Enforcement in free-flow systems: a case study

Dr. Heinz Schild
Research Consultancy Fellow with the Institute of Management Studies
at Witan International College in Reading / UK

Summary

Sometimes System Dynamics tools are not suitable for solving dynamic problems although on a high level of abstraction these problems can be modelled as "stock and flow" diagrams. A typical example for such a class of cases is the design of an enforcement system for a free-flow motorway toll system. This report will discuss how to deal with them by starting with a "stock and flow" diagram and then moving on to implement them as agent-based simulation.

The problem

Queues very often build up at the entrance points to systems with obligatory usage fees. This happens for instance at the box offices of theatres, at ticket counters in railway stations and at the gateways of toll motorways. Queues are unpleasant for the users because of the time losses incurred. They are also not desirable for the operators because they require space and produce operating costs for ticket offices. To overcome these problems it makes sense to permit a free flow into and out of such systems and not to collect fees at the entrances or exits. But then the question arises of how to enforce fee payment.

The primary goal of every enforcement schema, here for a truck toll motorway system, should be to ensure that it is rated as a deterrent in public opinion. Collection of the tolls and fines from toll violators comes second and is therefore of lesser interest here. An enforcement schema can be rated as sufficiently deterrent if the number of violators is significantly below 5% of all trucks and does not increase over time. Furthermore, the deterrent effect should be achieved with minimal obstruction to traffic on the toll road and with minimal costs for the installation and operation of the enforcement equipment and the personnel needed.

This paper is based on a study[*] which discusses the general design of enforcement schemas for a free-flow motorway on which trucks have to pay tolls. A motorway consists of seg-

[*] Heinz Schild: Evaluating enforcement schemas for free-flow toll systems, Bonn, September 2004, in http//:www.CaseLab.info

ments, i.e. the stretches between entrances and exits. The study assumes that DSRC (microwave) or GPS/GSM technology is used in combination with so-called onboard units to locate trucks travelling on segments of the motorway. Stationary or mobile enforcement is used to enforce tolls. The study is not concerned with either the technical implementation of the enforcement equipment or the ways in which violators might manipulate the onboard units of their trucks so as not to be registered when they pass through motorway segments. It concentrates more on questions as to which type of enforcement equipment - stationary or mobile – is to be used in what circumstances at what locations and which operational tactics can be applied to produce the best results.

The goals of this paper are threefold:

- It will describe the agent-based computer model which was used in the study to simulate different toll enforcement designs and briefly outline the main conclusions.
- It will explain some problems with validation which occur when an agent-based simulation is implemented and only a rudimentary empirical data base exists.
- It will further discuss what role system dynamics played during model building and why this approach was not sufficient for modelling this motorway toll problem.

The model

To date there exists only one comparatively large free-flow toll motorway system, namely that in Austria which started operations on 1st January 2004; there trucks and buses with a gross weight of 3.5t upwards must pay toll. Smaller free-flow toll systems are operational for instance in Toronto and in Australia. On Italian motorways a quasi free-flow system is in operation which allows entering and leaving motorways without stopping but forces the drivers to pass through a gate at the entrance and exit points. But their scope is limited and the experiences gained too specific for the purposes of this study.

This model developed uses relatively general but standard textbook relationships to describe the behaviour of the truck drivers and applies them to a slightly idealised toll motorway which resembles a stretch of the A2 motorway in south-eastern Austria.

The components of the enforcement schema are

- the amount of the fine which toll violators have to pay if caught,
- stationary enforcement gantries with equipment to detect and identify toll violators, installed on segments of the toll motorway,
- mobile enforcement devices in the form of either portable gantries or manned vehicles patrolling the motorway, both with equipment to detect and identify toll violators.

If a driver decides not to pay tolls at all, he must either switch off the onboard unit (if this is possible) or prevent it from communicating with the toll registering gantries.

Following the concept of Systems Thinking, first a chart (see figure 1) was drawn which shows the main elements and their relationships in an enforcement system in the form of stocks and flows. The pivotal element is the decision of the truck driver whether to pay a toll or not. The thick lines clearly define a feedback loop.

But there are other features, here printed in a magenta font, which require modelling of geometrical relationships, a concept not found in systems dynamics. For instance, journey length depends on a truck's home area.

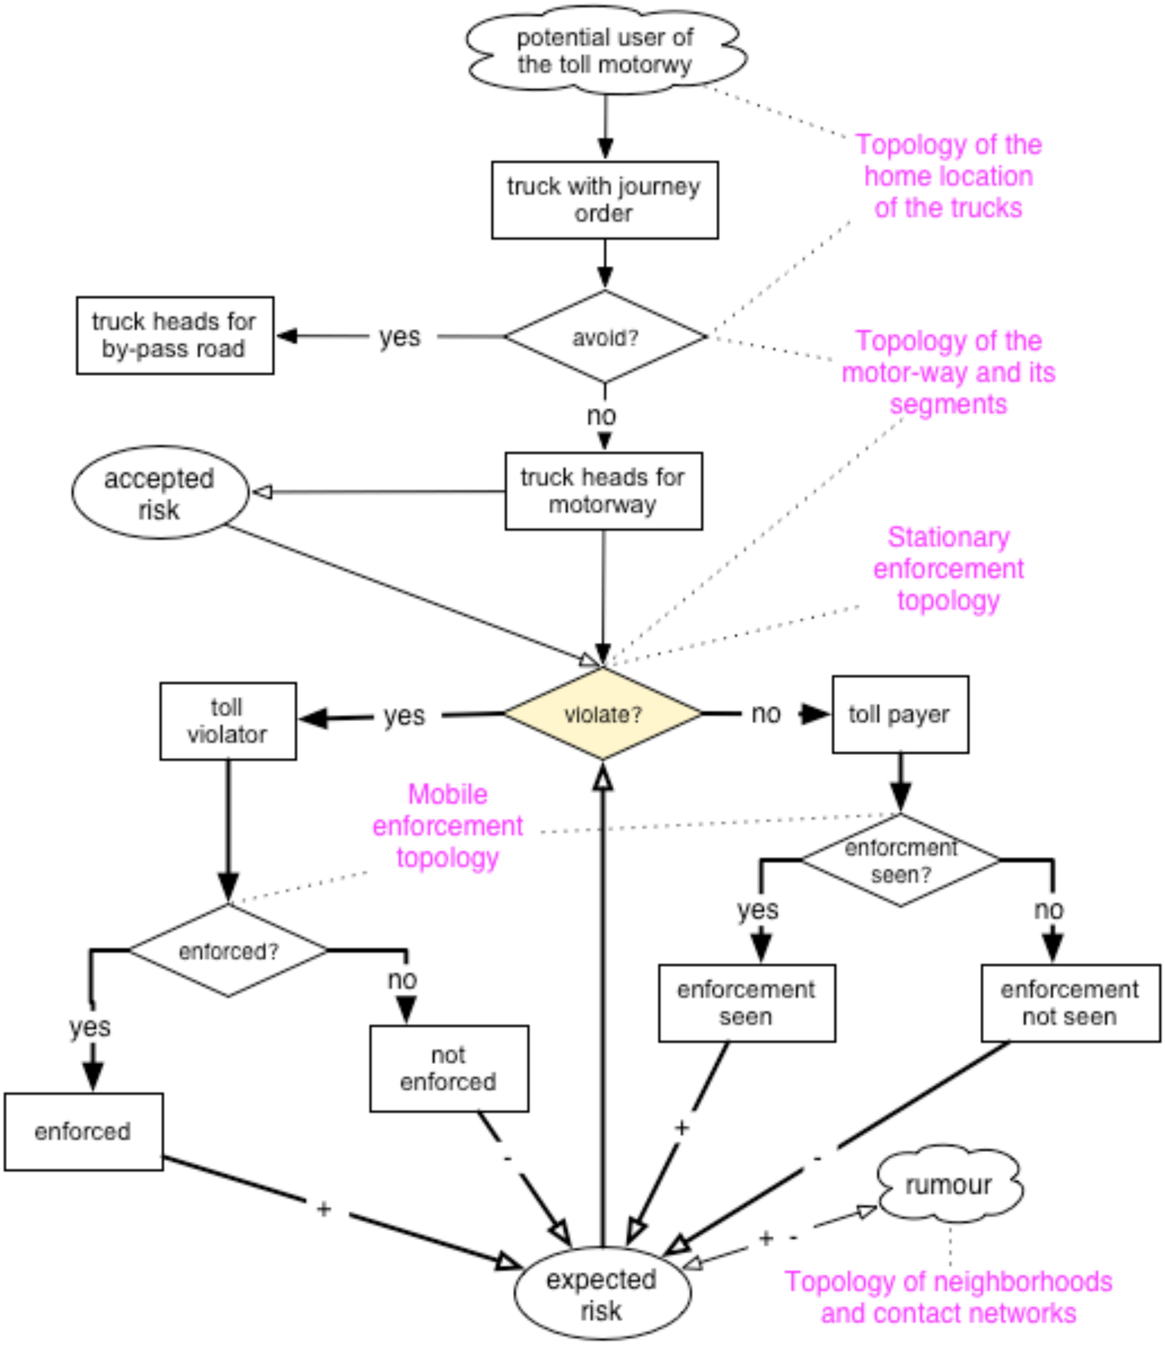

Fig. 1: Relationships of the enforcement model

The decision to avoid the motorway is influenced not only by the amount of toll but also by the length of the journey. The degree of inconvenience incurred by avoiding tolls when using the motorway depends on where a truck starts its journey on the motorway and how many of the stationary enforcement gantries it has to circumvent. The risk of being caught by enforcement depends again on where the journey starts, how long the journey is and where mobile enforcement units are operating unknown to truck drivers. The efficiency of the manner in which rumours about enforcement experiences are spread by a toll violator or a non-violator depends on the number of colleagues a truck driver has in his neighbourhood or further away.

One approach to overcoming these problems at least partly when using a System Dynamics tool would be to split up the users into a set of specific groups depending on their topological characteristics. To do this you would have to differentiate the stock variable in many sub-variables, introduce additional attributes to describe their geometric positions and define sub-procedures which update the geometric positions at every time step. This would substantially complicate the model development process and still only be an auxiliary solution.

The other approach is to use for the implementation an agent-based simulation tool which lends itself in a comprehensible way to modelling topological relationships. A suitable tool for this purpose is NetLogo*, a well documented and relatively easy-to-learn software package. NetLogo was applied for the case study in hand and found suitable. An alternative tool would have been Swarm, a powerful but complicated-to-use collection of C modules now hosted by the University of Michigan.

A simulation model implemented with NetLogo consists of 3 parts:

- Backcloth: a Graphic Window showing the agents, their geometrical relationships and the changes during a simulation run
- Program: LOGO-related code which sets up the backcloth and the behaviour of the agents and drives the simulation of their actions
- Input/output devices: Buttons, sliders and switches to control the simulation and to input values; charts and monitors to show the state of variables and their development during a simulation run.

Mapping the topology

The design of the enforcement system for a toll motorway is heavily influenced by the topology of the relevant stretch of motorway, the distribution of the home locations and the journey patterns of the trucks, which are the users of the motorway we are dealing with here.

The static and the dynamic topological aspects are mapped on to the backcloth of the simulation model. The static elements are the home locations of the trucks, the relevant motorway stretch with its segments and the locations of the stationary enforcement units. The dynamic elements are the positions of the trucks on their journeys and the mobile enforcement units during a simulation run. Figure 2 shows the screenshot of the graphic window at a given moment during a run.

The backcloth of the model is like a ribbon glued together at its left and right borders as well as its top and bottom borders, forming a torus or “doughnut”. An agent moving off on the left border comes back on the right side. Here the backcloth consists of 101 x 65 rectangles or

* Wilensky, U. 1999. NetLogo. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL.

patches. This granulation is sufficient to show the relevant features of the motorway and the movement of the trucks, and will still fit on a standard computer screen.

The backcloth is horizontally divided into two areas. The lower part is reserved for the motorway and the bypass road that avoiders will use to circumvent the motorway. The motorway consists of westbound and eastbound directions, lanes in both directions and segments. The upper area of the backcloth shows the home locations of the trucks in three differently coloured zones. The light-blue zone in the middle represents the core area of the city, flanked on both sides by an industrial belt in a khaki colour. The rest belongs to the rural and sparsely populated area between cities and is shown in green. The depth of the zones is chosen in such a way that the vertical densities of the locations conform to reality; this is needed to handle the spread of rumours properly.

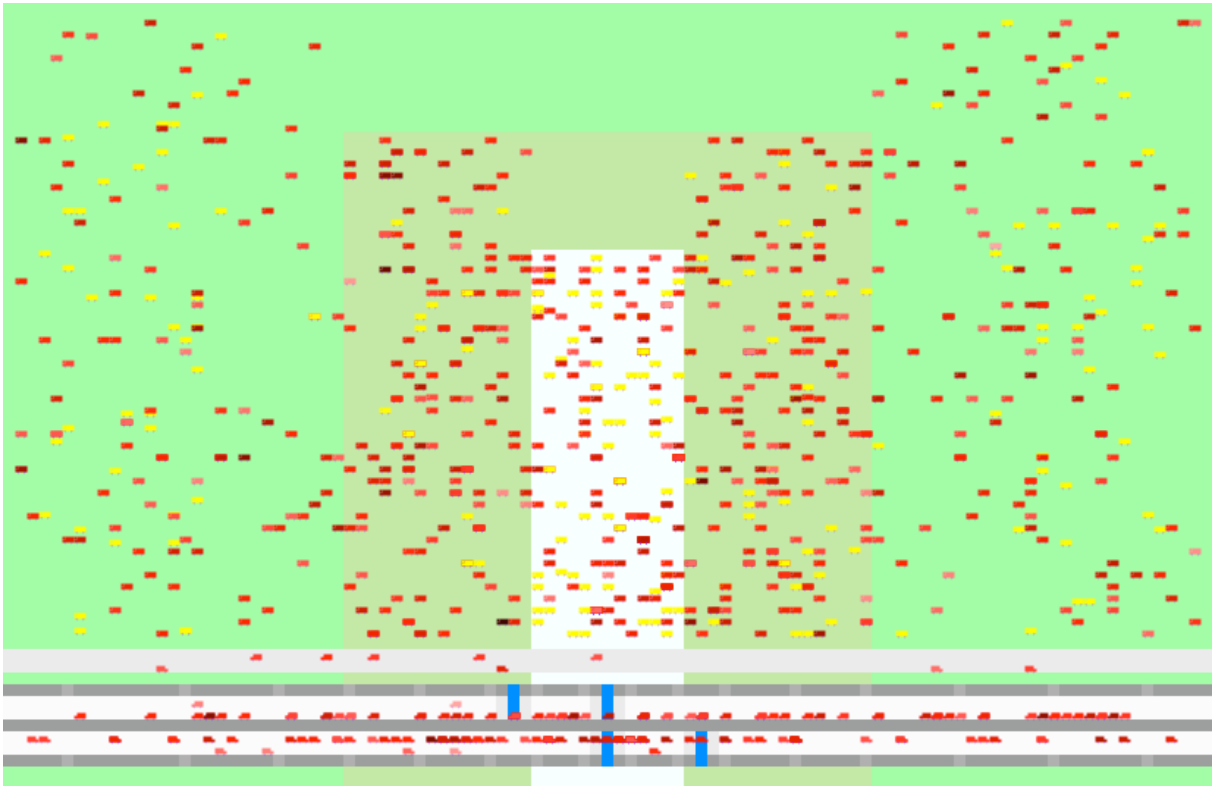

Fig. 2: Screenshot of the backcloth

The dots in the upper part of the diagram represent the trucks. The model includes 1000 trucks in total. Their distribution at the beginning of every simulation is shown in table 1.

Some of the trucks are parts of a fleet, of which there are 50 in the model. The frequency of the trucks over the fleets is approximated by a power function where the biggest fleet has 1 + 20 trucks, the smallest 1 + 1. Empirical data show that this approximation is sufficiently accurate. All trucks of a fleet are located in the same patch.

Zone	Proportion of trucks	Density of trucks
City	15 %	0.35 per patch
Industrial belts	2 x 27.5 %	0.4 per patch
Rural area	2 x 15 %	0.1 per patch

Tab. 1: Distribution of the truck home locations

At the beginning of every simulation run, values for some of the truck attributes are assigned for every truck. The values assigned to the accepted risk are normally randomly distributed with a mean value depending on the level of fines as discussed later. The values assigned to the expected risk are normally randomly distributed with a mean of 0.5. These values will change during the simulation run depending on experiences with enforcement and possible rumours. The value assigned to the number of acquaintances which a driver may have is approximated by a power function, which is in line with published experiences about human communication networks. It is assumed that 100 trucks out of the total may have between 1 and 8 acquaintances.

If a truck is in its home location, its colour represents the risk its driver expects of encountering mobile enforcement on his next journey. Dark-red coloured trucks have drivers with high risk expectations; drivers of light-red trucks expect a low risk. If more than one truck is located at one spot, the colour of the top truck is shown.

In the model a truck moves from its home location to the goal of a journey and after a while from there back home again. During its stay in the goal location the truck is coloured yellow.

The motorway shown in the lower part of the backcloth consists of segments with different lengths. The segment lengths are modelled reasonably realistically and are multiples of patches. A patch in the model is about 1.5 km long in reality. The modelled stretch of the motorway covers 101 patches or about 150 km. A truck travelling at 70–80 km/h will take about 1.2 minutes or 72 seconds to pass a patch.

Patches in blue represent stationary enforcement gantries, those in cyan, mobile enforcement units.

Starts of journeys

During each time step some of the trucks start their journey. The number of starts depends on the time of day. For the purposes of this model it is sufficient to use only one type of day and its profile. This profile is triangle-shaped, with the lowest value at 0:00 hrs and the highest at 9:00 hrs.

The mean number of starts per time step is 5 or 250 per hour. These numbers approximate reasonably accurately the situation on the stretch of motorway mapped in the model. Table 2 shows the values with which the mean number of starts must be multiplied to get the starts per time step.

Hour	Morning shift	Afternoon shift	Night shift
6:00 / 14:00 / 22:00	1.57	1.24	0.29
7:00 / 15:00 / 23:00	1.86	1.12	0.18
8:00 / 16:00 / 24:00	1.94	1.00	0.06
9:00 / 17:00 / 1:00	1.82	0.88	0.14
10:00 / 18:00 / 2:00	1.71	0.76	0.43
11:00 / 19:00 / 3:00	1.59	0.65	0.71
12:00 / 20:00 / 4:00	1.47	0.53	1.00
13:00 / 21:00/ 5:00	1.35	0.41	1.28

Tab. 2: Traffic profile

At every time step an appropriate number of trucks is randomly selected to start a journey. The selection mirrors the density of the locations in the 3 zones. Transit trucks are not explicitly considered. It is assumed that they behave like normal trucks on very long journeys. This is a reasonable assumption as long as the motorway stretch is long enough and the proportion of transit trucks low enough.

For every truck starting a journey, the direction (west- or eastbound) is randomly selected and the journey length calculated. If a truck is at its goal location, it will set off in the direction of its home location. The length of a journey depends on the zone where the truck is located; it is normally distributed and at least 4.5 km or 3 patches long, as shown in table 5.

Start area	Mean	Standard deviation
City	20	Mean * 2/3
Industrial belt	40	Mean / 2
Rural area	80	Mean / 3

Tab. 3: Journey length in km

Avoiders

When tolls are introduced on a motorway which has been in use for some time, normally about 5% to 15% of the drivers or owners of trucks may find the toll charge too high compared with the advantages gained, especially for short journeys, and start avoiding the motorways. The percentage of avoiders used in the case study depends on the level of tolls and fines. Only those whose journey length stays below a threshold that is modelled by a random exponential distribution with a fixed mean value of 45 km will actually avoid driving on the toll motorway.

Avoiders will increase the traffic on first, second and lower grade roads in unexpected ways, producing congestion at bottlenecks and inconvenience to residential areas. The authorities may undertake activities to curb the number of toll avoiders, like introducing speed limits or driving restrictions for transit trucks on circumvention routes. After some time these activities and the effects of habituation will bring back most of the avoiders.

Violators

Those drivers who do not avoid the motorway have to decide if they want to pay or to violate tolls. Following the microeconomic utility concept, a rational driver will consider violating tolls if his expected savings are higher than the expected losses due to enforcement and if he expects that the risk of being enforced is lower than his level of acceptance.

$$Toll * (1 - expectedRisk) \geq (Toll + Fine) * expectedRisk \mid expectedRisk \leq acceptedRisk$$

It is reasonable to assume that the fine is a function of the avoided toll. Here a linear function is used. An alternative approach would be to have a constant value plus an additional amount, depending on the toll that was violated.

In our model we consider 3 levels of fines. The definition of the 3 levels as used in the model and the accompanying mean accepted risks are shown in table 4.

Fine level	meanAcceptedRisk
Low	0.333
Medium	0.25
High	0.15

Tab. 4: Fine levels and mean accepted risk

Below the "low" fine level the deterrent effect of the fine will not be sufficient any longer and "law-abiding" drivers may feel let down. The "high" value for fines must bear a reasonable relationship to the damage done by violating a toll. If it is too high, drivers and related interest groups will start opposing the whole toll system. For instance, a violator might not cooperate in declaring his entry point if caught, contest any penalty which is slightly questionable in court, or complain publicly about minimal toll registration errors or delays in receiving receipts for post-paid tolls. This might cause a breakdown in the working relationship with the toll operator and damage the image of the toll and motorway operator.

The level of the accepted risk depends on the truck driver's personal degree of risk aversion -- a driver may be a risk seeker or a risk avoider-- and the mean accepted risk which is associated with the level of fines in the toll system. Drivers of trucks belonging to a fleet may not consider violating tolls at all. Then their accepted risk would be 0. But this option is not considered in the model.

In the model the accepted risk of a driver is calculated at the beginning of a simulation run using a left-truncated random normal distribution and does not change during a simulation run. In reality the levels of accepted risk are corrected if they are never reached or always violated. This does not affect the results of the simulation significantly and therefore is neglected here.

At the beginning of every journey a truck driver will estimate the risk of being caught if he violates tolls and the inconvenience which he would have to put up with to avoid being caught for sure.

The value of the estimated risk expectation depends on his direct or indirect experience with the enforcement system. He may have violated tolls in the past and been caught or not. Or he may have observed acts of enforcement or not. In addition, the risk expectation of a driver may be influenced by rumours about neighbours' or acquaintances' experiences with toll enforcement. The way in which the adaptation of the expected risk is handled in the model is discussed in the following two paragraphs.

If a driver wants to avoid tolls and has an onboard unit installed, he will definitely have to circumvent all segments on his journey in which stationary enforcement units are operating. Switching on the onboard unit only when passing a segment in which a stationary enforcement unit is operating and switching it off in all other segments will attract attention in the software of the computer system in the back-office and may trigger a visit from the enforcement agents to his home location. To circumvent a segment, the driver has to exit the toll road before reaching that segment and enter it again further on. A journey is considered as inconvenient if the difference between the total number of journey segments and the number of journey segments to be circumvented lies below a certain threshold.

In reality truck drivers may not behave as rationally and simply as assumed here [1]. Often violators are over-confident about not being caught. Parameters do not stay constant over time. But as long as empirical data relevant to the situation discussed here are not available or only sporadically so, it is reasonable to stay with the approach outlined above.

During a journey a truck driver may increase or decrease his risk expectation considerably, depending on his status (toll violator or not) and his encounters (or lack of such) with a mobile enforcement unit.

Rumour

On reaching their goal or home location some violators will tell neighbours and acquaintances about their experiences with enforcement. In the model it is assumed that a listener will adapt his risk expectation, but attach only half as much weight to the rumour as he attached to his own risk expectation before receiving the information from the rumour source.

In the model only drivers who have seen mobile enforcement units in operation during their last journey or violators who were caught by enforcement units spread rumours. Non-violators only very rarely spread rumours, if at all, and therefore are neglected here. Rumour is spread

[1] See Eisenstein, Franz, Martin Weber: Rationales Entscheiden, Berlin 2003, p. 359 ff

to all drivers whose trucks are located in a circle of given radius around their rumour source on the backcloth. It is assumed that a proportion of the drivers who spread rumours have acquaintances further away, to whom they will also recount their experiences. A proportion of these acquaintances will again spread the rumour to their neighbours.

In the model it is assumed that rumour spreading happens instantaneously in the same time step as when the driver who is the rumour source arrives at his goal or home location. This assumption allows using a simplified time regime for the simulation runs without seriously distorting the results.

Movements on the backcloth

When a truck starts a journey on the motorway, its icon on the backcloth is transferred from its home position vertically downwards either on to the bypass road if it wants to avoid the motorway or into one of the motorway directions. On the bypass road a truck moves forward on average one patch in every two time steps. On the motorway it moves forward one patch in every time step until it comes to the end of its journey, from where it is transferred vertically upwards on the backcloth to its goal location. The vertical coordinate is the same as in its home location, but adapted in proportion to the depth of the goal zone, to avoid disturbances in the spreading of rumours. When the truck starts its return journey, it is again transferred vertically downwards on to the motorway. Trucks moving off the left border enter the diagram again at the right border and vice versa, as explained previously.

Trucks on the motorways may move either in the right or left lane. Violators are shown in the right lane, non-violators in the left lane of each direction. Toll violators are hidden when they pass a segment with a stationary enforcement gantry, to make clear that they must circumvent this segment in order not to be caught.

A truck's colour may change during its journey on the motorway. When a toll-violating truck passes a mobile enforcement unit, it will be fined; this causes the driver to increase his risk expectation and the truck colour will turn a darker red.

Toll violators arriving at their goal or home location will circulate their experiences from their last journey, namely an adapted risk expectation, to neighbours. This may change the colour in which their trucks are shown in the diagram and influence their decision whether to violate tolls at the beginning of their next journey.

Stationary enforcement

The enforcement effects of stationary and mobile units are very different. Stationary enforcement units reduce the number of violators depending on their inconvenience threshold. Stationary enforcement units will only catch uninformed or inattentive violators, who are neglected in this study. "Professional" violators will circumvent them as explained earlier.

Stationary enforcement units are mounted on fixed gantries. Their purpose is to identify violators, take images of their number plates and send them to a back-office for further enforcement activities. In their effect they are comparable to fixed radar speed traps and like these have no, or only a marginal, influence on the general risk expectation of being caught. Bearing in mind that stationary enforcement units mainly have an effect of inconvenience, their installation makes sense primarily on motorway segments with heavy truck traffic and where circumvention would be highly inconvenient. This may be the case for some segments of motorways in city areas and their industrial belts or in special topographical situations like main river crossings or mountain passes.

Stationary Enforcement	Installations in the city zone	Installations in the industrial belts	Installations in the rural areas	Deployment density
non	0	0	0	0 % of segments
minimum	1 in each direction	0	0	7 % of segments
standard	1 in each direction	1 in each direction	0	13% of segments

Tab. 5: Stationary enforcement configurations

In the model two alternatives for stationary enforcement are considered. The positions of the stationary enforcement installations for the "standard" alternative can be seen in figure 2, the screen shot of the model backcloth. In the "minimum" alternative there is only one installation in the middle of the city zone. The deployment density of the "standard" alternative is comparable to that found on the Austrian toll motorways. See table 5.

The yearly costs of a stationary enforcement unit as installed in free-flow motorway toll systems today depend on the technology used, the installation site and the total number of gantries controlled by a given back-office team. Excluding administrative overheads, a stationary unit including its share of the back-office costs around € 120.000 – 150.000 a year. These costs are definitely too high to install stationary enforcement systems in more than 15% of all segments.

Mobile enforcement

The main goal of the mobile enforcement units is to increase the drivers' expected risk of being enforced if they violate tolls on their next journey.

Mobile enforcement units are either portable gantries equipped like stationary units or specially equipped enforcement vehicles which mingle with other traffic and are manned by at least two enforcement officers. The officers in the enforcement cars may be employees of the company that operates the toll system or police officers, depending on legal circumstances. With the help of installed equipment the enforcement officers can detect violators, register their number plates and send this information to a back-office for further enforcement activities. Under some circumstances it may be possible for mobile enforcement units to stop violators and enforce toll payment on the spot.

Providing mobile enforcement units do not operate for too long in the same spot, giving violators no time to be warned via CB radio or functionally equivalent devices, their positions are unknown to truck drivers and therefore cannot be circumvented by violators. From the point of view of deterrence, mobile units using manned cars would appear to be more efficient than portable enforcement units. Only these units and not portable ones are considered in the simulation model.

Mobile configuration	Active	Active in city zone	Active in industrial belts	Active in rural areas
non	0 %	0 %	0 %	0 %
fluc33%	60 % 33 % 24 %	20 % 11 % 8 %	20 % 11 % 8 %	20 % 11 % 8 %
fluc67%	86 % 67 % 55 %	29 % 22 % 18 %	29 % 22 % 18 %	29 % 22 % 18 %
fluc75%	86% 75 % 67 %	29 % 25 % 22 %	29 % 25 % 22 %	29 % 25 % 22 %

Tab. 6: Mobile enforcement configurations

Mobile enforcement units influence the risk expectation of violators and non-violators alike. Violators feel the effect when they are caught: they have to pay a fine and/or are penalised by other means. This will increase their risk expectation substantially at least for a while. Seeing mobile enforcement units in action increases the risk expectation of non-violators as well, although not as strongly as in the case of caught violators. Not being caught when violating tolls or not seeing mobile enforcement units in action for some time will reduce the risk expectation to a fraction of its previous value and influence a future decision to violate or not to violate tolls. Again, the effect is normally stronger on violators than non-violators.

Mobile enforcement alternatives are characterised by the proportion of time during which they are active on a motorway stretch and how they distribute their active time across the zones and throughout the shifts. 3 shifts of 8 hours each are considered (see table 2). In table 6 the first number in a cell stands for the first shift, the second for the second and the third for the third shift. The numbers are rounded. The distribution over the shifts reflects approximately the traffic profile in the course of a day. During a simulation run the position of the mobile enforcement units is changed randomly after every 25 time steps (or 30 minutes of real time).

The yearly costs of a portable enforcement unit deployed 24 hours a day every day of the year, again excluding administrative overheads, are about 2 to 3 times greater than those of a stationary unit. A manned mobile enforcement team, including car and equipment, costs around € 200.000 a year, overheads excluded. To operate one mobile enforcement unit 24 hours a day every day of the year requires at least 4 enforcement teams.

Inputs and Outputs

The parameter values that control the model are either part of the NetLogo program code or they can be selected interactively via buttons, choice and switch boxes. Choice boxes are used to input the stationary and mobile enforcement alternatives and alternative values for other variables, such as the type of fine, convenience threshold and rumour radius. A switch box allows the “avoid” option to be switched on and off.

The values of all global parameters and the attributes of every truck and every patch of the backcloth for each time step can be observed. This is very helpful when testing for, and identifying, bugs or implausibilities.

The results of a simulation run with the model are shown in the graphic window, in monitoring boxes and in charts. Agent-based simulation models are sometimes used to demonstrate emerging patterns and visualise their development on the backcloth. This is not an objective of this case study. Here the backcloth shown in the graphic window is used mainly to help set up the static elements and to visualise the agents’ movement and the development of some of their attributes over the course of the simulation so as to test the program and detect implausibilities.

Monitoring boxes show the values of reporter variables for every time step. Some allow comparing the effects of different alternatives; others are mainly used for plausibility checks. The main reporter variable for comparison purposes is the mean daily percentage of violators or the daily rate of violations. Another reporter variable is the mean value of the expected risk of enforcement. Monitors that show the proportion of trucks on a journey, the proportion of trucks waiting in their home locations, the mean journey length of avoiders and the percentage of trucks avoiding motorway journeys support plausibility checking.

Time series in the line chart showing the number of journey starts, number of violators and number of avoiders are mainly used to pick out oscillations or trends and are also needed for plausibility checking.

Simulation

A model, or the computer program which formalises and executes the model, is only useful for our purposes if the results which it calculates for a set of inputs are largely similar to those which we would observe in reality if the activities equivalent to the set of inputs were to be carried out. Due to the fact that the reality mapped in the model, here the toll system on Austrian motorways, has only been in existence for a short time and thus very little hard data about it have been published yet, the standard validation procedures are not applicable. We have to rely on a surrogate procedure.

Validation

The first step in the validation process is to model all aspects of the relevant reality in great detail and thereby to make sure that the elements and the relationships between them are in reality topologically similar to those in the model. It is well known that adding details increases the explanatory value of a simulation model and makes validity checking easier, in contrast to forecasting models where less detail in the model may increase the reliability of the forecast. Here we are more interested in the explanatory power than the accuracy of the forecasts produced by the simulation.

Agent-based simulation is per se details-oriented. The graphic interface of NetLogo facilitates the proper programming of geometrical relationships, always of importance in traffic models, and supports plausibility checking visually. Because every agent is treated as an entity, the behaviour equations can be formulated from an agent's viewpoint, making it easier to invoke one's own experiences as motorist. It is possible to imagine how one would behave as a truck driver on a toll road. In addition, one's own experiences and those of one's acquaintances with radar speed traps can provide inspiration in developing equations and making intelligent guesses about their algebraic form and parameter values.

Studying the description of the model as outlined above should have convinced the reader that all elements and relationships influencing the calculation of the percentage of drivers prepared to commit significant toll violations have been included in the model, and that the behaviour equations have been formulated in a way which does not contradict the text books or is at least not implausible.

In the second step of the validation process the results of simulation runs must be checked for plausibility. To do this, first some over-all results are compared with empirical data or at least reasonable estimates. In our case the proportion of trucks out on a journey in the model is compared with the known number of trucks out on a journey in reality, and this has been found to be approximately equal. The same is true for the mean journey length and other variables reported in the output monitors and output chart of the model, such as the percentage of avoiders.

The model includes parameters which have a significant influence on the main results variable, here the mean percentage of violators, but for which reliable empirical data are not known. To check the plausibility of the values of these parameters as applied in the model, sensitivity tests are performed. This is done by estimating 3 reasonable values for each of these parameters and then simulating their effects on the main output variable. If the deviations from the middle value of the results variable have the right sign and magnitude, the estimated parameter values are rated as plausible in a scale-variant sense. Scale-variant means here that the ranking of the alternative enforcement configuration remains unchanged when linearly transformed estimates are used for the input parameter values.

Time regime

NetLogo is a period-oriented simulation language. Therefore when implementing the model, one must decide how long each time period of the model should be in relationship to the reality mapped by the model and how many periods a simulation run must cover to allow deriving useful conclusions from the simulation results.

The length of a simulation period is about 72 seconds, as explained previously.

A simulation run should cover enough periods or time steps to reach a steady state or an oscillation of constant frequency and amplitude. In the model only a single day type is used. This makes it easier to set up the deployment of the mobile enforcement units. If weekly and seasonal patterns were considered, appropriate deployment schemas for the mobile enforcement units would be needed. This would only increase the length of the program code and the execution time without delivering new insights, because in the model neither the accepted risk nor the expected risk of the agents decreases over time and rumours spread instantaneously, as explained previously. Furthermore, neither data nor hints in the literature are available which clearly demonstrate that time-related decreases do exist. The spreading of rumours obviously takes time, but incorporating this into the model would only have required more time steps before reaching a stable state without triggering any further insights.

In our case a steady state is reached when the mean percentage of violators over a day cycle no longer changes. Tests show that this occurs with a simulation run covering 7 days or 8400 time steps.

Simulation runs

The main purpose of this case study is to answer the question as to how a truck toll enforcement system for a motorway should be configured. To do this the parameter space must be evaluated. Restricting the values of the parameters to 2 or 3 would still require too many simulation runs. Therefore the values of the results variables are only calculated for a sequence of subspaces and not for all cells of the whole parameter space simultaneously.

When interpreting the simulation runs, one must bear in mind what has been said about the calibration of the model. Only the relationships between the values of the results variable in the parameter space are relevant and not the numbers themselves.

avoid	*Mobile Enforcement*	*Stationary Enforcement*	Mean Violators %
"off"	*"non"*	*"non"*	52,68
"off"	*"non"*	*"minimum"*	36,65
"off"	*"fluc33%"*	*"non"*	10,99
"off"	*"fluc33%"*	*"minimum"*	8,67
"off"	*"fluc67%"*	*"non"*	4,33
"off"	***"fluc67%"***	***"minimum"***	**3,21 Base**
"off"	*"fluc75%"*	*"non"*	3,64
"off"	*"fluc75%"*	*"minimum"*	2,88
"on"	*"non"*	*"non"*	48,86
"on"	*"non"*	*"minimum"*	36,52
"on"	*"fluc33%"*	*"non"*	14,06
"on"	*"fluc33%"*	*"minimum"*	8,27
"on"	*"fluc67%"*	*"non"*	5,97
"on"	*"fluc67%"*	*"minimum"*	2,90
"on"	*"fluc75%"*	*"non"*	4,70
"on"	*"fluc75%"*	*"minimum"*	2,67

Tab. 7: Configuration alternatives

Table 7 shows the results of the subspace defined by the parameters *avoid, Mobile Enforcement* and *Stationary Enforcement*. All other parameters have default values.
The most surprising finding is that stationary enforcement units alone do not have a sufficiently deterrent effect to secure a fair operation of the toll system over time. This can only be achieved if mobile enforcement units are deployed. The reason is that “professional” viola-

tors circumvent stationary enforcement installations and are only caught by mobile units whose positions they cannot predict.

The benefits of stationary enforcement are inversely proportional to the intensity of mobile enforcement. The reason for this is that a high intensity of mobile enforcement discourages potential violators and reduces what is to be gained from stationary enforcement. Based on the values in this case study, a good solution would be to combine *"minimum"* stationary enforcement, installed in segments with high traffic or which are very inconvenient to circumvent, with the mobile enforcement alternative *"fluc67%"* marked in the table as base.

To operate the base alternative every day of the year, about 2.5 enforcement teams including those for backups are needed. The yearly costs would be less than € 700.000.

Avoiders reduce the violation rate because the majority of toll violations take place on short- and medium-length journeys, which are also the domain of the avoiders. As mentioned previously, the proportion of avoiders is relatively high when tolls are introduced but declines over the years.

Some other simulation runs for parameter subspaces are performed mainly to check relative sensitivities of the percentage of violators and to spot implausibilities. Their details are not reported in this paper.

Results

The results of the simulation runs did not show significant implausibility; therefore the model would appear to be valid and the following conclusions can be legitimately stated:

- In free-flow truck toll systems on motorways, stationary enforcement units, based on today's technology and costs, are not a sufficient deterrent on their own to secure fair toll operations with low violation rates over time. This can only be achieved if mobile enforcement units are deployed in combination with stationary enforcement installations, or even without them in some circumstances.
- Mobile enforcement units should patrol very visibly on motorways. Their deployment must be unpredictable, but carried out in such a way that it reflects the distribution of truck traffic on the motorway segments and takes into account suspected patterns of higher violation rates on some segments, such as near borders.
- To enforce tolls most efficiently, a monitoring system should be operated which comprehensively and permanently collects data about the activities of the enforcement units on duty, the traffic on the motorway segments and bypass roads, the enforcement cases and their geographical as well as their daily, weekly and seasonal patterns, and the results of opinion polls about the expected and accepted risk of being caught when violating tolls. In addition, the monitoring system should include a fast simulation model which uses the collected data to calculate deployment plans for all mobile enforcement units.
- Further research efforts should be undertaken to obtain and analyse available data concerning enforcement in free-flow toll systems wherever these are in operation. Based on such empirical data, realistic and robust equation systems modelling the behaviour of the agents can be developed using econometric methods. They provide the

basis for the implementation of comprehensive models which can be used in the design of new enforcement systems and the operation of existing ones. The agent-based simulation outlined in this case study could serve as a starting point.

Some final remarks

This report aims to show that for some dynamic problems System Dynamics thinking and its tools are indeed suitable for sketching a meta-model, although not for implementing it in sufficient detail. It is essential when working on such problems to provide detailed modelling of topological relationships. Concepts of, experience with, and data for their behaviour equations are available only for low-level agents and not in aggregated form. Agent-based simulation software packages like NetLogo, which was used in this case study, are appropriate tools for implementing simulation models to support the understanding of and solution to these problems.

It lies in the nature of agent-based simulation that a great many details have to be considered when implementing the backcloth and the behaviour of the agents. Without a System Dynamics meta-model it is easy to get sidetracked and overlook the dynamic feedback loops.

Agent-based simulation models produce a large amount of detailed data. This could be interpreted as pseudo-empirics and might be useful for generating and calibrating high-level behaviour equations, which could then be used to implement the problem as a fast System Dynamics model. To what extent such pseudo-empirical data are a substitute for real-world data and how statistical sampling theory methods would be applicable requires further study.