Saffre, F., Gianini, G., Hildmann, H., Davies, J., Bullock, S., Damiani, E., & Deneubourg, J.L. (2019). Long-term memory-induced synchronisation can impair collective performance in congested systems. *Swarm Intelligence*, 13(2), 95-114.
<https://doi.org/10.1007/s11721-019-00164-z>

Peer reviewed version

Link to published version (if available):
[10.1007/s11721-019-00164-z](https://doi.org/10.1007/s11721-019-00164-z)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Springer Nature at <https://link.springer.com/article/10.1007/s11721-019-00164-z>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Long-term memory-induced synchronisation can impair collective performance in congested systems

**F. Saffre · G. Gianini · H. Hildmann ·
J. Davies · S. Bullock · E. Damiani ·
J.-L. Deneubourg**

the date of receipt and acceptance should be inserted later

Abstract We investigate the hypothesis that long-term memory in populations of agents can lead to counter-productive emergent properties at the system level. Our investigation is framed in the context of a discrete, one-dimensional road-traffic congestion model: we investigate the influence of simple cognition in a population of rational commuter agents that use memory to optimize their departure time, taking into account congestion delays on a previous trips. Our results differ from the well-known minority game in that

F. Saffre
EBTIC (Emirates ICT Innovation Centre), Khalifa Univ., P.O. 127788, Abu Dhabi, UAE
BT Research and Innovation, Adastral Park, Martlesham Heath, IP5 3RE, United Kingdom
E-mail: fabrice.saffre@bt.com

G. Gianini
EBTIC (Emirates ICT Innovation Centre), Khalifa Univ., P.O. 127788, Abu Dhabi, UAE
Università degli Studi di Milano, Comp. Sci. Dept., via Bramante 65, 26013 Crema, Italy
E-mail: gabriele.gianini@unimi.it

H. Hildmann
BrainCreators, Prinsengracht 697, 1017 JV Amsterdam, The Netherlands
UC3M (Universidad Carlos III de Madrid), Léiganes (Madrid), 28911, Spain
EBTIC (Emirates ICT Innovation Centre), Khalifa Univ., P.O. 127788, Abu Dhabi, UAE
E-mail: hanno@cypherpunx.org

J. Davies
BT Research and Innovation, Adastral Park, Martlesham Heath, IP5 3RE, United Kingdom
E-mail: john.nj.davies@bt.com

S. Bullock
University of Bristol, Dep. of Comp. Sci., Woodland Rd., Bristol, BS8 1UB, United Kingdom
E-mail: seth.bullock@bristol.ac.uk

E. Damiani
Khalifa University of Science and Technology, P.O. Box 127788 Abu Dhabi, UAE
E-mail: ernesto.damiani@kustar.ac.ae

J.-L. Deneubourg
Université Libre de Bruxelles, Unité d'Ecologie Sociale, CP 231, B-1050 Bruxelles, Belgium
E-mail: jldeneub@ulb.ac.be

crowded slots do not carry any explicit penalty. We use Markov Chains analysis to uncover fundamental properties of this model and then use the gained insight as a benchmark. Then, using Monte Carlo simulations, we study two scenarios: one in which “myopic” agents only remember the outcome (delay) of their latest commute, and one in which their memory is practically infinite. We show that there exists a trade-off, whereby myopic memory reduces congestion but increases uncertainty, whilst infinite memory does the opposite. We evaluate performance against the optimal distribution of departure times (i.e. where both delay and uncertainty, are minimized simultaneously). This optimal but unstable distribution is identified using a Genetic Algorithm.

Keywords Multi-agent · Congestion · Synchronisation · Memory · Emergence · Optimisation · Markov Chain · Monte Carlo Simulation · Genetic Algorithms

PACS PACS code1 · PACS code2 · more

Mathematics Subject Classification (2010) MSC code1 · MSC code2 · more

1 Introduction

Movement is a fundamental aspect of the physical universe, from the entirely passive motion of molecules in a gas, through the “semi-active” displacement of microscopic life-forms such as bacteria or plankton, to the intentional movement of more complex organisms such as insects, fish, birds or mammals. One defining characteristic of movement is that, independently of the way in which it is achieved, it is necessarily mediated or constrained by the environment in which it takes place. Furthermore, when considering a number of entities in motion, even when such motion is governed by simple physical laws, complex and chaotic patterns can emerge as a result of their interactions, e.g. Rayleigh-Taylor instability (Sharp, 1984).

In biology, the various phenomena resulting from the interactive movement of several entities have been studied under the umbrella term of collective motion (Vicsek and Zafeiris, 2012). In both physics and biology, relevant and universally occurring global patterns can be expected to emerge as soon as the interaction among moving entities is governed by specific rules (Deutsch et al., 2012), and in the context of road-traffic, congestion is one such pattern.

Motivation

The motivation for this work is two-fold. On the one hand, we wish to demonstrate, in a concrete and socio-economically relevant case, that there are situations in which more advanced cognitive capabilities (long-term memory) are detrimental to collective performance as they promote the emergence of a sub-optimal stable-state. On the other, by documenting our investigations into this

phenomenon, we hope to improve our ability to model, predict and mitigate traffic congestion in a “smart city” application context (d’Aquin et al., 2015; Davies and Fisher, 2015), where it poses a significant challenge.

Any reduction in road congestion not only reduces journey times for travellers but, perhaps more importantly, reduces one of the major factors responsible for poor air quality, namely vehicle emissions (Wahle et al., 1999).

Contribution of this article

We consider *rational* agents that use stored information about past journeys as evaluated against some fitness function, such as “*being on time*”. We present a simple probabilistic model to study the influence of cognitive effects, i.e. learning from experience, over recurrent congestion dynamics (e.g. “rush hour”). We show that cognitive effects actually contribute to perpetuating congestion. Within our model, agents effectively displace the emerging congestion (instead of reducing it), and they do so by progressively learning and adapting to the average delay that they themselves experience. Our results clearly indicate that there is a trade-off between reducing congestion and increasing uncertainty about the outcome of a planned journey.

Our model is fundamentally different from e.g. the well-known minority game (presented in Section 2.2.3) in that crowded slots do not carry any explicit penalty. Indeed, slow travel during a congested period can be an optimal strategy with respect to the agents’ objective of arriving on time, provided they factor the delay into their calculations. In fact, this is exactly what we observe in the “infinite memory” scenario. The overall reduction of congestion through spreading the load is a secondary objective of an external entity (e.g. the city authorities) and is not explicitly shared by the commuter agents.

Our results open the way for subsequent investigations into how to mitigate congestion, which is caused indirectly by the agents’ ability to use their memory of past journeys to refine their strategy and meet their objective of arriving on time. In theory, it is possible to achieve this with a limited (and, depending on the cost attributed to congestion, acceptable) increase in journey time variability. Furthermore, departures can be scheduled according to an optimal distribution, the shape of which is determined by the relative weight of both costs (congestion and delay) and traffic dynamics.

2 State-of-the-art

2.1 Urn models

Combinatorics has been described by Gian Carlo Rota as the art of counting in how many ways one can put [coloured] balls in [coloured] urns (Damiani et al., 2009). Urn model representations provide a powerful way to study complex systems involving autonomous and semi-autonomous entities, as they draw a

clear conceptual link between the representation choices and the behavioral properties the model can express (often called the model's *Statistics*).

This kind of representation is at the core of many models (Young et al., 2015), including the model known as Maxwell-Boltzmann Statistics (Maxwell, 1879, 1890), for modeling non-interacting entities; it was also at the basis of the models known as Bose-Einstein Statistics (Bose, 1920; Einstein, 1924) and Fermi-Dirac Statistics (Fermi, 1926), modeling systems in which entities are, respectively, allowed or forbidden to occupy the same state / box at the same time. Comprehensive reviews of urn models can be found in (Johnson and Kotz, 1977; Mahmoud, 2008).

Another point in favor of “balls-in-urns” models for representing complex systems is that the temporal evolution of different arrangements can be represented by means of Markov Chains satisfying some basic properties:

1. Any balls-and-urns arrangement has a finite number of potential successors, called states, and the Statistics defines which are acceptable.
2. At each instant, the model can move from its present arrangement to any valid state or remain in the current one.
3. The probability of doing so is represented in a square transition matrix. Being probabilities, entries in each row must sum to unity.

2.2 Movement dynamics and congestion

Movement dynamics have been studied extensively (Chowdhury et al., 2000; Helbing, 2001; Kerner, 2009; Treiber and Kesting, 2012) and a host of models (cf. Bellomo and Dogbe (2011)) of various complexity has been proposed for numerous types of motion, e.g. Daganzo and UC Berkeley (1994). In the context of human behaviour, common examples are empirical studies on traffic networks (Schreckenberg and Selten, 2013), work on crowd motion, pedestrian flow and evacuation dynamics (Schadschneider et al., 2009) and investigations into road-traffic congestion (Quill, 2008). At its most general, congestion is the result of a flow exceeding the nominal capacity of its carrier (i.e. its bounding environment). With regard to road-traffic, the above is something of a simplification since jam formation is not entirely deterministic and, as such, cannot be guaranteed to occur when and only when a particular traffic capacity is reached. Indeed, tiny perturbation in the flow can already have repercussions leading to a drastic, localized drop in speed (Kerner and Rehborn, 1997).

In IP networks, congestion manifests itself as a need to drop packets in order to avoid exceeding buffering capacity which results in disruptive phenomena. In road networks, the most familiar effect of congestion is the traffic jam. On a fundamental level, all or most types of congestion have similar characteristics, most noticeably a strong non-linear component. This non-linearity implies that, for certain parameter values, small changes can have a large impact on the behaviour of the system.

Congestion effects will often be negligible below a certain threshold flow density, then rapidly increase in its vicinity, in a way reminiscent of a phase

transition. As a result of this non-linearity, controlling flow density in the traffic domain offers the possibility of significantly lowering congestion for a comparatively small price in terms of wider dispersion around expected transit travel time and therefore slightly increased transit travel time uncertainty.

2.2.1 Congestion

In their paper, Mahmassani and Chang (1986) use system dynamics approaches to study the effect of commuters deciding to alter their departure times in response to congestion delays experienced on previous trips. In particular, they investigate the conditions for the system to stabilise and show that the tolerance for error (difference between desired and actual arrival time), combined with the volume of traffic, are the most critical parameters. In short: if the error tolerance is low, the system struggles to find a stable regime. Of particular interest for the present study is their conclusion that, somewhat counter-intuitively, a weighted long-term memory does not have a clearly beneficial effect compared to what they refer to as a “myopic” learning algorithm in which only the outcome of the most recent trip is taken into account.

In their recent paper Xiao and Lo (2016) explore the possible influence of social networks on the process whereby commuters “learn” about the congestion delay they are likely to incur on their way to work, depending on their departure time. In their model the travellers progressively acquire very detailed knowledge, including not only the average delay associated with a certain departure time, but also its variance. The chosen utility function assigns specific weights to the three types of “costs” involved: the duration of the trip and penalties for early and late arrival. Memory is introduced in the form of weighted updates to the commuters’ knowledge (“belief” system), taking into account their own personal experience as well as that of those of their peers to which they are socially connected. The study of congestion dynamics is beyond the remit of Xiao and Lo’s paper. In their paper they use the classical “bottleneck” model (Vickrey, 1969) to calculate delays.

2.2.2 Cell Transmission Model

We model a stretch of road as a sequence of cells (*urns*) representing road segments. Traffic (*balls*) move through these cells along predefined paths. As such, our model is a simple version of the Cell Transmission Model (CTM), originally proposed by Daganzo (1994) in which road-traffic is modelled on a macroscopic level, as opposed to microscopic models, considering individual cars and the interactions among them (Wahle et al., 1999). Cells are connected in a restricted number of ways (Daganzo and UC Berkeley, 1994) and the traffic flow through them is governed by two functions defining how the traffic may enter and leave each cell. One benefit of using this model is the fact that it is comparatively easy to update and implement (Chuo et al., 2016).

2.2.3 Minority Game

A model related to our own, but different in many aspects, is the Game Theory model known as Minority Game (Challet and Zhang, 1997, 1998) (for a comprehensive treatment see for instance (Challet, 2006; Challet et al., 2013)).

In the inspiring prototypical example of the game - the El Farol bar setting (Arthur, 1994) - each member of a set of selfish agents endowed with bounded rationality tries to avoid attending the bar in overcrowded evenings. The agents do not have the opportunity to communicate with one another and do not have ahead information about the behaviour of the other agents. The players make decisions based on the common knowledge of the past record. There are two strategies / sides: the *goers* and the *non-goers*. The agents that end up on the minority side win. The setup can be easily reformulated so as to include the avoidance of time slots in place of space slots and so on. Variants have been studied that included differently structured rewards (binary reward 0–1 structure or reward inversely proportional to the number of winners), different memory depths (in number of rounds $m = 1, 2, \dots, +\infty$) and number of strategies available to each player (in the binary reward structure, a subset of the 2^m possible strategies). The distinctive features of the game have been found insensitive to the memory depth (Cavagna, 1999). The game has been investigated extensively through numerical simulations (Challet and Zhang, 1998; Johnson et al., 1999; Savit et al., 1999), then solved analytically - in the infinite number of agents limit - through a mapping onto a spin-glass model (Challet and Marsili, 1999; Challet et al., 2000; Marsili et al., 2000).

The main difference between the Minority Game definition and our definition is that, in the former, the objective function - i.e. the one on which the agents are basing their decision - is explicitly and exclusively linked to the avoidance of crowded time-slots or locations. By contrast, in our model, the primary objective function, is directly related only to their timely arrival at destination. More specifically, in our setting, the agents choose in which time slot to start the journey: it is true that ending up in a crowded time slot for the start increases the likelihood of incurring in longer delays due to congestion; however delays incurred due to congestion have no direct explicit cost to the agents. The agent will receive a positive feed-back from a given choice even if (s)he has spent a long time in a crowded road, provided that (s)he has arrived on time. This is why, in the unbounded memory scenario, the emerging stable-state is a synchronised departure which perpetuates congestion. The irrelevance of memory in the Minority Game (Cavagna, 1999), is another remarkable difference. In the context of our model, unbounded memory perpetuates congestion, whilst “myopic” short-term memory has a benign influence. Since the avoidance of crowded slots is not the explicit objective of the decision-making agents, memory does play a critical part in promoting synchronisation and therefore increasing congestion.

3 The model

3.1 Description

In line with the Cell Transmission Model (CTM), we employ a one-dimensional discrete model in order to study rigorously the dynamics of congestion. We restrict our investigations to the simplest possible case: a single traffic lane is represented as a chain of urns, i.e. the lane does not merge or divide and there are no junctions. Furthermore, there is only one lane of traffic. In this scenario, all entities (packets, vehicles, ...) have to travel from the same origin point at one end of a chain of discrete locations to the same destination point at the other end of the chain (using the same route), and share the same target arrival time. This can be conceived of as a number of balls moving through a sequence of “discrete” urns, i.e. adjacent locations, until all of them have reached the last one. Congestion is introduced by making the probability of one ball being moved from its current urn to the next urn an inverse non-linear function of the number of balls n in its current urn (however, entering a cell is always possible). For simplicity, and because our primary concern is the modelling of cognitive effects (as opposed to the modelling of road-traffic flow) we used a slightly modified version of the Hill function to simulate stochastic movement:

$$P_i = \frac{1}{1 + \left(\frac{i-1}{n_u}\right)^\alpha} \quad (1)$$

giving us the probability P_i of a ball i moving from one urn to the next. Parameter n_u is the threshold occupancy value for which the probability of a ball moving from one urn to the next is one half ($P_i = 0.5$), and α is the non-linearity. On every time-step, all balls present in the system are allowed one attempt to move to the next urn. Each attempt consists in drawing a random real number in the interval $[0; 1[$ (uniform distribution). If it is lower than the value of P_i calculated for each specific ball characterised by its index i (position in the urn), then it is transferred to the next urn, otherwise it does not move. Thus when $(i-1) > n_u$, the higher the value of α , the lower the probability P_i that the corresponding ball passes into the next urn. Furthermore the above expression ensures that at least the first ball from each urn ($i = 1$) moves to the next one on every time-step, since $P_{i=1}$ is always 1, independently of non-linearity. To avoid any permanent bias, the order in which the balls in a urn are considered (their index) is randomized on every time-step.

Finally, so as to avoid interaction between urns within the same time-step, they are considered one by one in the opposite direction to ball movement (i.e. from last to first urn).

Example

To illustrate the process, let us consider the case where: the total number of balls is $N = 3$, there are only three urns, and the parameter values for

equation 1 are $n_u = 2$ and $\alpha = 2$. On the first time-step, all three balls are in the first urn, written as $[3, 0, 0]$. Their probability to move to the second urn are respectively $P_1 = 1$, $P_2 = 0.8$, and $P_3 = 0.5$. At the end of the first step, the most likely configuration of the system is $[1, 2, 0]$, which is observed with frequency 0.5, followed closely by $[0, 3, 0]$, which is observed with frequency 0.4. Configuration $[2, 1, 0]$ will be observed with frequency 0.1.

Thus, in the first and most common outcome of round 1, one ball “falls victim” to congestion effects. Figure 1 shows all possible configurations reachable at the end of each round and all the possible transitions between them.

Fig. 1 The reachable configurations for a system comprised of 3 balls and 3 urns, and the possible transitions between them according to the rules of our model applied for four steps (S0 to S4) and without distinguishing between balls (Bose-Einstein Statistics).

If congestion is impossible ($n = 1$) or unlikely ($n \ll n_u$) throughout the process then, if there are U urns in total, all balls will complete their journey in $U - 1$ steps. If they are governed by “rational” agents sharing the same target arrival time T , they should leave the first urn at $T - U + 1$. For the remainder of this paper we will assume that $T = U - 1$, so that the best time to leave the first urn in the absence of congestion is always zero.

3.2 Model validation

We studied the model introduced above both using Monte Carlo simulation and Markov Chain modeling: the latter was used to investigate the general characteristics of the dynamics, the former was used to investigate more complex scenarios, such as the ones described in the next sections. The two approaches complement one another and provide a cross check of the consistency of one another’s results in the common domains of study.

Markov Chains are appropriate for modeling stochastic processes endowed with the so called Markov property (i.e. the future state of the system depends only on the present state, e.g. the probability governing the transition of the balls from one urn to the next). This allows the representation of a process in terms of a directed weighted graph. In such a graph, the nodes are the states and the arcs represent transitions with associated probabilities (called a Markov Chain). The algebraic properties of the corresponding state-to-state transition probability matrix allow the characterization of the relevant features of the process, including the asymptotic behavior of the system. This facilitates the study of short / medium term behavior of a system with high accuracy.

In the Monte Carlo simulation the process is basically a random walk over the state and transition graph; the quantities obtained by repeated runs of the simulation are therefore subject to statistical fluctuations. In the Markov Chain model probabilities are propagated directly and consequently fluctuations are avoided by construction.

In summary, we have the following parameters for the model:

- N , the total number of balls in the system;
- U , the number of discrete locations (urns) in the path;
- n_u , the threshold occupancy value and
- α , the non-linearity parameter.

In the remainder of this section we present results of the Markov Chain analysis in terms of general features of the dynamics (§3.3) and in terms of the distribution of the arrival times for simultaneously departing balls (§3.4) as a function of the parameters of the model. The cases of non-simultaneously departing balls is dealt with later in the context of the different models of adaptation and memory (§4); those results are obtained by running Monte Carlo simulations. With regard to departure, we assume that all the balls depart from the starting urn at the same time.

3.3 Propagation dynamics

Hereafter we will use the following naming convention for the different space and time granular elements: the path to be travelled consists in an array of U contiguous urns – that we call the *urn-array* – characterized, in our case, by the same dynamic constants n_u and α (no bottlenecks are present) determining $P(n|n_u, \alpha)$, except the last urn, for which the propagation probability is $P = 0$.

3.3.1 States at different granularities of the Markov Chain

From the point of view of the Markov Chain modeling, the evolution process is a discrete-time discrete-state-space process. The time and space elements can be considered at various nested resolutions / granularities: the transition probabilities among states at one resolution can be used to compute the transition probabilities among states at a coarser resolution.

The minimal resolution corresponds to the most elementary process, namely a ball attempting to pass to the next urn. Here, initial and final state correspond to the initial and final position (success or failure) for the ball, a process which is governed by equation 1. This process is repeated for all the balls from the penultimate urn to the first in the chain.

One can also describe the process at a coarser granularity by studying the update of a whole urn, as all the contained balls attempt to pass to the next urn; here the initial *urn-states* are the possible initial occupancy numbers of the current urn and the possible final occupancy numbers of the current urn (which influences the occupancy of the next urn).

The transition probabilities are computed as products of conditional probabilities of individual attempts based on equation 1. This is repeated starting from the penultimate urn in the array and going backward.

A full *pass* over all the urns of the array corresponds to the next upper granularity description: here the initial states are the possible configurations of the urn array. Figure 1 describes the process at the pass-granularity for the case where $N = 3$ and $U = 3$. For the sake of clarity, annotations of the transition probabilities from one pass-state to the next are not reported.

A process going from a departure state to the final state (all the balls at destination) is called a *round*. In this section we do not have any interest in distinguishing between rounds, since all the rounds, here, use the same dynamics and the same starting pass state (all the balls in the starting urn); that distinction will, however, be relevant for the following sections, where memory effects will be introduced, which modify the starting pass state of a round based on the outcomes of the previous rounds.

3.3.2 Pass state characteristics and transition dynamics

We analyze the behaviour of the system at *pass state* granularity. The pass states and their transition probabilities have been obtained by composition of the probabilities at finer granularities. We start by providing a list of connected observations along with the supporting evidence. The findings discussed are the following: the set of urns hosting the balls, during their propagation, always form a *connected chain*; its length increases gradually with time; given enough running space this relaxation process leads to a one-ball-per-urn procession-like configuration that is stable and moves forward at a rate of one urn per time-step. The speed of this relaxation process is controlled by the parameters n_u and α : the larger their values the slower the relaxation and the lower the variance of the arrival times.

To start, we recall that by the definition of the dynamics, at least one ball in each urn succeeds in moving on. This has an important consequence: if all the balls start from the same urn, the balls can adopt different configurations but during their journey they will always occupy contiguous urns. In other words the balls will form a wave or chain containing no gaps because even if all the balls in a urn leave that urn, at least one ball from the preceding urn will advance to fill the gap. As a consequence, in principle, the sequence

of urns s spanned by this chain can go from a single urn ($s = 1$, e.g. when all the balls are located in the starting urn) to $s = N$ where each of a set of contiguous urns hosts a single ball.

The stability of the distinct configurations varies. It is worth analyzing the two extreme cases. Consider the configuration with $s = 1$: it is easy to demonstrate, by using the chain product and equation 1, that the case in which all the N balls succeed in passing to the next urn is the least likely; consequently the lifetime of such an extreme chain configuration is rather short; soon the configuration will give rise to an $s = 2$ chain. At the opposite extreme, the configuration with $s = N$ is absolutely stable: the first ball in the leading urn will deterministically succeed in passing to the next urn, the same will happen to each of the other balls. The effect of a pass (while there is room ahead) will always be a translation of this configuration one urn forward. Should the system reach this configuration, then it would proceed deterministically: this chain configuration / shape would correspond to a *soliton*.

Although the no-gap condition limits the number of configurations that the wave / chain can assume, there is still a large number of configurations accessible to the propagating wave. To avoid repeatedly counting configurations that differ only by a simple translation, let us pin our reference frame to the leading urn (the foremost non-empty urn). It can be shown that the number of possible translation-discounted configurations is $F = 2^{N-1}$. For example those with $N = 5$, ordered lexicographically are shown in Figure 2.

The study of the process with unconstrained propagation by Markov Chain can be greatly simplified by reducing the process to a set of transitions among such translation-discounted configurations. The dynamics expressed by equation 1 have some other interesting consequences, which become apparent when the transition matrix is built. These interesting consequences are:

- Not all states can be reached for any starting state. Rather, states form a partial ordering with respect to the state-to-state transition relation. States representing chains of length $s = l$ can either transition to states with the same value of s or to states where $s = l - 1$. This property can be expressed compactly by ranking the states in lexicographic order: if we do so, transitions take place only between a state of length $s = l$ and a state with lower ranking either with $s = l$ or with $s = l - 1$. This is illustrated in Figure 2 by the transition matrix for $N = 5$.
- As the transition matrix shows, higher states can always reach lower-ranking states through one or more transitions, i.e. the set of states is connected “downward” (though not necessarily to *all* lower ranking states).

The system will eventually relax into the state with the lowest ranking and the maximum length with $s = N$. Since this state is an absorbing state of the Markov Chain it reproduces itself deterministically during propagation of the wave (i.e. it represents a soliton). This implies that, given enough time and space, the system will relax into a soliton of length $s = N$ for *any* starting connected chain configuration, independent of the parameter values.

Fig. 2 The transition probabilities (top) for $N = 5$, $n_u = 1$ and $\alpha = 2$ and all 16 formations (bottom right) expressed numerically (top) and as a heat-map (bottom left). All probabilities are below the leading diagonal - i.e. there is a lexicographic ordering over transitions. Transition probabilities on the leading diagonal are each less than or equal to the other non-zero values in their row - i.e. no states (apart from the absorbing soliton state) are likely to be maintained. The table (bottom right) represents the 16 possible variations the 5 balls can occur in, represented as the stretch s of urns they cover (flushed right).

The relaxation time, however, depends significantly upon the values of parameters n_u and α and increases with them: i.e. the larger their values the slower the relaxation and the lower the variance of the arrival times. We computed the transition matrix for several values of the parameters n_u and α .

As n_u and α increase the off-diagonal elements of the matrix decrease. For high enough values of the two parameters the sub-matrix of the states with $s > 1$ becomes approximately diagonal: as n_u and α increase the $s = 1$ state has a larger lifetime, but it is still ephemeral (its survival time is a negative exponential in the number of passes), while the other states tend to live longer (for the sake of exemplification, with $N = 5$, $n_u = 4$ and $\alpha = 10$ all the self-transition probabilities of the states with $s > 1$ are greater than 0.999). Thus if the parameter values are sufficiently high, the system transitions from one near-soliton to another near-soliton state: the lifetime of each of these soliton-like states being finite but large.

Fig. 3 Frequency distributions of the time of completion (last ball arriving) for different values for n_u and α , for $N = 5$ balls starting simultaneously at time $t = 0$ from urn $u = 1$. Increasing n_u positively affects the likelihood of a ball moving forward while larger α reduces the spread of arrival times.

3.4 Distribution of the time of arrival

So far, to capture the general features of the dynamics, we assumed that the wave could undergo a free unconstrained propagation. However, we are interested in the situation where the array is of finite length, i.e. the case when there exists a final urn that the balls cannot leave. By introducing this special urn we can model the arrival times.

We studied the time of arrival of a chain of balls when starting simultaneously, for different values of N, n_u, α and path length U . The results obtained by Markov Chain modeling were validated by identical outcomes from Monte Carlo simulations. Each of the plots in Figure 3 show the probability distributions of the time of completion for different values of the parameters n_u and α , for a chain of $N = 5$ balls starting simultaneously at time $t = 0$ from urn $u = 1$: the time of completion here is defined by latest time at which a ball reaches the final urn, $u = 10$. As α increases, one can observe the progressive “leaning back” of the completion-time probability distributions (there is a regression to the mean), and the progressive reduction of their variance: the higher n_u the more apparent is the effect. Thus, for higher n_u the increase of the non-linearity, in effect represented by α , is beneficial to the time of com-

pletion.

This is due to the combined effect of two previously mentioned mechanisms:

1. the higher the value of the parameter n_u , the faster the balls evacuate their current urn when its cardinality falls below a given threshold; and
2. for a given fixed n_u , the higher α , the longer the lifetime of short-length chains (i.e. the longer it takes for the chain to stretch out). In other words, if, early in its journey, the start of the chain reaches the arrival position, and the chain is still short, then it will take a short time for the whole chain to reach the destination.

4 The effect of memory on adaptive behaviour

4.1 Adaptive behaviour

Up to this point we considered each round (from departure to arrival of N balls) in isolation with all the balls departing at the same time and did not consider any cognitive effects. In this section we introduce adaptive mechanisms in the individual choice of the departure time. Because we are no longer focusing on fundamental combinatorics but on our findings' implications for traffic congestion patterns, we also switch vocabulary to refer more explicitly to commuters and vehicles. When there is congestion, it seems reasonable to expect that rational agents will attempt to correct for it so as to reach their destination on time (or as close as possible) on subsequent journeys. Assessing the likely delay caused by congestion may however be difficult, not least because one agent's decision about when to start their journey will affect the congestion experienced by other travellers.

One possible strategy is to rely on one's own experience of previous journeys, i.e. on memory. For instance, in reference to the simple combinatorics example described in section §3.1, the first ball might have needed 3 steps to reach the final urn. A rational decision would be to allow 3 time-steps instead of two for the next journey. This is something that many travellers probably do routinely (if perhaps subconsciously) when they are planning their daily commute. In our model, agents only remember the duration of their journey (thus not the time when they departed and the resulting arrival time).

If memory is infinite, i.e. if agents remember perfectly how many steps it took them to reach the final urn on each attempt and weight them equally, then all will eventually converge towards the average congestion delay. If they all add this average delay to the minimal duration of the journey ($U - 1$), they will all arrive more or less on time, sometimes early, sometimes late depending on random chance. The downside is that congestion, which can be measured, e.g. as the total number of failed attempts to move forward across the entire population, remains roughly unchanged. It may be worth emphasising that this is a stable-state, as long as the parameter values (number of commuters, length of the journey, non-linearity, etc.) do not change.

On their first attempt (without any memory to rely on), agents start their journey at time zero ($t = 0$), resulting in a certain level of congestion. Once this level has been correctly assessed by all the agents through averaging over many journeys, all of them will still depart simultaneously (or, more accurately, within one time-step of each other since the system is discrete and the average congestion delay might not be). Instead of this time being at $t = 0$ (as was the case on the first attempt and would continue to be the case without memory), it will be $t = -\delta$ (the average delay), but congestion effects will be statistically the same.

Fig. 4 Evolution of the average delay caused by congestion with infinite memory. Parameter values are: $N = 32$, $U = 32$, $n_u = 8$, $\alpha = 3$. The plotted variable is an “average of averages” over the whole population and for 1000 independent realisations.

Figure 4 illustrates this by showing how average delay first decreases as the spread of departure times peaks (leading to lower congestion), a result of each agent correcting for the individual delay it experienced in the first round. It then progressively creeps back towards its initial value (when all commuters departed at $t = 0$) as the agents converge towards the same average estimate and the spread narrows again around an earlier departure time ($t = -\delta$). The oscillatory signature is the result of discretisation effects.

The apparent conclusion is that a shorter memory span is beneficial, since congestion is substantially reduced when agents only have recent information at their disposal. As illustrated in Figure 5, this is indeed what was found when running the corresponding simulations. The dampened oscillation signature particularly visible for the shortest memory span is explained by the fact that, due to leaving at the same time-step in the first round (worst case), agents first tend to “overestimate” congestion. The resulting “overcompensation” leads to lower-than-average delays in the second round, which in turn fosters more congestion on round three etc. until the system stabilises on the equilibrium distribution of departure times (steady-state).

Fig. 5 Evolution of the average delay caused by congestion with memory span 1, 2 and 4 (rounds). Parameter values are: $N = 32$, $U = 32$, $n_u = 8$, $\alpha = 3$. The plotted variable is an “average of averages” over the whole population and for 1000 independent realisations.

4.2 Cost of reducing congestion

Of course, the reduced congestion / delay for short memory spans comes at the expense of increased uncertainty about actual arrival time. However, for the arbitrarily chosen parameter values considered here, this effect appears to be relatively weak. When comparing the two extremes (single-round memory versus infinite memory), we found an increase of less than 10% in the average “error” (absolute difference between target and actual arrival time) for more than a factor of two in congestion reduction (equivalent to average delay, see Figures 4 and 5). Furthermore, the frequency distributions of the difference

Fig. 6 Frequency distribution of the difference between target and actual arrival times for myopic and non-myopic agents for $N = 32$ (A) and $N = 64$ (B), recorded at the last (256th) round of 1000 independent realisations, i.e. 32,000 and 64,000 data-points for (A) and (B) respectively. Negative values correspond to early arrivals. Other parameter values are identical: $U = 32$, $n_u = 8$, $\alpha = 3$.

between target and actual arrival time are not dissimilar, the most noticeable

difference being an increase in early arrivals for the shorter memory span (see Figure 6a). This result appears quite robust, at least with respect to population size (Figure 6b). However, when plotting the absolute difference between target and actual arrival time against departure time, the difference between infinite and short memory spans becomes more apparent (see Figure 7).

When memory is limited to one round, early departures ($t = -6$ or -7), prompted by a previous late arrival, can be observed. This typically leads to a very early arrival, an effect attributable to the rarity of this strategy and the correspondingly low congestion levels experienced by the agent choosing it, who has “the road to itself” for the entire journey. At the opposite end of the spectrum, late departures ($t = 0$ to -2) are characterised by high variability. In other words: a short memory span favours a stable spread of departures, which in turn leads to lower congestion levels, but the outliers of the distribution are penalised. Simplifying this we can say that “early birds” arrive long before their target time, whilst those who leave it to the “last minute” are not only likely to be late, but also experience increased uncertainty.

Fig. 7 Correlation between departure time and the absolute difference between target and actual arrival times, for the last (256th) round of 1000 independent realisations. Histograms are the average value, error bars the standard deviation. $N = 32$, $U = 32$, $n_u = 8$, $\alpha = 3$.

4.3 Individual experience

A further aspect of the self-organisation process whereby the system moves to a steady-state characterised by stable frequency distributions of departure times and delays (whatever the memory span) is the experience of individual agents. Indeed, “macroscopic” statistical stability does not necessarily translate into a similarly repeating pattern on a “microscopic” scale. For instance, for long memory spans, a narrow frequency distribution of departure times indicates that individual agents tend to leave roughly together and at the same time round after round (at or close to steady-state), but this does not hold for the “single-round memory” scenario. As per the rules described previously,

an agent who leaves at, e.g. $t = -7$ and, having experienced no congestion, arrives 7 time-steps early (leftmost bar of Figure 7), will be scheduling its next departure at $t = 0$, likely following an “oscillatory” pattern, alternating between early and late departures.

Considering this, the “single-round memory” scenario is arguably not very plausible, at least not if the model is meant to emulate the decision processes of truly “intelligent” agents, such as commuters trying to determine when to leave home to arrive on time. The “infinite memory” hypothesis, as an approximation of a learning process through which the agent progressively estimates an average delay, seems more accurate, not the least because it does not lead to high amplitude oscillations (which are intuitively unrealistic). The fact that it also fosters high congestion levels appears to add to its credibility.

If we accept the “infinite memory” algorithm as realistic, then an interesting question is: can the system be steered so as to reap some of the benefits of a short memory span in terms of congestion reduction, knowing that this would require a change in observable behaviour even though decision rules and the value of the parameters governing congestion are both fixed? This effectively means identifying another distribution of departure times that remains stable, not because agents alternate between extremes (a by-product of the unrealistic “single-round” memory), but because different departure times lead to correspondingly different journey durations. In other words: commuters leaving X minutes before the target arrival time must on average take X minutes to reach it, just as in the “natural” scenario, but for different values of X .

5 Performance Evaluation

5.1 Searching for the optimal distribution of departures

Since there doesn’t seem to be a principled way to look for an optimal distribution of departure times, we used a genetic algorithm (GA) to identify it. As in any GA-based approach, two key steps are: (1) to find a suitable representation of system configuration to use as “genome” and (2) to formulate an adequate fitness function.

For representation, we simply used a sequence of 8 integers (“genes”), each one corresponding to the number of agents departing within the associated time-slot (indexed from 0 to -7, i.e. up to seven steps prior to what would be the correct departure time in the absence of congestion). At initialisation, a uniform distribution is selected (i.e. for $N = 32$, each integer is set to $\frac{32}{8} = 4$). In an attempt to simultaneously minimise congestion / delay and the inconvenience resulting from early / late arrival, the fitness of an “individual” (gene sequence) is set to be inversely proportional to the sum of these two values (absolute in the case of the latter), averaged over all $N = 32$ participating agents and 128 independent realisations. Note that this measure of fitness is essentially equivalent to the cost calculation method used in (Xiao and Lo, 2016), which was itself taken from (Small, 1982), only giving the same weight

to the three relevant variables (travel time, penalty for early and late arrival). The fittest 25% of the current population are carried forward to the next generation and are copied at random with mutation in order to generate the remaining 75% of this new generation. Mutation takes the form of one unit / agent being moved at random from one departure time-slot to another.¹

Examining the fittest individual in the 64 final generations of 64 independent runs of 1024 generations each (parameters: $N = 32$, $U = 32$, $n_u = 8$, $\alpha = 3$), the most frequently observed values for each gene, from $t = 0$ to $t = -7$, are 11, 4, 4, 5, 4, 4, 0 and 0. This result suggests that, for the chosen parameter values and performance criterion, no agent should leave earlier than $t = -5$, and that a “peak” at $t = 0$ combined with an “even spread” of departures over the other time-slots is preferable. This seemingly optimal distribution is shown in Figure 8. Imposing this distribution on the agents, one finds an average delay of 1.68 time-steps and an average difference between target and actual arrival time of 1.85 time-steps. The corresponding values were 3.49 and 1.88 respectively in the “infinite memory” scenario, 1.56 and 2.04 in the “single-round memory” scenario. In other words: the departure schedule corresponding to the distribution shown on Fig. 6 does achieve something of a “best of both worlds” performance, with a minimal average penalty (difference between target and actual arrival times) similar to that observed for the “infinite memory” scenario, but much lower congestion (only marginally longer delays than in the “single-round memory” scenario).

5.2 Stability considerations

Unfortunately, as it turns out, the optimal distribution, shown in Figure 8, is also unstable. Even when initialising the system with the corresponding values, which would be equivalent to assigning a departure time to all agents prior to their first attempt, when they have no individual experience to guide their choice, as soon as (infinite) memory starts governing their behaviour they quickly revert to the schedule produced by natural system dynamics.

In other words: the answer to the question formulated in the closing paragraph of §4.3 on page 18 (asking whether the system be coaxed to exploit the benefits of short term memory) is “no”.

In summary: although it is possible to identify an optimal schedule, steering a system comprised of rational “selfish” agents with infinite memory towards it does not appear to be feasible, at least not through some “one-off” external intervention designed to fix initial conditions.

¹ It may be worth emphasising that minimising only one of the two performance indicators, either congestion *or* absolute difference between target and actual arrival time, is best achieved by either spreading departures as much as possible (a strategy approximated in the “single-round memory” scenario) or, on the contrary, by concentrating them around the expected average delay (as observed in the “infinite memory” scenario), respectively.

Fig. 8 “Optimal” distribution of departures identified by genetic algorithm. Parameter values remain unchanged $N = 32$, $U = 32$, $n_u = 8$, $\alpha = 3$.

6 Conclusion and Future Work

In this paper, we have presented a simple probabilistic model designed to study the influence of cognitive effects (learning) over recurrent congestion dynamics (e.g. “rush hour”). Although this was not our main objective, we hope to have also made a small contribution to the field of traffic flow modelling by showing that congestion phenomena can be conceived as a combinatorial problem and tackled as such in simulation or using formal methods such as Markov Chains. However, we also wish to acknowledge that we chose this approach primarily because it was simpler to use and yet fit for our purpose, which didn’t require the level of accuracy afforded by more detailed and established traffic models.

Our main finding is that it is very plausible that cognitive effects actually contribute to perpetuating congestion: as agents (drivers) progressively learn and adapt to the average delay, they simply “displace” peak traffic back in time so that it doesn’t cause them to arrive late. Unfortunately, this ensures that the other costs of congestion (pollution, waste of time and energy, stress ...) remain the same, which should be of concern to relevant authorities.

By contrast, “forgetful” or “myopic” agents, i.e. those that only remember the outcome of their last journey, experience a reduction in congestion because random daily fluctuations result in them spreading their departures over a period of time, therefore limiting competition for the contended resource.

It does seem important to emphasise again that our work is only indirectly related to the minority game. Although both models deal with rational agents and the impact of overcrowding (or congestion), said impact is explicitly and negatively affecting the agents’ performance in the minority game, leading them to actively try to avoid overpopulated slots (with complex dynamics resulting from their concurrent decisions). By way of contrast, in our model, congestion is “neutral” in the sense that increasing the duration of a commute does not necessarily result in the agent arriving late if it had chosen to leave early, but precisely in the opposite (i.e. timely arrival). In other words:

since the duration of the journey is not explicitly factored into the objective function determining the performance of the rational agents, avoiding crowded departure slots simply plays no direct part in their decision. This is why the steady-state in the infinite memory scenario is a congested state, which could not be the case in the equivalent minority game. In fact, our main contribution to the field might be to have identified a possible explanation as to how cognitive effects can contribute to the perpetuation of congestion, despite the obvious inconvenience of extended journey times.

Our results clearly indicate a trade-off between reducing congestion and increasing uncertainty about journey time on the one hand, and accepting the former in order to reduce the latter on the other. It is not entirely clear where human drivers sit between the two extremes (“myopic” vs. infinite memory), but common sense suggests that we are closer to the second option. Humans possess the ability to remember several journeys and introspection suggests that we use this knowledge to estimate an average delay. The question of whether we can significantly reduce the adverse effects of memory (recurrent congestion) for an acceptable increase in journey time variability seems therefore both legitimate and largely open. Our findings suggest that this is feasible in theory, by scheduling departures according to an optimal distribution, the shape of which is determined by the relative weight of both costs (congestion and delay) and traffic dynamics. However, even if we had the means to identify and enforce such an optimal distribution of departures at a given point in time, we would still be facing the problem that memory cannot be “turned off” and that this distribution is unstable, i.e. quickly destroyed by “natural” learning dynamics in the infinite memory approximation. Accordingly, we will focus our planned future work along two main axes.

The first axis will consist in developing methods to influence departure times to generate the desired (optimal) distribution profile over time. In principle, it seems feasible to use information and communication technology to (1) identify commuters frequently sharing the same “space-time” coordinates and therefore adversely impacting each other’s journey times, and (2) use this data to issue a personalised advisory designed to de-synchronise their trips. One relevant aspect of human trip preferences that has not so far been considered here is that travellers might differ in terms of how they treat the cost of being late vs. being early. If a population of travellers differ in this respect then it may be possible, through personalised advice about departure time, to encourage a stable and efficient set of departure times.

The second axis will consist in acquiring a better understanding of the interplay between the kind of advisory scheme described above and the learning dynamics that appear to govern normal system behaviour. This particular work will have a theoretical component, which will revolve around quantifying the “drift” away from the optimal distribution and how frequently it needs resetting to retain beneficial properties. Depending on our ability to test our proposed approach experimentally, it could also involve a behavioural psychology component, in which the response of drivers to advisory recommendations (and to their effect on journey duration and arrival time) is measured. Evi-

dently, such an experiment would need to be conducted on a sufficient scale for the test population to have a significant impact on congestion.

We suggest that this could be achieved by testing the proposed advisory scheme on a large industrial site, to and from which enough people commute to represent a major component of the local traffic.

References

- W. Brian Arthur. Inductive reasoning and bounded rationality. *The American economic review*, 84(2):406–411, 1994.
- Nicola Bellomo and Christian Dogbe. On the modeling of traffic and crowds: A survey of models, speculations, and perspectives. *SIAM Review*, 53(3):409–463, 2011.
- Satyendra Nath Bose. Plancks Gesetz und Lichtquantenhypothese. *Zeitschrift für Physik (in German)*, (26):178—181, 1920.
- Andrea Cavagna. Irrelevance of memory in the minority game. *Physical Review E*, 59(4):R3783, 1999.
- Damien Challet. Coolen, acc: The mathematical theory of minority games. statistical mechanics of interacting agents, 2006.
- Damien Challet and Matteo Marsili. Phase transition and symmetry breaking in the minority game. *Physical Review E*, 60(6):R6271, 1999.
- Damien Challet and Y-C Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica A: Statistical Mechanics and its Applications*, 246(3-4):407–418, 1997.
- Damien Challet and Yi-Cheng Zhang. On the minority game: Analytical and numerical studies. *Physica A: Statistical Mechanics and its applications*, 256(3-4):514–532, 1998.
- Damien Challet, Matteo Marsili, and Riccardo Zecchina. Statistical mechanics of systems with heterogeneous agents: Minority games. *Physical Review Letters*, 84(8):1824, 2000.
- Damien Challet, Matteo Marsili, Yi-Cheng Zhang, et al. Minority games: interacting agents in financial markets. *OUP Catalogue*, 2013.
- D. Chowdhury, L. Santen, and A. Schadschneider. *Statistical Physics of Vehicular Traffic and Some Related Systems*. Physics reports. Elsevier, 2000.
- H. S. E. Chuo, M. K. Tan, B. L. Chua, R. K. Y. Chin, and K. T. K. Teo. Computation of cell transmission model for congestion and recovery traffic flow. In *2016 IEEE International Conference on Consumer Electronics-Asia (ICCE-Asia)*, pages 1–4, Oct 2016. doi: 10.1109/ICCE-Asia.2016.7804769.
- C. Daganzo and UC Berkeley. *The Cell Transmission Model: Network Traffic*. Institute of Transportation Studies, UC Berkeley, 1994.
- Carlos F. Daganzo. The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory. *Transportation Research Part B: Methodological*, 28(4):269–287, 1994. URL <http://EconPapers.repec.org/RePEc:eee:transb:v:28:y:1994:i:4:p:269-287>.

- Ernesto Damiani, Ottavio D'Antona, Vincenzo Marra, and Fabrizio Palombi. *From Combinatorics to Philosophy: The Legacy of G.-C. Rota*. Springer Publishing Company, Incorporated, 1st edition, 2009. ISBN 0387887520, 9780387887524.
- M. d'Aquin, J. Davies, and E. Motta. Smart cities' data: Challenges and opportunities for semantic technologies. *IEEE Internet Computing*, 19(6): 66–70, Nov 2015. ISSN 1089-7801. doi: 10.1109/MIC.2015.130.
- John Davies and Mike Fisher. Internet of things - why now? *Institute of Telecommunications Professionals*, 7(3), Sept. 2015.
- Andreas Deutsch, Guy Theraulaz, and Tamas Vicsek. Collective motion in biological systems. *Interface Focus*, 2(6):689–692, 2012. ISSN 2042-8898. doi: 10.1098/rsfs.2012.0048. URL <http://rsfs.royalsocietypublishing.org/content/2/6/689>.
- Albert Einstein. Quantentheorie des einatomigen idealen gases. *Sitzungsber. Kgl. Preuss. Akad. Wiss.(in German)*, (261), 1924.
- Enrico Fermi. Sulla quantizzazione del gas perfetto monoatomico. *Rendiconti Lincei (In Italian)*, (3):145–149, 1926.
- D. Helbing. *Traffic and Related Self-driven Many-particle Systems*. American Physical Society, 2001.
- Neil F Johnson, Michael Hart, and Pak Ming Hui. Crowd effects and volatility in markets with competing agents. *Physica A: Statistical Mechanics and its Applications*, 269(1):1–8, 1999.
- Norman Lloyd Johnson and Samuel Kotz. Urn models and their application; an approach to modern discrete probability theory. 1977.
- B. S. Kerner and H. Rehborn. Experimental properties of phase transitions in traffic flow. *Phys. Rev. Lett.*, 79:4030–4033, Nov 1997. doi: 10.1103/PhysRevLett.79.4030. URL <https://link.aps.org/doi/10.1103/PhysRevLett.79.4030>.
- B.S. Kerner. *Introduction to Modern Traffic Flow Theory and Control: The Long Road to Three-Phase Traffic Theory*. Springer Berlin Heidelberg, 2009. ISBN 9783642026058.
- Hani S. Mahmassani and Gang-Len Chang. Experiments with departure time choice dynamics of urban commuters. *Transportation Research Part B: Methodological*, 20(4):297 – 320, 1986. ISSN 0191-2615. doi: [http://dx.doi.org/10.1016/0191-2615\(86\)90045-7](http://dx.doi.org/10.1016/0191-2615(86)90045-7). URL <http://www.sciencedirect.com/science/article/pii/0191261586900457>.
- Hosam Mahmoud. *Pólya urn models*. Chapman and Hall/CRC, 2008.
- Matteo Marsili, Damien Challet, and Riccardo Zecchina. Exact solution of a modified el farol's bar problem: Efficiency and the role of market impact. *Physica A: Statistical Mechanics and its Applications*, 280(3-4):522–553, 2000.
- James Clerk Maxwell. *On Boltzmann's theorem on the average distribution of energy in a system of material points*. Printed at the University Press; and sold by Deighton, Bell and Company and Macmillan and Company Cambridge; Bell and Sons, London, 1879.

- James Clerk Maxwell. Illustrations of the dynamic theory of gases (1867). In J. C. Maxwell, editor, *The Scientific Papers of James Clerk Maxwell*, volume 1. 1890.
- Elizabeth Quill. Unclogging urban arteries. *Science*, 319(5864):750–751, 2008. ISSN 0036-8075. doi: 10.1126/science.319.5864.750b. URL <http://science.sciencemag.org/content/319/5864/750.2>.
- Robert Savit, Radu Manuca, and Rick Riolo. Adaptive competition, market efficiency, and phase transitions. *Physical Review Letters*, 82(10):2203, 1999.
- Andreas Schadschneider, Wolfram Klingsch, Hubert Klüpfel, Tobias Kretz, Christian Rogsch, and Armin Seyfried. *Evacuation Dynamics: Empirical Results, Modeling and Applications*, pages 3142–3176. Springer New York, New York, NY, 2009. ISBN 978-0-387-30440-3.
- M. Schreckenberg and R. Selten. *Human Behaviour and Traffic Networks*. SpringerLink : Bücher. Springer Berlin Heidelberg, 2013. ISBN 9783662078099.
- D. H. Sharp. An overview of Rayleigh-Taylor instability. *Physica D: Nonlinear Phenomena*, 12(1):3–18, 1984. ISSN 0167-2789. doi: [http://dx.doi.org/10.1016/0167-2789\(84\)90510-4](http://dx.doi.org/10.1016/0167-2789(84)90510-4). URL <http://www.sciencedirect.com/science/article/pii/0167278984905104>.
- Kenneth A. Small. The scheduling of consumer activities: Work trips. *The American Economic Review*, 72(3):467–479, 1982. ISSN 00028282. URL <http://www.jstor.org/stable/1831545>.
- M. Treiber and A. Kesting. *Traffic Flow Dynamics: Data, Models and Simulation*. Springer Berlin Heidelberg, 2012. ISBN 9783642324604.
- William S. Vickrey. Congestion theory and transport investment. *The American Economic Review*, 59(2):251–260, 1969. ISSN 00028282. URL <http://www.jstor.org/stable/1823678>.
- Tamás Vicsek and Anna Zafeiris. Collective motion. *Physics Reports*, 517(3–4):71 – 140, 2012. ISSN 0370-1573. doi: <http://dx.doi.org/10.1016/j.physrep.2012.03.004>. URL <http://www.sciencedirect.com/science/article/pii/S0370157312000968>. Collective motion.
- J. Wahle, L. Neubert, and M. Schreckenberg. Modeling and simulation of traffic flow. *Computer Physics Communications*, 121:402 – 405, 1999. ISSN 0010-4655. doi: [http://dx.doi.org/10.1016/S0010-4655\(99\)00367-7](http://dx.doi.org/10.1016/S0010-4655(99)00367-7). URL <http://www.sciencedirect.com/science/article/pii/S0010465599003677>. Proceedings of the Europhysics Conference on Computational Physics CCP 1998.
- Yu Xiao and Hong K. Lo. Day-to-day departure time modeling under social network influence. *Transportation Research Part B: Methodological*, 92, Part A:54 – 72, 2016. ISSN 0191-2615. doi: <http://doi.org/10.1016/j.trb.2016.05.006>. URL <http://www.sciencedirect.com/science/article/pii/S0191261516302788>. Special issue: Day-to-Day Dynamics in Transportation Networks.
- H.D. Young, R.A. Freedman, and A.L. Ford. *University Physics with Modern Physics*. Pearson, 2015. ISBN 9780133977981.