of advanced traveler information systems for route choice. Transp. Res. Part C Emerg. Technol. 11, 75–87. http://dx.doi.org/10.1016/S0968-090X(02)00023-2.

Levy, N., Ben-Elia, E., 2016. Emergence of system optimum: a fair and altruistic agent-based route-choice model. Procedia Comput. Sci. 83, 928–933. http://dx.doi.org/10.1016/j.procs.2016.04.187.

Lindsey, R., Daniel, T., Gisches, E., Rapoport, A., 2014. Pre-trip information and route-choice decisions with stochastic travel conditions: theory. Transp. Res. Part B Methodol. 67, 187–207. http://dx.doi.org/10.1016/j.trb.2014.05.006.

Lu, X., Gao, S., Ben-Elia, E., Pothering, R., 2014. Travelers' day-to-day route choice behavior with real-time information in a congested risky network. Math. Popul. Stud. 21, 205–219. http://dx.doi.org/10.1080/08898480.2013.836418.

Lu, X., Gao, S., Ben-Elia, E., 2011. Information impacts on route choice and learning behavior in a congested network: experimental approach. Transp. Res. Rec. 2243, 89–98.

Mak, V., Rapoport, A., 2013. The price of anarchy in social dilemmas: traditional research paradigms and new network applications. Organ. Behav. Hum. Decis. Process. 120, 142–153. http://dx.doi.org/10.1016/j.obhdp.2012.06.006.

Mayeres, I., Ochelen, S., Proost, S., 1996. The marginal external costs of urban transport. Transp. Res. Part D Transp. Environ. 1, 111–130.

Meijer, S.A., Hofstede, G.J., 2003. Simulation games for improving the human orientation of production management. In: Current trends in production management, European series in industrial management, vol. 6. Shaker-Verlag Aachen, Aachen, pp. 58–64.

Nash, J., 1951. Non-cooperative games. Ann. Math. 54, 286–295.

O'Hare, S.J., Connors, R.D., Watling, D.P., 2016. Mechanisms that govern how the Price of Anarchy varies with travel demand. Transp. Res. Part B Methodol. 84, 55–80. http://dx.doi.org/10.1016/j.trb.2015.12.005.

Prashker, J.N., Bekhor, S., 2004. Route choice models used in the stochastic user equilibrium problem: a review. Transp. Rev. 24, 437–463. http://dx.doi.org/10.1080/0144164042000181707.

Rapoport, A., Gisches, E.J., Daniel, T., Lindsey, R., 2014. Pre-trip information and route-choice decisions with stochastic travel conditions: experiment. Transp. Res. Part B Methodol. 68, 154–172. http://dx.doi.org/10.1016/j.trb.2014.06.007.

Rapoport, A., Kugler, T., Dugar, S., Gisches, E.J., 2009. Choice of routes in congested traffic networks: experimental tests of the Braess Paradox. Games Econ. Behav. 65, 538–571. http://dx.doi.org/10.1016/j.geb.2008.02.007.

Rawls, J., 1999. A Theory of Justice. Harvard University Press.

Rothengatter, T., 1982. The effects of police surveillance and law enforcement on driver behaviour. Curr. Psychol. Rev. 2, 349–358.

Roughgarden, T., 2005. Selfish Routing and the Price of Anarchy. MIT press, Cambridge.

Rouwendal, J., Verhoef, E.T., 2006. Basic economic principles of road pricing: from theory to applications. Transp. policy 13, 106–114.

Schrank, D., Eisele, B., Lomax, T., 2012. TTI's 2012 Urban Mobility Report. Texas A&M Transp. Institute. Texas A&M Univ, Syst.

Schulz, A.S., Stier-Moses, N., 2003. On the performance of user equilibria in traffic networks. In: Proc. 14th Symp. Discrete Algorithms. p. 86.

Schulz, A.S., Stier-Moses, N.E., 2006. Efficiency and fairness of system-optimal routing with user constraints. Networks 48, 223–234. http://dx.doi.org/10.1002/net.20133.

Selten, R., Chmura, T., Pitz, T., Kube, S., Schreckenberg, M., 2007. Commuters route choice behaviour. Games Econ. Behav. 58, 394–406. http://dx.doi.org/10.1016/j.geb.2006.03.012.

Shalizi, C.R., 2006. Methods and techniques of complex systems science: an overview. Complex Syst. Sci. Biomed. 33–114.

Shang, W., Han, K., Ochieng, W., Angeloudis, P., 2017. Agent-based day-to-day traffic network model with information percolation. Transp. A Transp. Sci. 13, 1–29. http://dx.doi.org/10.1080/23249935.2016.1209254.

Shatz, D., Waxman, C.I., Diament, N.J., 1997. Tikkun Olam: Social Responsibility in Jewish Thought and Law. Jason Aronson (Incorporated).

Stark, H.U., 2010. Dilemmas of partial cooperation. Evolution (N.Y.) 64, 2458–2465. http://dx.doi.org/10.1111/j.1558-5646.2010.00986.x.

Steg, L., 2005. Car use: lust and must. Instrumental, symbolic and affective motives for car use. Transp. Res. Part A Policy Pract. 39, 147–162. http://dx.doi.org/10.1016/j.tra.2004.07.001.

Thaler, R.H., Sunstein, C.R., 2008. Nudge. Yale University Press.

The World Bank, 2016. Urban Population Data [WWW Document]. URL http://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS?display=map (accessed 12.22.16).

Timms, C., 2013. Has London's Congestion Charge Worked? Br. Broadcast. Coop.

Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: heuristics and biases. Science (80-.) 185, 1124–1131.

Varga, L.Z., 2014a. On intention-propagation-based prediction in autonomously self-adapting navigation. In: 2014 IEEE Eighth International Conference on Self-Adaptive and Self-Organizing Systems Workshops (SASOW), 2014, IEEE, pp. 38–43.

Varga, L.Z., 2014b. Online routing games and the benefit of online data. In: Proc. Eighth International Workshop on Agents in Traffic and Transportation, at 13th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2014). pp. 88–95.

Vickrey, W.S., 1969. Congestion theory and transport investment. Am. Econ. Rev. 59, 251–260.

Viegas, J.M., 2001. Making urban road pricing acceptable and effective: searching for quality and equity in urban mobility. Transp. Policy 8, 289–294. http://dx.doi.org/10.1016/S0967-070X(01)00024-5.

Wahle, J., Bazzan, A.L.C., Klügl, F., Schreckenberg, M., 2002. The impact of real-time information in a two-route scenario using agent-based simulation. Transp. Res. Part C Emerg. Technol. 10, 399–417. http://dx.doi.org/10.1016/S0968-090X(02)00031-1.

Wardrop, J., 1952. Some theoretical aspects of road traffic research communication networks. Proc. - Inst. Civ. Eng. 2 (2), 325–378.

Wei, F., Jia, N., Ma, S., 2016. Day-to-day traffic dynamics considering social interaction: from individual route choice behavior to a network flow model. Transp. Res. Part B Methodol. 94, 335–354. http://dx.doi.org/10.1016/j.trb.2016.10.002.

Zhao, J., Szilagyi, M.N., Szidarovszky, F., 2008. An n-person battle of sexes game. Phys. A Stat. Mech. Appl. 387, 3669–3677. http://dx.doi.org/10.1016/j.physa.2007.09.053.

Zhu, F., Ukkusuri, S.V., 2017. Efficient and fair system states in dynamic transportation networks. Transp. Res. Part B Methodol. 104, 272–289. http://dx.doi.org/10.1016/j.trb.2017.07.010.