Measurement of uncertainty costs with dynamic traffic simulations

A. de Palma* F. Marchal†‡

August 1, 2007

Number of words including abstract and references: 5692

Number of tables and figures: 7

TOTAL: 7472

*ENPC and THEMA, Université de Cergy-Pontoise, 33 bd. du Port, F-95011 CERGY-PONTOISE, France. Phone/FAX: +33-1-3425-6337/6233, andre.depalma@eco.u-cergy.fr

†Laboratory of Transportation Economics (LET), CNRS, Ave. Berthelot 14, F-69363 LYON, France. Phone/FAX: +33-47272-79444/6448, fabrice.marchal@let.ish-lyon.cnrs.fr

‡Corresponding author

Abstract

Non-recurrent congestion in transportation networks occurs as a consequence of stochastic factors affecting demand and supply. Intelligent Transportation Systems such as Advanced Traveler Information Systems (ATIS) and Advanced Traffic Management Systems (ATMS) are designed in order to reduce the impacts of non-recurrent congestion by providing information to a fraction of users or by controlling the variability of traffic flows. For these reasons, the design of ATIS and ATMS requires reliable forecast of non-recurrent congestion. This paper proposes a new method to measure the impacts of non-recurrent congestion on travel costs by taking risk aversion into account. The traffic model is based on the dynamic traffic simulations model METROPOLIS. Incidents are generated randomly by reducing the capacity of the network. Users can instantaneously adapt to the unexpected travel conditions or can also change their behavior via a day-to-day adjustment process. Comparisons with incident-free simulations provide a benchmark for potential travel time savings that can be brought in by a state-of-the-art information system. We measure the impact of variable travel conditions by describing the willingness to pay to avoid risky or unreliable journeys. Indeed, for risk averse drivers, any uncertainty corresponds to a utility loss. This utility loss is computed for several levels of network disruption. The main results of the paper is that the utility loss due to uncertainty is of the same order of magnitude as the total travel costs.

1 Introduction

A potential benefit of Advanced Traveler Information Systems (ATIS) is the reduction of unexpected traffic delays. The travel time for a given journey often exhibits variability due to various sources of stochasticity. The travel demand may vary from one day to the next because of the rhythms of human activities (e.g. sport events, week-end and holiday departures) but also because the capacity of the infrastructure undergoes changes that are partially unpredictable (e.g. road repairs, weather conditions and accidents). The term *non-recurrent congestion* is often used to describe various types of unexpected delays. Although it is difficult to find an accurate definition for this concept in the literature, empirical studies such as (11; 12) and more recently (9; 19) point out that non-recurring traffic conditions are responsible for a large share of congestion in metropolitan areas (the figures vary from 15% to 60%). The primary goal of this paper is to measure non-recurrent congestion by a dynamic traffic assignment simulator - METROPOLIS, and check whether the values pointed out by the empirical studies lead to situations where ATIS would be efficient at reducing congestion. The key idea is that drivers are risk averse so that travel time variability induces a psychological cost to travelers that can be given a dollar measure as explained later on in the paper. One could define recurrent congestion as the travel delays caused by the fact that the (stationary) demand for mobility exceeds the available capacity of the transportation system. However, the demand for mobility is itself defined for a given transportation system that in reality undergoes exogenous shocks (e.g. weather) but also endogenous shocks (e.g. accidents) which depend on demand and supply in a complex way. We adopt here the same assumptions as that of (7): in a simulation environment, exogenous and endogenous shocks cannot be distinguished and are generated in a similar way. By doing so, the distinction between recurrent and non-recurrent congestion is easier since travel times can be compared between shock and shock-free simulations. Identifying the location of recurrent and non-recurrent congestion is often ambiguous because they influence each other mutually. Indeed, if some incidents are well-known to occur frequently at about the same location and under given traffic conditions (e.g. delivery trucks blocking a particular urban street), users may anticipate the corresponding delays and select another itinerary from the start, thereby causing congestion elsewhere on a regular basis. Conversely, non-recurrent congestion can also be induced by recurrent congestion. For instance, the probability of an accident on a road might increase with its traffic load.

One could argue that the *frequent* occurrence of an incident does not qualify it as a non-recurrent or unexpected shock. For these reasons, any evaluation done by simulation should include (a) sensitivity analysis based on varying the probability distribution of incidents and (b) measurements of the travel time delays at the metropolitan level as opposed to local measurements restricted to a subset of the network (e.g. corridors or motorways).

Travelers react to *recurrent* congestion by departing at a time and on a route that take into account their own schedule constraints and their willingness to incur congestion delays. This implies often a trade-off between facing long travel times and arriving close to a desired arrival time. Travelers react also to *non-recurrent* delays by altering their pre-trip decisions such as mode, departure time and route but they can also divert en-route from their planned route and alter their destination. Obviously, a broad panel of reactions are possible, depending on their expectations

about the level of service of the transportation system. The building of the users' expectations is a day-to-day dynamic process where the presence of ATIS can play a role. However, even without any ATIS, travelers are likely to learn about the stochastic nature of the failures of the network if external shocks are applied to the system on a random but regular basis. For instance, if drivers tend to double park on an urban arterial, it leads to undue congestion, and despite the stochastic nature of the situation, travelers taking that route are wary of the situation and may decide to change their itinerary. It has been reported in (18) that *"travelers appear to have an intuition about which trips are highly variable and thus where information acquisition would be of greatest value"*. This internalized information about network reliability reduces considerably the potential benefit of an ATIS system on the long run: if drivers start to be used to random incidents occurring at given locations with certain probabilities, they will adapt by changing their travel behavior. Consequently they would need less severely external information systems since they could rely on their own experience. Conversely, tourists or occasional drivers would still benefit from the provision of external information. Our first set of simulations alongside measuring non-recurrent congestion and proving the short-term efficiency of ATIS, will also show that the long-term efficiency of ATIS is indeed small if travel time reduction is the only objective.

However, travel time savings are not the only benefit of ATIS: the same report (18) shows that car drivers are also willing to reduce the uncertainty of their journey. A recent empirical study (16) underlined the importance of the reliability of the predicted travel times. For instance, drivers knowing beforehand that they will arrive late for a meeting can call up their secretary and postpone it. The users' behavior with regard to unreliable travel times depends on their attitude toward risk. Although the reduction of uncertainty has been widely accepted as a potential benefit of ATIS in the literature, risk aversion is still absent from operational transportation models. A formulation for analytical DTA has been proposed by (13). Similarly, measuring the cost of uncertainty itself in a given transportation system subject to stochasticity would give a meaningful benchmark of the maximum benefit that ATIS can achieve at reducing uncertainty. Preliminary work in that direction has been carried out by (17) in the context of mean-variance and without explanatory variables (see Section 2). Empirical works by (2) have shown that the Value Of Reliability (VOR) can be as important as the Value Of Time (VOT). Our second set of simulations propose to measure the cost of uncertainty in a framework compatible with micro-economic theory and, more specifically, with expected utility theory. For that purpose, the concept of risk aversion is introduced in the simulations to model the impact of travel time reliability on generalized user cost. We show that, even for reasonable levels of severity of the incidents, the cost of uncertainty can be of the same order of magnitude as that of the travel cost.

The remainder of this paper is organized as follows: first, risk aversion is introduced along with the methodology to measure the utility loss due to uncertainty; second, the simulation framework is explained, in particular the day-to-day dynamics and how incidents are generated; lastly, numerical results are provided on realistic examples.

2 Measuring the cost of uncertainty

The most widely used model to capture risk behavior in Economics is that of expected utility theory (14). It relies on two separate assumptions. Firstly, it assumes that preferences can be described by a utility function that is known to the modeler. Secondly, it assumes that the attitude toward risk can be rationalized by an expected utility function. This latter function depends on a parameter θ referred to as the *risk aversion*. In this paper, travelers are assumed to be mostly risk averse (i.e. to have a positive degree of risk aversion $\theta > 0$). Therefore, the variability of the driving conditions corresponds to a loss of utility. Different expected utility function have been proposed in the literature. The Constant Absolute Risk Aversion (CARA) utility function is used here. The utility of a journey started at time t with travel time τ is given by:

$$U_\theta\left(\tau\left(t\right)\right) = \frac{1 - e^{\theta\tau(t)}}{\theta}, \tag{1}$$

where θ denotes the coefficient of absolute risk aversion which is constant in the case of the CARA utility. The limit case $\theta \to 0$ corresponds to the risk neutral individual with $U\left(\tau\left(t\right)\right) = -\tau\left(t\right)$. We only consider the day-to-day variability of τ for a given O-D pair and for a given departure time t. For the sake of simplicity we drop the time-dependent index t as well as $O-D$ dependencies . The variability is described by a probability distribution $f\left(\tau\right)$. Travel times are assumed to be stochastic during a time period T that spans several days. The support of $f\left(\tau\right)$ over T is denoted by Ω_T. According to the theory developed by (20), the expected utility of an individual who departs from O to D at time t, in the case of the CARA utility function, is given by:

$$E\left\{U_\theta\right\} = \frac{1}{\theta} - \frac{1}{\theta}\int_{\Omega_T} e^{\theta\tau} f\left(\tau\right) d\tau$$

while the utility of the expected travel time is given by:

$$U_\theta\left(\left\langle\tau\right\rangle\right) = \frac{1 - e^{\theta\left\langle\tau\right\rangle}}{\theta},$$

where $\left\langle\tau\right\rangle$ is the expected travel time:

$$\left\langle\tau\right\rangle = E\left\{\tau\right\} = \int_{\Omega_T} \tau f\left(\tau\right) d\tau\,.$$

For instance, if we have a lottery two equally likely outcomes τ_1 and τ_2 , we have $E\left\{U_\theta\right\} = \frac{1}{\theta} - \frac{1}{2\theta}\left[e^{\theta\tau_1} + e^{\theta\tau_2}\right]$ while the utility of the expected travel time is $\frac{1}{\theta}\left(1 - e^{\frac{\theta}{2}(\tau_1+\tau_2)}\right)$. The latter utility is not stochastic and certain to happen (if the user prefers that situation to the lottery), and hence we call it the risk-free choice in our context. For risk averse users $\theta > 0$ and $U_\theta\left(\left\langle\tau\right\rangle\right) > E\left\{U_\theta\right\}$ by Jensen's inequality. Therefore, the risk-free solution is always more attractive. The travel time compensation χ is introduced to measure the discrepancy between the stochastic and the

risk-free situations (the concept is reminiscent of the risk premium as introduced in the finance literature). It is defined as the amount of additional time that a user would be willing to spend in order to reduce the utility loss caused by uncertainty. That is, the travelers are indifferent between a stochastic system described by $f(\tau)$ (with an average travel time $\langle\tau\rangle$) or a risk-free system with an average travel time $\chi + \langle\tau\rangle$. Therefore,

$$U_\theta\left(\langle\tau\rangle + \chi\right) = E\left\{U_\theta\right\} \ ,$$

and

$$e^{\theta(\langle\tau\rangle+\chi)} = \int_{\Omega_T} e^{\theta\tau} f(\tau)\, d\tau \ ,$$

and

$$\chi = \frac{1}{\theta} \ln\left\{\int_{\Omega_T} e^{\theta\tau} f(\tau)\, d\tau\right\} - \langle\tau\rangle \ . \tag{2}$$

So far we have considered that the utility was only dependent on the travel time $\tau(t)$. In practice it is often desirable to use a travel cost specification $C(t)$ that depends on other components: boarding costs, access and egress costs, schedule delay costs, etc. In that case, the *monetarization* of the utility loss can be computed using eq. (2) by replacing τ by C, χ and θ^{-1} have then the dimension of money. A meaningful dimensionless parameter to measure the relative impact is $\phi = \frac{\chi}{\langle C\rangle}$ where $\langle C\rangle$ is the average travel cost.

3 Simulation of non-recurrent congestion

So far, nothing has been assumed concerning the probability distribution of travel times (or travel costs) $f(\tau)$. We propose to compute $f(\tau)$ by performing explicit Monte-Carlo simulations where the stochasticity is induced by random incidents. The incidents disrupt the capacities of the simulated road network. The corresponding travel time delays are computed using state-of-the-art mesoscopic traffic with METROPOLIS (4; 5).

3.1 Traffic simulation

As it can be seen in the existing literature, the value of information systems is usually measured by modeling the information systems themselves via simulation, and comparing it with a base case. We use a new approach of modeling a simulation without an explicit information system, but a *departure time* feedback, and *historical* and *instantaneous* information available to the users. This is explained in the following sections.

3.2 Information-based traffic simulation

METROPOLIS is a simulation tool intended to be a fully dynamic model that features within-day as well as day-to-day traffic dynamics. Some of its features are summarized below that are relevant to our current study. Traffic models have usually two main components: (a) a supply model that describes how the traffic conditions evolve in the road network given the users' driving choices and (b) a demand model that describes the users' behavior given their driving environment and other drivers' decisions. When the description is time-dependent, the supply model is often a DTA (Dynamic Traffic Assignment) model. The architecture of METROPOLIS considers information as a third component of transportation models (see Fig. 1). Information means, in the broad sense, any piece of knowledge that can play a role in the users' travel decisions: experienced travel times, shortest routes, congestion levels on downstream links, etc. Information is user specific since it corresponds to the users' perception of traffic conditions. For the supply side, the mesoscopic level of details assumes that congestion is described by functions that depend on aggregate variables such as density and flow. The congestion is located on links and all the data about the network state is embedded in the *information* block. Some external information can also take place in that block, like information provided by ATIS devices, radio broadcast or other technologies (though this is not yet implemented). In turn, users are assumed to be disaggregated and make use of the available information (which can be user specific) to perform their travel decisions. The supply consists of the coded car network: zones, intersections and links. Both supply and demand data interact in the simulator that computes sequentially the mode choice, the departure time choice and eventually, the route choice for each simulated user. Destination choice is not yet implemented.

The traffic assignment procedure uses an event-based approach. As the traffic simulator proceeds, the *within-day* time evolves. Supply outputs are collected as dynamic level of services: time-dependent travel times and time-dependent traffic flow patterns on each road section. On the demand side, actual users choices are collected: departure time, route and mode taken. These results may be aggregated to compute Measures Of Effectiveness (MOEs) or performance indexes. When the simulation of the considered period (e.g. morning peak) is over, the learning process uses these results to update the historical information. The learning process corresponds to the *day-to-day* dynamics. The departure time model and the mesoscopic DTA are responsible for the *within-day* dynamics. In absence of exogenous incidents, the system converges toward a stationary state. Indeed, users improve continuously their knowledge of the traffic condition. At some point, they will not be able to improve anymore their travel choices in order to minimize their generalized travel cost. The simulator computes a generalization of Wardrop's first principle: at equilibrium, no user can strictly decrease his generalized travel cost by changing either his mode, his departure time or his route.

3.3 Learning process

The learning process can be seen as a black box that models day-to-day dynamics: it captures the fact that users take into account their past driving experiences in future travel decisions. The data

processed here is the *information* that must be understood in a broad sense. It groups together historical drivers' information and instantaneous information. It consists of any data relevant to user travel choices, like data required to estimate time-dependent shortest routes, but also the expected variability of travel times under non-recurrent situations or data relevant to road tolls.

Input *Historical* information denotes the accumulated knowledge that users gain from using the transportation system days after days. Should traffic conditions never change (which is never the case in real life), this historical information would reach a stationary state and remain constant in a well behaved model. Instantaneous information refers to what each user may learn during the within-day simulation: perception of actual traffic conditions, forecasts or exogenously simulated information provided to the drivers, etc. *Instantaneous* information is continuously updated as the dynamic assignment proceeds, while historical information is only updated daily. Information is partially shared among the drivers (public information) and partially specific to individual drivers (private information). Information determines all the travel choices modeled in the simulator: Mode, departure time and route choices. The flexibility of this information handling allows, for instance, to provide different users segments with different sets of instantaneous information. *External* information consists of any information provided from outside of the transportation system by means of information technology like radio, variable messages signs, on-board computers, etc. In principle, *external* information can also be included in the framework, but as of now the travel choices are driven only by *historical* and *instantaneous* information.

Output Once a simulation day (or morning rush) is completed, the historical information is updated according to the driving conditions experienced by the simulated drivers. The updating process is based on a Bayesian mechanism that combines the historical information and the last day changes in the network usage according to the user characteristics (i.e. perception and cognitive abilities). Since these mechanisms are typically very expensive to implement literally at the numerical and computing memory point of view (see (1)), we decided to resort to heuristic laws similar to (15). The historical information available to users on day $\omega + 1$ is the output of the learning process on day ω, that is, the accumulated knowledge of the ω previous days (i.e. a Markov process of order 1). For a given O-D pair, the expected travel time when departing at time t_d on day $\omega + 1$ is computed as follows:

$$E^{\omega+1}\left(\tau\left(t_d\right)\right) = \left(1-\lambda\right) E^{\omega}\left(\tau\left(t_d\right)\right) + \lambda\tau^{\omega}\left(t_d\right) \tag{3}$$

Numerous experiments with METROPOLIS have shown that the updating process (3) is fairly robust and converges in a few dozen iterations. A typical value used in practice is $\lambda = 0.1$. It can be interpreted as if 10% of the users are revising their decisions every day. For a discussion on the optimal strategy of partially informed drivers in this context, see (3).

3.4 Incidents and impacts

A possible solution to evaluate the impacts of non-recurrent traffic incidents would be to use a static traffic assignment approach. Given the probability distribution of occurrence of an incident $f_I(t)$, links could be characterized by a stochastic capacity and fully stochastic assignment algorithms can then be used to solve the problem (see (10)). This approach lacks however an important aspect which is within-day departure time choice adjustment, i.e. if users expect to encounter unreliable traffic conditions on certain part of the network at certain moment of the day, they might decide to schedule their trip at another time. Also, the impacts measured by the static approach are those of a supposedly long-term situation where the users have somehow discovered $f_I(t)$ by experiencing traffic conditions. But the transient states might exhibit larger impacts that cannot be measured without taking into account explicit day-to-day dynamics, and as stated earlier, we intend to show that ATIS is most effective in dispersing congestion during these transient states. Hence we use a dynamic traffic assignment approach for the simulation. Note that even a DTA, but without departure time feedback capability is not sufficient, since some mechanism has to be responsible for the adaptation of (exogenous) departure time profiles as shown in (15). The architecture of METROPOLIS overcomes this as explained in Fig. 1.

Exogenous traffic hazards are straightforward to implement in event-based models such as METROPOLIS. At the beginning of each day (or morning peak), a random number $R \equiv U(0,1)$ is drawn for each link where a potential incident can happen. If $R < p$ then an incident happens on that link. In our experiments, the probability of occurrence p is the same for a selected subset of important links of the network (e.g. motorways) and $p = 0$ for the rest of the network. The incidents are characterized by a capacity drop of 50% that lasts for the whole morning peak. The linear bottleneck congestion model is applied throughout the whole simulated period. Note that incidents could also be defined with a specific duration (e.g. half an hour for a stopped vehicle blocking a lane). Nevertheless, these incidents still affect the departure time choice since roads are usually under-used (i.e. below capacity) at the beginning (and end) of the peak period. If the capacity drops often during the peak period, some users will consider departing earlier or later to avoid congestion. Vehicles that reach an intersection are informed if an incident has happened on a downstream link so that en-route diversion is possible. Note that users gather the information of the incident only when they reach the origin intersection of the link on which the incident has happened and are not informed of the incident beforehand. Hence they make a decision to divert their route based on this local information (i.e. time delay on the incident-ridden link) and historical information (i.e. travel times of the alternative routes possible).

Therefore, the non-recurrent congestion scenario can serve as a benchmark for comparing the effects of non-recurrent congestion with a scenario having an information system, where users would be informed of the incidents occurring on far-off links. The increase in travel cost $\triangle TC$ due to incidents (compared to an incident-free scenario), is assumed to be contributed by the global effect of the reduction of capacity $\triangle Cap$, and the effect of a the lack of information system $\triangle IS$.

$$\triangle TC = \triangle Cap + \triangle IS$$

In the absence of any Braess-like paradox, we can assume that $\Delta Cap > 0$ and therefore $\Delta IS \leq \Delta TC$. Our methodology consists in measuring ΔTC as an upper bound of the potential benefits ΔIS from any information systems. We measure ΔTC through out first set of simulations as a proxy to the maximal travel cost savings that can be squeezed out of a sophisticated information system.

4 Simulation results

4.1 Control runs

A first set of experiments is performed on the well-known network of Sioux Falls to understand the day-to-day adjustment process. About $50,000$ individual car trips are simulated. Mode choice is disabled. Incidents are introduced on four links that are located on the north-south corridors identified on Fig. 2. Three simulations are performed:

1. for the base case scenario without incidents,
2. for a scenario with incidents in the corridors from day #20 on,
3. for a scenario with incidents in the corridors between day #20 and day #50.

Fig. 3 presents the travel costs for the first hundred iterations of the three cases. During the first twenty iterations the three curves are overlapping which is consistent with the fact that the random number generator is initialized with the same seed. The base case system exhibits some oscillations. The peak at iteration #10 shows that the system is not yet completely stabilized. This is due to the fact that the simulation model is stochastic and that the exponential smoothing process does not lead here to a unique situation but rather to a set of stationary states. We will show below that this does not hinder the evaluation of non-recurrent congestion. It can be seen that the impacts of the incidents are quite large on day #20 since users are unaware of the incidents and did not yet had the opportunity to adapt their travel habits. This would be the impact reported by a pure DTA model *without* feedback on departure time choice. Two things are noteworthy in this example: the adaptation in the case of the standing incidents is such that the travel cost follows a decreasing trend. Also in the case of the third scenario, the decrease in travel cost does not happen as soon as the exogenous shocks cease (i.e. on the 50th day) because of the inertia of the system. The comparison between the second scenario and the base case provides an upper bound of the potential travel cost savings of an information system: about 50 cent per trip (equivalent to a travel delay of 3 minutes). The situation on day #20 is strikingly different: the difference is tenfold larger.

4.2 Impacts of incidents

The second set of experiments is performed on a real-world example for the Paris area. The coded network consists of about $17,000$ links and more than $3,000,000$ individual trips are simulated for each morning period (the computation of a single iteration takes around 20 minutes on an Apple G5 with a clock speed of 2Ghz). Both commuting trips and non-commuting trips are simulated to take off-peak congestion into account. The incidents are introduced on the major roads of the area , defined as the roads that have at least three lanes (see Fig. 4). Several simulations are run for different probabilities of occurrence p ranking from 0 to 1. In each case, 50 iterations (days) are performed. The corresponding global indicators are reported in Tab. 1. The impacts are important because each traveler uses on average at least one road section that belongs to the major roads (there are about $1,000$ major-road links and the average trip length is 16 links). The schedule delay cost measures the penalties incurred by travelers arriving too early or too late at their destination. Obviously, drivers arrive later than expected when the level of incidents increases. Note that the indicators show that the case $p = 0.5$ is worse than the risk-free case $p = 1$, even if the average capacity of the overall system is higher. As explained earlier, by risk-free we refer to a scenario which is certain (definite incidents: $p = 1$, or no incidents at all: $p = 0$). This discrepancy between the stochastic and the risk-free situation is evaluated in the next section.

4.3 Utility loss

A survey (6) was administered by one of the authors to estimate the level of risk aversion of drivers in the Paris area. Respondents are asked lottery-type questions where they have to rank different lotteries assigned to different level of risk. The same method has been used also to compute the distribution of risk aversion of private investors (see `www.RiskDynaMetrics.com`). In the transportation area, the users are asked to compare route with different variability of travel times, while in the finance application, respondents are asked to compare different financial products which differ according to their level of risk and return. In both case the alternative are naturally ranked and an ordered probit model is used to estimate the level of risk aversion of the users and to determine the different factors and socio-economic characteristics which influence the level of risk aversion. The results show for instance that men are less risk averse than women, and that blue collars are more risk averse than white collars. Basically, the socio-economic characteristics and the purpose of the trip do influence the level of risk aversion. The major outcome of the survey (6) is the provision of the distribution of risk aversion in the population of drivers.

In the simulations reported in this paper, users still base their travel decisions on expected costs updated using eq. (3). The travel cost includes free flow cost, queuing cost and schedule delay cost. Users attempt to maximize their utility $U(t) = -C(t)$ and do not use the CARA utility specification given by (1). Nevertheless, the utility loss can be computed for the stochastic system at hand. All users are assumed to be risk neutral except for an infinitesimal fraction that are risk-averse. Their utility loss can be computed using (2). Since it is an infinitesimal

fraction, their decision does not affect the overall system. The purpose here is to demonstrate first the importance of the utility loss. The long term goal is to develop a completely consistent framework that would use (1) as the user objective. In that case users would also adapt their departure time and route choices according to the variability of the travel conditions.

The compensation χ introduced above for travel *time* variability is measured on the Sioux Falls example. This compensation is extended here to the compensation of travel *cost* variability instead. This value can then be interpreted as the utility loss due to uncertainty (or *monetarization value* of uncertainty). The variability $f\left(C\left(t\right)\right)$ is computed by recording the total travel costs C_i for T =300 iterations. Therefore, the monetarization of uncertainty (per trip) is computed as follows:

$$\chi_\theta = \frac{1}{\theta} \ln \left\{ \frac{1}{T} \sum_{i=1}^{i=T} e^{\theta C_i} \right\} - \frac{1}{T} \sum_{i=1}^{i=T} C_i$$

Incidents are introduced during the T iterations for a given level of probability p. Several sets of simulation are performed for values of p ranking from 0 to 1. Fig. 5 presents the results of the evaluation of χ_θ as a function of the level of risk aversion θ. As shown in the control runs (Fig. 3), travel costs can oscillate even without incidents because of the stochastic nature of the traffic simulation. Nevertheless, those oscillations are rather small, which explains why the risk-free situations ($p = 0$ and $p = 1$) correspond to the bottom curves associated with an (almost) null cost for uncertainty. As for the Paris example, the case $p = 0.5$ yields the highest cost of uncertainty. Obviously it is not a symmetric dependence (see the discrepancies between $p = 0.4$ and $p = 0.6$). If we assume $\theta = 10/\$$ (a value compatible with the survey results of (6)), we get the relative impacts is $\phi = \frac{\chi}{\langle C \rangle}$ reported in Tab. 2. The travel time equivalent of the compensation, τ_χ, is computed by assuming a VOT of $10\$/h$.

5 Concluding comments

We have developed in this paper a method to improve the measurement of the cost of uncertainty by modeling explicitly recurrent and non-recurrent congestion in transportation systems. The methodology is built upon METROPOLIS, a dynamic traffic simulation tool that is able to handle very large realistic networks, within-day dynamics and day-to-day dynamics. Random incidents are introduced in the system and users react to them by adapting their departure time and route on a within-day and day-to-day basis. We have stressed the importance of day-to-day adjustment process by showing that the distinction between recurrent and non-recurrent congestion is blurred: if the same incidents occur with a given probability on a long period, users might eventually learn how to adapt and take into account travel time variability as an additional cost. This additional cost or monetarization of uncertainty has been estimated using both empirical measures and simulation results. This evaluation is crucial to the designers and managers of ATIS and ATMS since it gives a benchmark of the potential benefit that can be obtained with such technologies.

The paper provides two important insights. First, the long-term efficiency of ATIS technology in term of travel time savings is an order of magnitude smaller than its short-term efficiency which has been the main focus in the literature. This is consistent with previous work (8) but depends largely on the probability distribution of incidents. Therefore, one could but recommend to take into account the likeliness of the incidents when performing a cost-benefit analysis. Second, the long-term efficiency of ATIS might be found elsewhere in the reduction of uncertainty. Our results show that the utility loss due to uncertainty is of the *same* order of magnitude as that of the travel cost. The cost of uncertainty is as important as the cost of congestion. Therefore, simulation models should integrate risk aversion using concave utility functions.

Acknowledgments

The authors are grateful to Nathalie Picard and Robin Lindsey for their wise comments on a preliminary version of the paper. We would like to thank Dany Nguyen from IAURIF and Kiarash Motamedi who kindly helped us to setup the data for the simulations. Computing resources were provided by CoLab, ETHZ. This work was partially sponsored by the ANR project RiskAttitude and by the Institut Universitaire de France. We benefited from the careful reviews of two anonymous referees and from Saurabh Kasturia.

References

[1] ARENTZE, T., AND TIMMERMANS, H. Representing mental maps and cognitive learning in micro-simulation models of activity-travel choice dynamics. *Transportation 32* (2005), 321–340.

[2] BROWNSTONE, D., AND SMALL, K. A. Valuing time and reliability: Assessing the evidence from road pricing demonstrations. *Transportation Research Part A: Policy and Practice 39* (2005), 279–293.

[3] CHANCELIER, J.-P., DE LARA, M., AND DE PALMA, A. Risk aversion, road choice, and the one-armed bandit problem. *Transportation Science Vol. 41, No. 1* (2007), pp. 1–14.

[4] DE PALMA, A., AND MARCHAL, F. Dynamic traffic analysis with static data: some guidelines from an application to Paris. *Transportation Research Record 1756* (2001), 76–83.

[5] DE PALMA, A., AND MARCHAL, F. Real case applications of the fully dynamic METROPOLIS tool-box: an advocacy for large-scale mesoscopic transportation systems. *Networks and Spatial Economics 2(4)* (2002), 347–369.

[6] DE PALMA, A., AND PICARD, N. Route choice decision under uncertainty. *Transportation Research Part A: Policy and Practice 39* (2005), 295–324.

[7] EMMERINK, R., AXHAUSEN, K., PIJKAMP, P., AND RIETVELD, P. The potential of information provision in a simulated road transport network with non-recurrent congestion. *Transportation Research Part C 3*, 5 (1995), 293–309.

[8] HALL, R. Non-recurrent congestion: How big is the problem? Are traveler information systems the solution? *Transportation Research Part C 1*, 1 (1993), 89–103.

[9] HALLENBECK, M. E., ISHIMARU, J., AND NEE, J. Measurement of recurring versus non-recurring congestion. Tech. Rep. WA-RD 568.1/568.2, Washington State Transportation Center, 2003.

[10] HAZELTON, M. Total travel cost in stochastic assignment models. *Networks and Spatial Economics 3* (2003), 457–466.

[11] LINDLEY, J. A. Urban freeway congestion: quantification of the problem and effectiveness of potential solutions. *ITE Journal 57*, 1 (1987), 27–32.

[12] LINDLEY, J. A. Urban freeway congestion problems and solutions: an update. *ITE Journal 59*, 12 (1989), 21–23.

[13] LIU, H., BAN, X., RAN, B., AND MIRCHANDANI, P. Analytical dynamic traffic assignment model with probabilistic travel times and perceptions. *Transportation Research Record 1783* (2002), 125–133.

[14] MACHINA, M. Choice under uncertainty : Problems solved and unsolved. *Journal of Economic Perspectives 1* (1982), 281–296.

[15] MITHILESH, J., MADANAT, S., AND PEETA, S. Perception updating and day-to-day travel choice dynamics in tra c networks with information provision. *Transportation Research Part C 6* (1998), 129–212.

[16] MUIZELAAR, T., AND VAN AREM, B. Drivers preferences for traffic information for non-recurrent traffic situations. In *TRB 2007 Annual Meeting CD-ROM* (2007).

[17] NOLAND, R. B., AND SMALL, K. Travel time uncertainty, departure time choice, and the cost of the morning commute. *Transportation Research Record 1493* (1995), 150–158.

[18] PEIRCE, S., AND LAPPIN, J. Acquisition of traveler information and its effects on travel choices: evidence from a seattle-area travel diary survey. Research paper 13813, Volpe National Transportation Systems Center, 2003.

[19] SKABARDONIS, A., VARAIYA, P., AND PETTY, K. Measuring recurrent and non-recurrent traffic congestion. *Transportation Research Record 1856* (2003), 118–12.

[20] VON NEUMANN, J., AND MORGENSTERN, O. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1980.

List of Tables

Table 1: Impacts of non-recurrent congestion.

p	0	0.25	0.4	0.5	0.75	0.9	1
Travel cost [$]	7.6	11.3	11.5	12.0	12.6	12.0	11.9
Schedule delay cost [$]	2.2	3.2	3.3	3.4	3.5	3.3	3.3
Travel time [min.]	25.0	37.4	38.2	40.0	42.1	40.0	39.8
Change in consumer surplus [$]	0	-1.9	-2.0	-2.3	-2.3	-2.2	-2.1
Congestion index [%]	28.9	55.4	58.2	60.8	66.9	64.0	65.0
Mileage [10^6km]	56.5	59.8	59.8	60.2	60.6	60.0	59.8
Early arrivals [%]	49.5	48.2	47.8	48.1	45.6	46.6	46.1
Late arrivals [%]	29.7	34.7	35.1	35.0	37.5	36.2	36.5

Table 2: Evaluations of the cost of uncertainty for an individual with a risk aversion level of $\theta = 10/\$$ and a VOT of $10\$/h$.

p	0.25	0.40	0.50	0.60	0.75	0.90
χ[\$]	2.10	2.55	3.35	1.35	1.15	0.45
τ_χ[min.]	12	15	20	8	7	3
ϕ	0.37	0.45	0.57	0.23	0.20	0.08

List of Figures

Figure 1: Architecture of METROPOLIS. Blue dashed arrows correspond to within-day dynamics. Green solid arrows correspond to day-to-day dynamics.

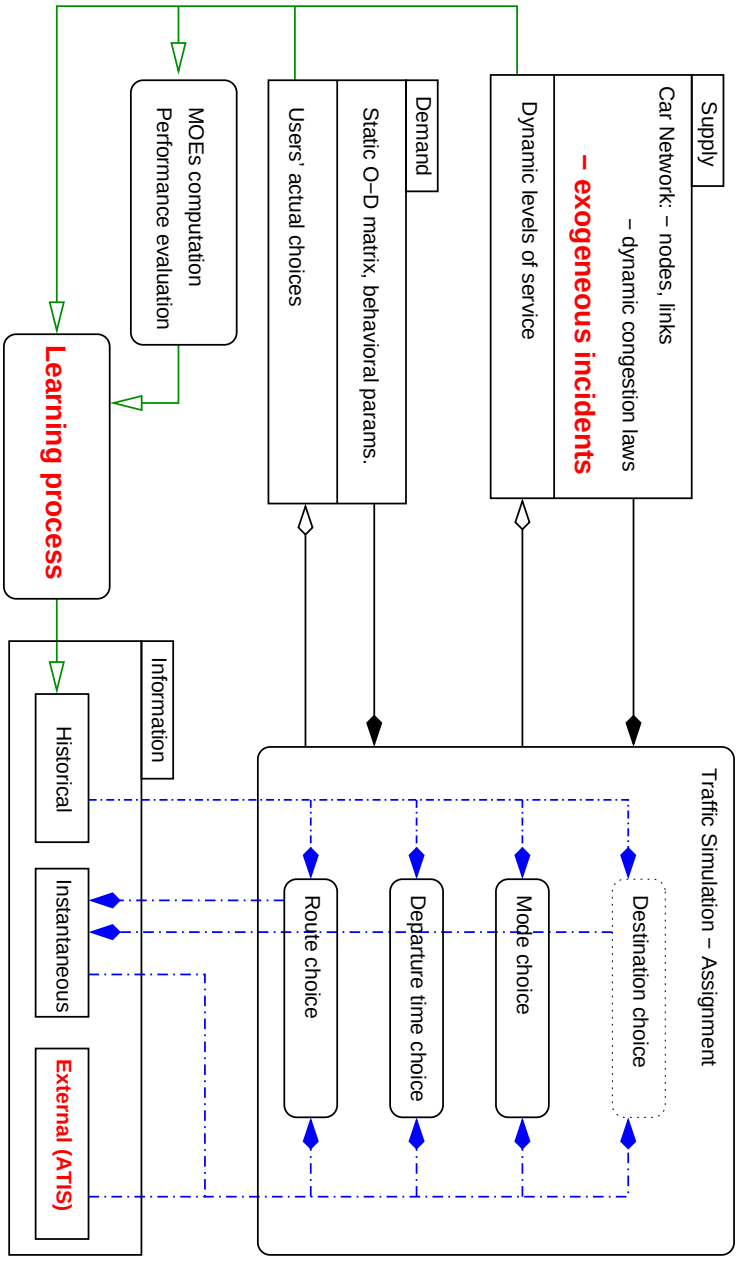

Figure 2: Congestion index on the Sioux Falls network in the base case scenario (without incidents). Incidents are introduced on the most loaded north-south corridors (with index ≥ 0.6).

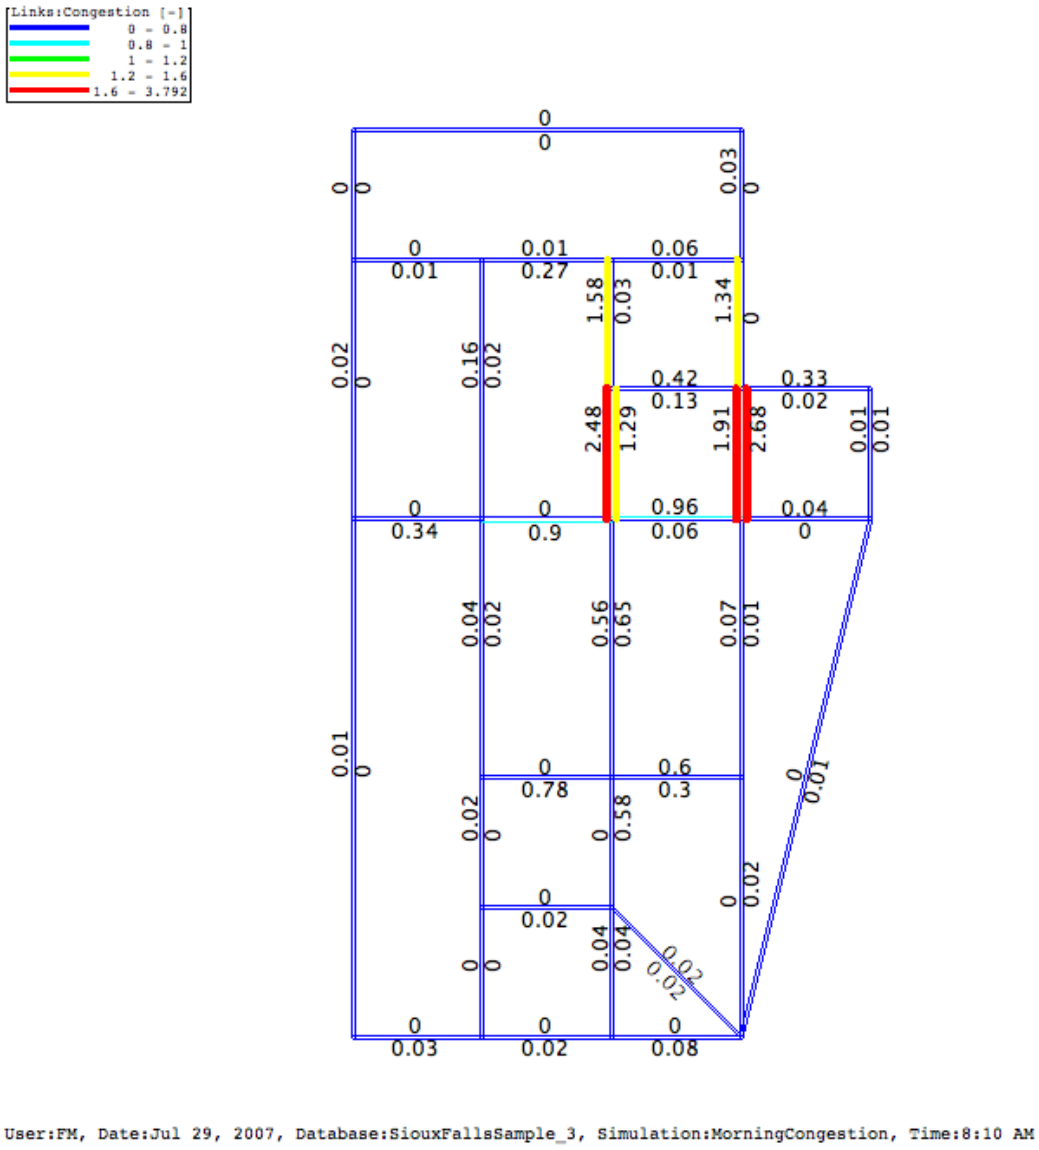

Figure 3: Control runs. Three simulations with and without incidents. Day-to-day evolution of the total travel cost.

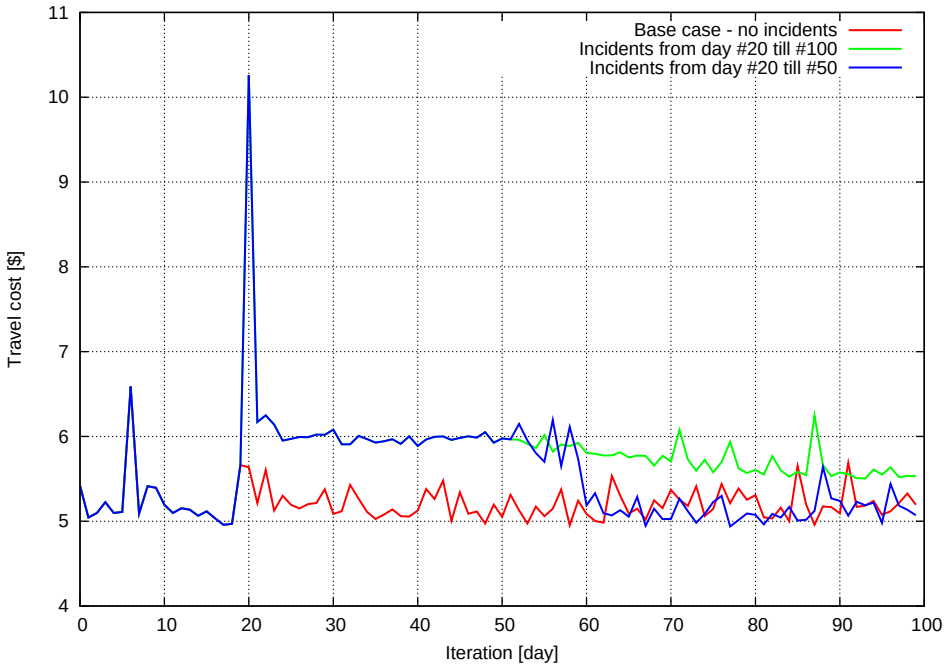

Figure 4: Ile-de-France area surrounding Paris. Incidents are introduced on the major roads (red).

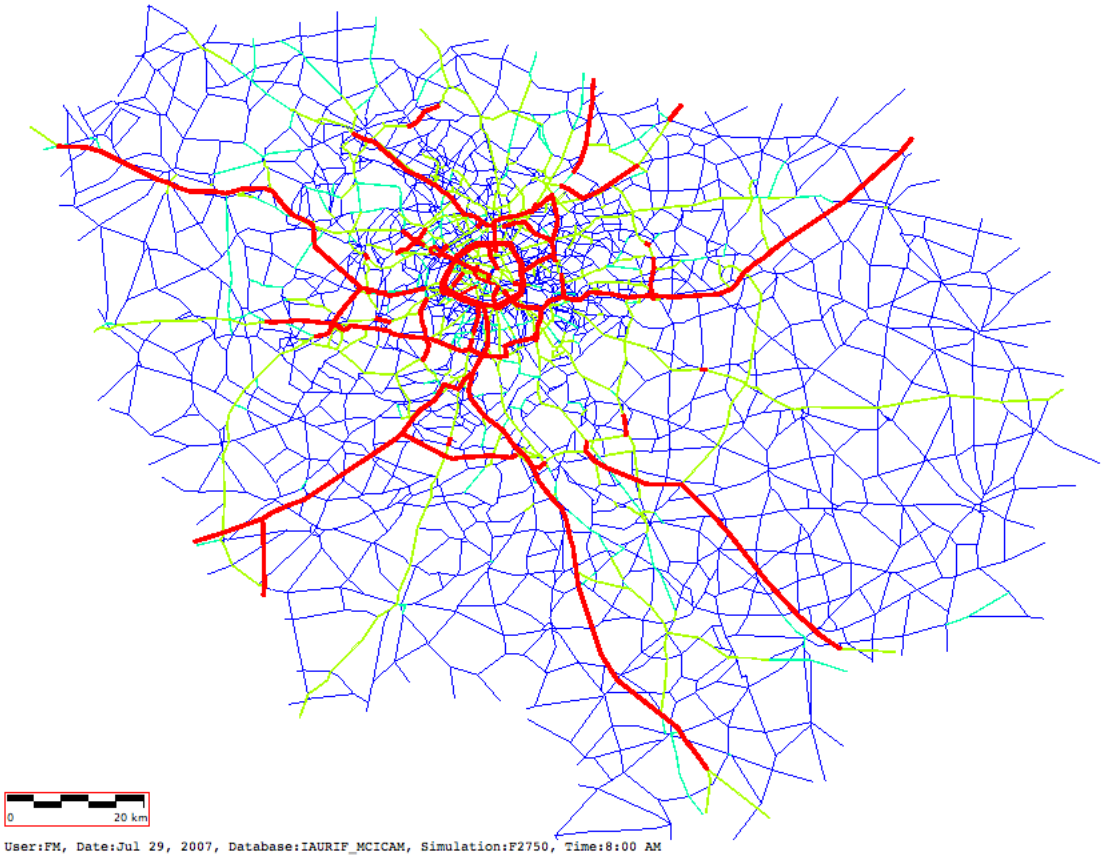

Figure 5: Monetarization of uncertainty. Measurements for different probability p of occurrence.