Inertia in travel choice:
The role of risk aversion and learning

Caspar G. Chorus[1],*, Benedict G.C. Dellaert**

Department of Technology, Policy & Management, Delft University of Technology
**Erasmus School of Economics, Erasmus University Rotterdam*

Abstract

This paper contributes to literature by showing how travellers that make normatively rational choices exhibit inertia during a series of risky choices. Our analyses complement other studies that conceive inertia as the result of boundedly rational or even non-deliberate, habitual decision-making. We start by presenting a model of risky travel choice based on Bayesian Expected Utility maximization premises. We show how inertia emerges due to a learning-based lock-in effect: travellers learn about risky attributes such as travel times and costs by observing the performance of a chosen alternative. Given risk aversion this implies that repeatedly choosing the same alternative from an initial set of equally risky alternatives is a rewarding strategy. We then extend our model to capture forward-looking behaviour and the availability of travel information: we show how inertia grows somewhat slower among forward-looking travellers, and how the provision of multimodal pre-trip travel information may somewhat reduce inertia growth, to the extent that it is reliable. Combining our findings with the large body of literature on inertia emerging from boundedly rational and habitual behavior, we argue that expectations regarding the potential of travel demand measures to counter inertia should be modest at best.

Keywords: Inertia, Bayesian learning, risk aversion, lock-in effects, mode choice.

1.   Introduction

For many years, there has been a considerable interest among travel behavior researchers in understanding inertia (see Gärling & Axhausen (2003) for a relatively recent overview). As a result, many studies have illuminated the role of inertia in the context of route-choices (e.g. Mahmassani & Chang, 1987; Jou et al., 2005; Srinivasan & Mahmassani, 2000; Bogers et al., 2005) and especially mode-choices (e.g. Aarts et al., 1997; Aarts & Dijksterhuis, 2000; Cantillo et al., 2007; Eriksson et al., 2008; Gardner, 2009). This amount of effort put into understanding traveler inertia reflects transportation policy-makers'

---

ambitions to change travel behavior away from established patterns with an aim to increase the efficiency and sustainability of transport network usage (e.g. Ministry of Transport, Public Works and Water Management, 2002; Department of Transport, 2004; Commission of the European Communities, 2007). Such behavioral adaptation is by definition difficult to achieve when travelers are inert (e.g. Gärling et al., 2002; Chorus et al., 2006a).

Predominantly, the travel behavior research community frames inertia as resulting from bounded rationality and/or habitual behavior: exploring and testing new travel options consumes time, effort and attention. Since time, effort and attention are scarce resources (Simon, 1978; Shugan, 1980; Payne et al., 1993), most inertia-models postulate that it is a good decision-strategy to stick with an alternative that one knows to perform reasonably well, whereas one could also try to find the best performing option for each new trip.

In this paper we show that the emergence and growth of inertia can be explained without making these assumptions of bounded rationality or effort-accuracy trade-offs: we show that even travelers that (behave as if they) consider alternative travel options for each trip, maximize Expected Utility and learn from past observations in a strictly Bayesian manner, exhibit inertia, as long as i) they dislike risk and ii) part of the quality of travel alternatives is only revealed upon usage. In other words: even travelers that live up to high standards of 'unbounded' rationality[2] are shown to prefer sticking to an alternative that is chosen before, purely because it has been chosen before. The intuition behind this result can be put as follows: travelers learn about a travel mode's quality by observing the performance of a chosen alternative (see, for example, Chorus et al. (2007) for an empirical underpinning of this intuitive claim). Given risk aversion this implies that repeatedly choosing the same alternative from an initial set of equally risky alternatives is a rewarding strategy. We formally derive this learning-based cognitive lock-in effect within the context of Bayesian, Expected Utility maximizing travel-choices. We also show that inertia grows faster when risk and risk aversion increases.

In addition, we illustrate how forward-looking travelers are somewhat less inclined to develop inertia than those who are only focused on the current trip. The intuition behind this finding can be put as follows: learning by means of observation is most effective when the level of knowledge about an alternative is relatively low. As a result, a forward looking traveler knows that exploring relatively unknown alternatives (being the opposite of inertia) pays off as it may lead to substantial gains in utility (note that a similar result was obtained by Arentze & Timmermans (2005) in the context of destination choices). However, we also show that the presence of even small levels of risk aversion will strongly diminish the negative effect of forward-looking behavior on inertia strength. We proceed by illustrating how the costless provision of multimodal travel information may slow down inertia emergence to the extent that the information is considered reliable: when reliable information about all alternatives is available, the role of making observations for the sake of learning becomes irrelevant, and the learning-based cognitive lock-in effect described above vanishes.

---

[2] We define a choice process as unboundedly rational, or from here on, rational, when it involves i) the consideration of all alternatives available in the choice set, ii) Expected Utility maximization-based evaluation and iii) Bayesian learning. A more detailed discussion of what constitutes rationality (a hotly debated topic throughout the social sciences) is outside the scope of this paper, and can be found in, for example, March (1988).

Finally, as a case study, we show how our model predicts that the forced temporary abandonment of a current travel mode may help break inertia to the extent that a) the quality of the current alternative is uncertain in the eyes of the traveler, and b) an alternative travel mode performs satisfactory. This result is fully in line with empirical literature. We will argue that this suggests that our model of rational inertia can help explain empirically observed inertia patterns that have predominantly been interpreted in light of boundedly rational or habitual behavior.

The paper's scope is determined as follows: first, in line with most inertia-related research and policy-making, we focus on a travel mode choice between a car and a train option. However, obtained results are informative for other, multinomial, contexts as well – for example involving travelers' departure time- and route-choices. Second, we assume that travelers hold subjective probabilities with respect to the performance, or quality, of an alternative. Quality is conceived as a composite function of tastes, and of different quality-aspects, some of which are 'tangible' (e.g. travel times, costs) whereas others are less tangible (e.g. scenery, crowdedness in a train). By making observations, the traveler gets an increasingly good idea of how much he likes an alternative. Specifically, we consider the situation where a traveler faces a trip towards a new destination for the first time, so that uncertainty about the quality of travel alternatives is due to a lack of experience with the alternatives, rather than being due to day-to day variability. Third, as discussed above, we focus on rational behavior. This does not imply that we believe that boundedly rational or habitual mechanisms are less important for the understanding of traveler inertia. The motivation underlying our focus is that we wish to contribute to the abundant body of literature relating to inertia by adopting a different, more normative perspective. As will be seen further below, our findings provide additional support for the general result obtained in most of the studies that adopt a boundedly rational or habitual perspective: inertia is difficult to avoid, and difficult to break.

The remainder of this paper is organized as follows: section 2 presents a model of Bayesian Expected Utility-maximizing travel mode-choice behavior, and formally shows how inertia emerges when travelers dislike risk. Section 3 extends this model to incorporate forward-looking behavior, and shows how the balance between forward-looking and risk aversion determines the level of inertia. In section 4, we discuss how the presence of multimodal travel information impacts inertia strength, depending on the anticipated reliability of the information. Section 5 presents a case study concerning the impact of a forced highway closure on car inertia. Section 6 presents conclusions, reflections, and possible avenues for further research.

2.   Inertia among Bayesian, Expected Utility-maximizing travelers

2.1.   A model of Bayesian, Expected Utility-maximizing travel mode-choice behavior
Consider a traveler that has changed jobs and faces a choice between two travel modes – car and train – for his daily commute towards his new job location. The quality x of each mode – being a function of tastes and attributes that are relevant to the traveler – is anticipated by the traveler as a risky variable given the absence of experience with both modes in the context of the changed commute destination. Assume that car and train quality are both anticipated in terms of a normal distribution whose mean ($\hat{x}$) and variance (VAR)

represent expected quality and quality uncertainty: $f\left(x_{car}^{t}\right) \;=\; N\left(\hat{x}_{car}^{t}, \sqrt{VAR_{car}^{t}}\right)$, $f\left(x_{train}^{t}\right) = N\left(\hat{x}_{train}^{t}, \sqrt{VAR_{train}^{t}}\right)$, where t denotes trip number (or day number, when we assume that each day a trip is made). Note that although many other distributions may be used here, the normal distribution is convenient when specifying a Bayesian learning process. In line with previous studies in transportation (Allen et al., 1985; Harker & Hong, 1994; Lam & Small, 2001; Liu et al., 2007), and to keep the tractability of subsequent derivations at a reasonable level, we assume that the traveler evaluates travel modes based on a mean-variance linearization[3] of Expected Utility (EU). In the context of non-forward-looking behavior, EU is equated with *instantaneous* Expected Utility ($E\tilde{U}$: with the term instantaneous, we mean to reflect that the utility of a particular travel mode is only based on its anticipated performance during the current trip):

$$E\tilde{U}_{car}^{t} = \beta_{x} \cdot \hat{x}_{car}^{t} - \beta_{VAR} \cdot VAR_{car}^{t}, \;\; E\tilde{U}_{train}^{t} = \beta_{x} \cdot \hat{x}_{train}^{t} - \beta_{VAR} \cdot VAR_{train}^{t} \qquad (1).$$

Here, β_{x} and β_{VAR} are nonnegative, the latter reflecting the level of risk aversion[4]. In line with empirical evidence (e.g. Hey, 1995) and theoretical arguments (e.g. Manski, 1977), we assume that the traveler's choice behavior is – to some extent – unstable. In other words, the traveler possibly chooses differently in two consecutive choice situations, even when they are exactly the same in terms of levels of quality (uncertainty) associated with the two modes. We capture this behavioral volatility by adding to the travel modes' utility iid error terms ε_{car}^{t} and ε_{train}^{t}, drawn from an Extreme Value Type I distribution with variance $\pi^{2}/6$. This specification results in the standard binary logit formulation of probabilities $P\left(y^{t} = car\right)$ and $P\left(y^{t} = train\right)$, where y^{t} denotes the option chosen during trip t.

Upon choosing one of the travel modes (e.g. the car mode) and executing the trip, the traveler makes an observation ($\breve{x}_{car}^{t}$) of the mode's quality. The traveler knows that this observation will help him provide a more accurate assessment of the mode's quality. However, the traveler also believes that he is unable to make a perfectly reliable assessment of an alternative's quality by making only one observation (remember that quality is defined as a function of the traveler's tastes and a number of tangible and less tangible quality-aspects). Formally, assume that a quality observation is a noisy signal of actual quality x_{car} in the sense that $f\left(\breve{x}_{car}^{t} \mid x_{car}^{t}\right) = N\left(x_{car}, \sqrt{VAR_{obs}}\right)$. That is, the traveler believes that, if the actual quality-level equals x_{car}, observed quality during trip t is normally distributed with mean equaling x_{car} (in other words: he believes observations provide

---

[3] Note that in the context of normally distributed quality-levels, and assuming that the function mapping quality to utility is exponential (implying Constant Absolute Risk Aversion), Expected Utility is in fact given by a mean-variance formulation (de Palma and Picard, 2006a, b).

[4] As a referee pointed out, it has been shown that risk aversion may be affected by learning processes (e.g. Erev & Barron, 2005). However, in this paper we assume that risk aversion parameters are static over time, and have an effect on learning dynamics. The study of the inverse relationship is considered an interesting avenue for further research.

unbiased measurements of actual quality), and variance equaling VAR_{obs}. The magnitude of VAR_{obs} reflects the extent to which the traveler believes that an observation of quality during the execution of an alternative is an unreliable measurement of actual quality. In other words: higher levels of VAR_{obs} reflect that the traveler distrusts his own observations and believes that it takes time to 'get to know' the alternative and appreciate its quality. We assume that VAR_{obs} does not differ between the car and the train mode, hence the absence of a mode-specific superscript.

Given these assumptions, the travelers updated perception of quality of the car mode, after having made trip t using the car mode and having observed a quality level $\breve{x}_{car}^{t}$, denoted $f\left(x_{car}^{t+1}\middle|\breve{x}_{car}^{t}\right)$, is given by applying Bayes' Theorem[5] (e.g. Edwards et al., 1963):

$f\left(x_{car}^{t+1}\middle|\breve{x}_{car}^{t}\right) = N\left(\hat{x}_{car}^{t+1}, \sqrt{VAR_{car}^{t+1}}\right)$, where:

$$\hat{x}_{car}^{t+1} = \frac{\left(VAR_{car}^{t}\right)^{-1}\cdot\hat{x}_{car}^{t} + \left(VAR_{obs}\right)^{-1}\cdot\breve{x}_{car}^{t}}{\left(VAR_{car}^{t}\right)^{-1} + \left(VAR_{obs}\right)^{-1}} \text{ and } VAR_{car}^{t+1} = \frac{VAR_{car}^{t}\cdot VAR_{obs}}{VAR_{car}^{t} + VAR_{obs}} \qquad (2).$$

In words: updated quality perceptions are a weighted average of prior beliefs and observed quality. Weights reflect perceived reliability of prior beliefs and observations, respectively: when the traveler distrusts (trusts) his own observations, updated perceptions of quality are relatively close to initially anticipated quality (observed quality).

2.2.   Inertia: underlying behavioral mechanisms

Before analyzing the behavioral mechanisms underlying inertia among rational travelers, it is of much importance to clearly define what we mean by the term itself. Although it is tempting to define inertia in terms of a decision-maker repeatedly choosing the same alternative, the reason for this repetition may in fact be that the anticipated (expected) quality of an alternative is much higher than that of its competitors. As a result, defining inertia in terms of repetition alone is not very meaningful. Intuitively, we want to define inertia in a way that acknowledges that the mere *action of choosing* a particular alternative makes it more probable that the alternative is chosen again on the next day. In the context of the model presented above we can formalize this intuition as follows: a traveler exhibits inertia when the probability of choosing car over train during day $t+1$ is higher than the same probability during day t, under the condition that a) car is chosen during trip t and b) the level of car quality that was observed during the trip matches the expected level of quality. Inertia strength is defined in terms of the difference between these two choice probabilities. In notation (in the remainder of this paper, unless stated otherwise, we consider car-inertia):

---

[5] Note that since the early 1990s, Bayes' Theorem has been used extensively to model how travelers learn from making observations (e.g. Kaysi, 1991; Jha et al., 1998 ; Chen & Mahmassani, 2004; Sun et al., 2005, Arentze & Timmermans, 2005; Chancelier et al., 2007).

Definition 1a: A traveler exhibits inertia when $P\left(y^{t+1} = car \middle| y^t = car, \breve{x}^t_{car} = \hat{x}^t_{car} \right) > P\left(y^t = car \right)$ or, in expected utility terms, when $\left(EU^{t+1}_{car} \middle| y^t = car, \breve{x}^t_{car} = \hat{x}^t_{car} \right) - \left(EU^{t+1}_{train} \middle| y^t = car, \breve{x}^t_{car} = \hat{x}^t_{car} \right) > \left[EU^t_{car} - EU^t_{train} \right]$.

Definition 1b: Inertia strength equals $P\left(y^{t+1} = car \middle| y^t = car, \breve{x}^t_{car} = \hat{x}^t_{car} \right) - P\left(y^t = car \right)$. Alternatively, inertia strength, defined in terms of expected utilities, equals: $\left[\left(EU^{t+1}_{car} \middle| y^t = car, \breve{x}^t_{car} = \hat{x}^t_{car} \right) - \left(EU^{t+1}_{train} \middle| y^t = car, \breve{x}^t_{car} = \hat{x}^t_{car} \right) \right] - \left[EU^t_{car} - EU^t_{train} \right]$.

Note that by providing these definitions we contribute to the literature on traveler inertia, where inertia is mostly defined rather loosely in terms of repetitive behavior, or lack of willingness to switch to alternative routes and modes. Having defined inertia (strength), we can now explore when, to what extent and why the rational traveler presented in section 2.1 exhibits inertia. We do so by deriving two results.

Result 1: Under the condition that β_{VAR}, VAR^t_{car} and VAR_{obs} are all strictly positive, the rational traveler presented in section 2.1 exhibits inertia. Inertia strength – defined in terms of expected utilities – equals $\beta_{VAR} \cdot \left(VAR^t_{car} \right)^2 \middle/ \left(VAR^t_{car} + VAR_{obs} \right)$.

See the Appendix for a derivation of this result. The intuition behind the first part of this result is as stated in the introduction: the rational traveler learns about risky quality by observing the quality of a chosen travel mode. Given risk aversion this implies that choosing a travel mode, and observing that quality is as expected, leads to less quality uncertainty (and: higher utility) during the next trip: inertia arises as a cognitive lock-in effect. The intuition behind the second part of result 1 is as follows: it is obvious that a higher degree of risk aversion (β_{VAR}) and increased uncertainty (VAR^t_{car}) will both lead to more pronounced lock-in effect. The effect of observation unreliability (VAR_{obs}) on inertia strength can be explained as follows: remember that the key driver of inertia among rational travelers is the opportunity to learn from observations (in combination with risk aversion). Now, higher values of VAR_{obs} imply that observed quality during trip t is perceived by the traveler as a relatively unreliable signal of actual quality, which in turn implies that less weight is attached to these observations in the traveler's Bayesian learning process. As a result, to the extent that the traveler perceives observations to be unreliable signals, there is little opportunity for learning and the cognitive lock-in effect causing inertia is suppressed.

Result 2: Under the condition a) that β_{VAR}, VAR^t_{car} and VAR_{obs} are all strictly positive, and b) that the observed level of quality is Δ units lower than expected, a choice for the car mode during day t implies an increased choice probability for the car mode during day $t+1$ as long as $\beta_{VAR} > \dfrac{\beta_x \cdot \Delta}{VAR^t_{car}}$.

6

See the Appendix for a derivation of this result. The intuition behind it is as follows: even when quality is lower than expected, the fact that quality uncertainty is reduced by learning from observations may imply a net gain in utility. Higher levels of initial uncertainty and of risk aversion imply higher gains. However, when the disappointment in terms of quality – or the traveler's marginal valuation of quality – are too large, the net gain becomes a net loss and the probability that the traveler chooses the car-mode is lower during trip t+1 than it was during trip t.

3. Inertia among forward-looking travelers

3.1. A model of (myopic) forward-looking travel mode-choice behavior
Until now, we have assumed that the expected utility (*EU*) associated with choosing a travel mode-alternative is a function of anticipated mean and variance during the current trip t alone (hence: instantaneous Expected Utility). We now present a formulation of forward-looking travel choice-behavior that acknowledges that (the traveler anticipates that) choosing a particular mode during the current trip has consequences (because of *anticipated* learning dynamics) for the anticipated instantaneous Expected Utility that may be derived from the next trip. In short: we assume that a traveler, when planning trip t, knows that the observation of the chosen mode's quality during that trip may help him derive more utility from his mode choice during trip t+1. We choose to adopt this myopic (i.e. focused on the next day only) perspective instead of the perspective of a traveler that is infinitely forward-looking to keep the tractability of the model at a reasonable level. Note that the myopically forward-looking perspective is often used in travel behavior research (e.g. Arentze & Timmermans, 2005; Chorus et al., 2006b) as well as in choice models used in adjacent fields (e.g. Gabaix et al., 2006; Chorus & Timmermans, 2008), and can be interpreted as an extreme case of hyperbolic discounting (Laibson, 1997) as well.

The myopically forward-looking traveler maximizes (linearized) Expected Utility, where a travel mode's Expected Utility is the sum of the instantaneous Expected Utility defined in (1), and the product of a forward-looking parameter γ and the anticipated instantaneous Expected Utility associated with the next trip (which in turn is conditional on having chosen the considered travel mode during the current trip):

$$EU_{car}^t = E\tilde{U}_{car}^t + \gamma \cdot E\tilde{U}^{t+1} \big| \big(y^t = car \big)$$

(3).

$$EU_{train}^t = E\tilde{U}_{train}^t + \gamma \cdot E\tilde{U}^{t+1} \big| \big(y^t = train \big)$$

When γ approaches zero, the traveler is only concerned with the current trip, and (3) reduces to (1). When γ approaches one, the traveler is concerned with making a good choice during the next trip as much as he is concerned with the current trip's utility. The anticipated instantaneous Expected Utility associated with the next trip, conditional on having chosen a particular travel mode during the current trip, is denoted as:

7

$$EU^{t+1} \Big| \left(y^t = car \right) = \int\limits_{\breve{x}^t_{car}} \left[\max \left\{ E\tilde{U}^{t+1}_{car} \Big| \left(\breve{x}^t_{car} \right), E\tilde{U}^t_{train} \right\} \cdot f \left(\breve{x}^t_{car} \right) \right] d\breve{x}^t_{car}$$

$$(4).$$

$$EU^{t+1} \Big| \left(y^t = train \right) = \int\limits_{\breve{x}^t_{train}} \left[\max \left\{ E\tilde{U}^t_{car}, E\tilde{U}^{t+1}_{train} \Big| \left(\breve{x}^t_{train} \right) \right\} \cdot f \left(\breve{x}^t_{train} \right) \right] d\breve{x}^t_{train}$$

Here, $E\tilde{U}^{t+1}_{car} \big| \left(\breve{x}^t_{car} \right) = \beta_x \cdot \hat{x}^{t+1}_{car} - \beta_{VAR} \cdot VAR^{t+1}_{car}$ and $E\tilde{U}^{t+1}_{train} \big| \left(\breve{x}^t_{train} \right) = \beta_x \cdot \hat{x}^{t+1}_{train} - \beta_{VAR} \cdot VAR^{t+1}_{train}$,
where updated perceptions of quality (uncertainty) are as defined in (2). In words, the
traveler knows that when choosing a travel mode (e.g. the car mode) during trip t, he will
make a noisy observation of its quality $\breve{x}^t_{car}$ – of course, he does not know what quality
level he will observe, hence the integration over $f \left(\breve{x}^t_{car} \right)$. He also knows that he will use
this noisy observation to update his beliefs about car quality, and that he will base his
choice between the car and train mode during the next trip on these updated beliefs. The
traveler's beliefs regarding what level of quality he will observe during trip t equal his
initial beliefs about car quality, that is: $f \left(\breve{x}^t_{car} \right) = f \left(x^t_{car} \right)$.

Finally, note that the right-hand-side of (4) implies that it is assumed that a travel
mode's anticipated instantaneous Expected Utility during trip $t+1$, conditional on choosing
the other mode during day t, equals the instantaneous Expected Utility during trip t. This
assumption follows from the notion that travelers are assumed to only learn about a mode's
quality by means of direct observation[6]. Choice probabilities are given by using the binary
logit model presented in section 2.1. Upon choosing one of the travel modes, a noisy
observation is made of the quality of the chosen mode, leading to updated beliefs regarding
the chosen mode's quality.

3.2.   *Inertia among myopically forward-looking travelers (simulation)*
Because the integrals over $f \left(\breve{x}^t_{car} \right)$ and $f \left(\breve{x}^t_{train} \right)$ do not have a closed form solution we
discuss inertia among forward-looking travelers by means of a numerical simulation.
Importantly, because shown numerical simulation outcomes partly depend on arbitrarily
chosen values for relevant variables, we will be as careful and conservative as possible
when interpreting obtained simulation results. Assume the following settings:
$\hat{x}^t_{car} = \hat{x}^t_{train} = 0$, that is: expected quality of both the car and the train mode, as anticipated
when planning trip t, equals zero. Furthermore, $VAR^t_{car} = VAR^t_{train} = 5$, implying that during
trip t, both car and train quality are anticipated to fall, with a probability of 67% (95%),
within the interval [-5,5] ([-10,10]). Quality is evaluated in terms of one util per unit
($\beta_x = 1$). Finally, $VAR_{obs} = 2.5$. Given these settings, the following result is obtained (note
that a series of sensitivity analyses show that the result is robust with respect to varying one
or more of the above parameter settings):

---

[6] The next section shows how this assumption may be relaxed to acknowledge the role of secondary learning
from travel information.

Result 3a: For small magnitudes of β_{VAR}, higher values of γ imply lower levels of inertia strength, potentially leading to negative inertia.

Result 3b: For higher values of β_{VAR}, inertia strength increases, and the effect of γ on inertia vanishes.

Figure 1 illustrates these results (β_{VAR} is varied from 0 to 1 (this latter extreme being the value of β_x), while simultaneously γ is also varied from 0 to 1). The dependent variable, inertia strength, is measured (in line with definition 1) in terms of the probability that car is chosen during trip t+1, given a choice for the car during trip t and an associated observation of car quality that exactly matches expectations, minus the probability of choosing the car during trip t (which equals 50%). Integration over the density functions that represent what observations the traveler expects to make when choosing a particular mode is performed by means of Monte Carlo simulation (100 pseudo-random draws are made from each density function – sensitivity analysis showed that this number is sufficiently high). Software package GAUSS 7.0 is used for performing the simulation.

It is immediately seen that, as stated in result 3, the negative effect of γ on inertia strength only becomes noticeable when the degree of risk aversion approaches zero. For these low levels of risk aversion, positive values of γ may lead to negative levels of inertia strength (a choice for the car mode during trip t leads to a reduction in probability of choosing the car mode again during the next trip). However, when risk aversion grows, the negative effect of γ on inertia strength rapidly decreases.

The intuition behind these results is as follows: in the absence of (substantial levels of) risk aversion, the traveler is predominantly concerned with choosing the mode with highest expected quality. Given the Bayesian learning process defined in section 2.1, the observation of car quality during trip t results in an improved estimate of expected car quality (although expected quality itself is unchanged, given that $\breve{x}_{car}^t = \hat{x}_{car}^t$). As a result, when planning trip t+1, the traveler is still indifferent between the two modes in terms of expected quality, but he does face a choice between on the one hand a further decrease in car quality uncertainty (implied by a choice for the car mode), and on the other hand a decrease in train quality uncertainty (implied by a choice for the train mode). Intuition, as well as Bayes' Theorem, state that the latter decrease in uncertainty will be larger than the former, because a second observation of an uncertain phenomenon provides less information than the first observation. As a result, when planning trip t+1, a traveler that has chosen the car mode during trip t knows that he should choose the train mode during trip t+1 if he wants to be as sure as possible that he will make the right choice, in terms of maximizing expected quality, during trip t+2.

In the presence of substantial levels of risk aversion, the argumentation changes. First, as was shown in section 2, higher levels of risk aversion lead to higher levels of inertia strength due to the learning-based cognitive lock-in effect. Second, the intuition behind the result that when risk aversion is non-negligible the effect of γ on inertia strength rapidly approaches zero, is as follows: a risk averse traveler, having chosen a particular travel mode during trip t, has developed a preference for this mode as a result of the learning-based cognitive lock-in effect. When planning trip t+1, the traveler anticipates

(to the extent that he is forward-looking) that a choice for the other mode will bring the uncertainty level of that mode to the same level as that of the mode chosen during trip *t*. However, he also knows that he is able to further decrease the uncertainty associated with the mode chosen during trip *t*, by again choosing this mode during trip *t*+1. To the extent that risk aversion is present, this additional decrease in uncertainty counterbalances the potential gains in terms of getting a better estimate of expected quality, resulting from choosing a different mode during trip *t*+1. As a result, in the presence of risk aversion, increasingly forward looking behavior does not lead to lower inertia strength.
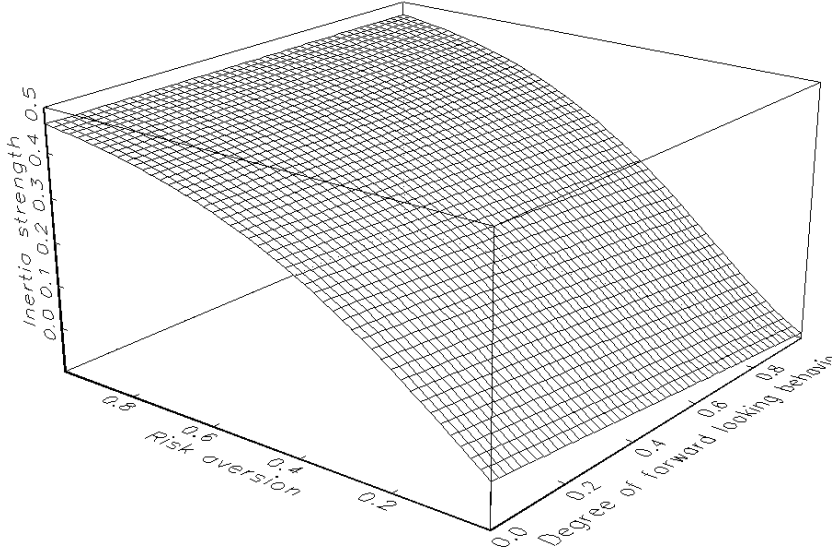

Figure 1: Inertia among forward looking travelers

Note that this interplay between the presence of forward-looking behavior and the presence of risk aversion can also be interpreted in terms of the combination of two competing perspectives on the value of information (the information resulting from making an observation (primary learning)). One perspective postulates that the value of information lies in its potential to help make better choices in future choice situations: information value being conceptualized as the difference between the anticipated expected utility of future choice situations with and without having received the information. A different perspective postulates that information value lies in its potential to reduce risk: information value is then conceptualized, in the context of our mode-choice example, as the difference between the expected utility associated with risky quality and the utility associated with expected quality. The former perspective implies that (anticipated) learning from observations leads to variety seeking, while the latter perspective implies that learning from observations leads to inertia. When combined, as is the case in this section, it is their mutual interplay that determines whether or not inertia prevails.

4.   Inertia in the presence of multimodal travel information

Until here, we have assumed that travel choices are made in the absence of information. Clearly, broadening this assumption towards acknowledging the presence of information

would lead to a better correspondence with most actual choice situations faced by travelers nowadays. Specifically, recent years have seen a great increase in personalized multi-modal information, referring to multiple quality-related attributes such as travel times and costs by car, and travel times, waiting times, costs and seat availability in Public Transport (e.g. Kenyon & Lyons, 2003; Chorus et al., In Press). In this section, we study how the presence of pre-trip personalized multimodal information about travel mode-quality impacts inertia emergence.

4.1.   A model of travel choice behavior in the presence of multimodal travel information
Assume, without loss of generality, that pre-trip information becomes available (or: noticed by the traveler) after the traveler has made the first trip t towards his new working location. Assume that he has chosen to use the car mode for this first trip. When planning trip $t+1$, the traveler receives multimodal quality information in the form of messages $\overline{x}_{car}^{t+1}$ and $\overline{x}_{train}^{t+1}$. As a result, his beliefs concerning the quality of both the car and the train mode (the former of these already updated once as a result of observed car quality during trip t) are updated.

Assume that the traveler believes that the information provider is unable to faultlessly assess quality (e.g. because the information provider is only partially able to correctly assess the traveler's tastes). Specifically, we postulate that $f\left(\overline{x}_{car}^{t+1}\middle|x_{car}\right)=N\left(x_{car},\sqrt{VAR_I}\right)$ and that $f\left(\overline{x}_{train}^{t+1}\middle|x_{train}\right)=N\left(x_{train},\sqrt{VAR_I}\right)$, VAR_I being a measure of anticipated information unreliability. This formulation implies that information is perceived to be unbiased (i.e.: the expected quality message equals expected quality). Note that we assume that there are no differences between car and train information in terms of anticipated unreliability. As a result, reception of information leads to the following updates in quality anticipations: $f\left(x_{car}^{t+1}\middle|\breve{x}_{car}^{t},\overline{x}_{car}^{t+1}\right)=N\left(\hat{x}_{car}^{t+1},\sqrt{VAR_{car}^{t+1}}\right)$ and $f\left(x_{train}^{t+1}\middle|\overline{x}_{train}^{t}\right)=N\left(\hat{x}_{train}^{t+1},\sqrt{VAR_{train}^{t+1}}\right)$, where:

$$\hat{x}_{car}^{t+1}=\frac{\left(VAR_{car}^{t+1}\right)^{-1}\cdot\hat{x}_{car}^{t+1}+\left(VAR_I\right)^{-1}\cdot\overline{x}_{car}^{t+1}}{\left(VAR_{car}^{t+1}\right)^{-1}+\left(VAR_I\right)^{-1}} \text{ and } VAR_{car}^{t+1}=\frac{VAR_{car}^{t+1}\cdot VAR_I}{VAR_{car}^{t+1}+VAR_I}$$

$$\text{(5)}.$$

$$\hat{x}_{train}^{t+1}=\frac{\left(VAR_{train}^{t}\right)^{-1}\cdot\hat{x}_{train}^{t}+\left(VAR_I\right)^{-1}\cdot\overline{x}_{train}^{t+1}}{\left(VAR_{train}^{t}\right)^{-1}+\left(VAR_I\right)^{-1}} \text{ and } VAR_{train}^{t+1}=\frac{VAR_{train}^{t}\cdot VAR_I}{VAR_{train}^{t}+VAR_I}$$

Here, $\hat{x}_{car}^{t+1}$ and VAR_{car}^{t+1} are as defined in (2). Note the difference in superscripts between the car and train-mode: for the car-mode, the updated quality perceptions are a combination of the received message and perceived quality after having chosen the car mode during trip t (which itself is a combination of initial anticipations and observed quality). For the train mode, the updated quality perception is simply a combination of the received message and initial quality anticipations.

For a non-forward looking traveler, the story ends here. However, to the extent that the traveler is concerned – when planning trip t – with the instantaneous utility to be

11

derived from trip $t+1$, the presence of information has a second order effect: the traveler, when planning trip t, anticipates that he will receive travel information again before trip $t+1$ and that this information will lead him to once again update his perceptions before planning trip $t+1$. Of course, he does not know beforehand what messages he will receive. However, he does know that there is a formal relationship between on the one hand the probability of receiving particular messages $\overline{x}_{car}^{t+1}$ and $\overline{x}_{train}^{t+1}$ when planning trip $t+1$ and on the other hand his anticipation of car and train quality after having made trip t, in combination with his anticipation of information reliability.

This relationship can be illustrated as follows: take for example the situation where quality equates travel time. Given a lack of experience, the traveler is uncertain about the travel time for a given trip, but he thinks it will probably be somewhere between 50 and 75 minutes. When acquiring travel time information, and given that he believes the information to be very reliable, the traveler will attach a very low probability to the occurrence of a message saying that the actual travel time is, say, 10 minutes or 280 minutes, but he will attach a far higher probability to the occurrence of messages from the 50-75 minutes interval. On the other hand, when the traveler believes that the information is very unreliable, he will indeed anticipate the reception of messages that saying that the travel time is (much) lower than 50 minutes, or (much) higher than 75 minutes.

As a result of this relationship, the traveler's anticipations of what message he might receive when planning trip $t+1$ differ between the two modes, since his anticipation of messages is conditional on the anticipated observation to be made during trip t. In notation, the traveler knows that, should he for example choose the car mode during trip t and make an observation $\breve{x}_{car}^{t}$, his anticipations of what messages $\overline{x}_{car}^{t+1}$ and $\overline{x}_{train}^{t+1}$ he will receive when planning trip $t+1$ are:

$$f\left(\overline{x}_{car}^{t+1}\middle|\breve{x}_{car}^{t}\right)=\int\limits_{x_{car}^{t+1}}\left[f\left(\overline{x}_{car}^{t+1}\middle|x_{car}^{t+1}\right)\cdot f\left(x_{car}^{t+1}\middle|\breve{x}_{car}^{t}\right)\right]dx_{car}^{t+1}$$

$$(6).$$

$$f\left(\overline{x}_{train}^{t+1}\middle|\breve{x}_{car}^{t}\right)=\int\limits_{x_{train}^{t}}\left[f\left(\overline{x}_{train}^{t+1}\middle|x_{train}^{t}\right)\cdot f\left(x_{train}^{t}\right)\right]dx_{train}^{t}$$

Here, $f\left(x_{car}^{t+1}\middle|\breve{x}_{car}^{t}\right)$ is as presented in (2). Furthermore, in line with the argument presented right before equation (5), $f\left(\overline{x}_{car}^{t+1}\middle|x_{car}^{t+1}\right)=N\left(x_{car}^{t+1},\sqrt{VAR_I}\right)$ and $f\left(\overline{x}_{train}^{t+1}\middle|x_{train}^{t}\right)=N\left(x_{train}^{t},\sqrt{VAR_I}\right)$. Finally, note that $f\left(\overline{x}_{car}^{t+1}\middle|\breve{x}_{train}^{t}\right)$ and $f\left(\overline{x}_{train}^{t+1}\middle|\breve{x}_{train}^{t}\right)$ are derived in the same fashion as $f\left(\overline{x}_{car}^{t+1}\middle|\breve{x}_{car}^{t}\right)$ and $f\left(\overline{x}_{train}^{t+1}\middle|\breve{x}_{car}^{t}\right)$.

Given these anticipated probabilities of receiving particular messages, we can derive the forward-looking traveler's anticipation of instantaneous Expected Utility associated trip $t+1$, conditional on having made a particular observation during trip t:

$$EU\tilde{}^{t+1}\big|\big(\breve{x}_{car}^{t}\big) = \int_{\overline{x}_{car}^{t+1}} \int_{\overline{x}_{train}^{t+1}} \left[\max\begin{Bmatrix} EU\tilde{}_{car}^{t+1}\big|\big(\overline{x}_{car}^{t+1},\breve{x}_{car}^{t}\big), \\ EU\tilde{}_{train}^{t+1}\big|\big(\overline{x}_{train}^{t+1}\big) \end{Bmatrix} f\big(\overline{x}_{car}^{t+1}\big|\breve{x}_{car}^{t}\big) f\big(\overline{x}_{train}^{t+1}\big|\breve{x}_{car}^{t}\big) \right] d\overline{x}_{car}^{t+1} d\overline{x}_{train}^{t+1}$$

$$EU\tilde{}^{t+1}\big|\big(\breve{x}_{train}^{t}\big) = \int_{\overline{x}_{car}^{t+1}} \int_{\overline{x}_{train}^{t+1}} \left[\max\begin{Bmatrix} EU\tilde{}_{car}^{t+1}\big|\big(\overline{x}_{car}^{t+1}\big), \\ EU\tilde{}_{train}^{t+1}\big|\big(\overline{x}_{train}^{t+1},\breve{x}_{train}^{t}\big) \end{Bmatrix} f\big(\overline{x}_{car}^{t+1}\big|\breve{x}_{train}^{t}\big) f\big(\overline{x}_{train}^{t+1}\big|\breve{x}_{train}^{t}\big) \right] d\overline{x}_{car}^{t+1} d\overline{x}_{train}^{t+1}$$

(7).

Conditional instantaneous Expected Utilities are derived by applying the updating rule given in (5) and entering updated perceptions in (1). Substituting (7) in (4) and subsequently substituting (4) and (1) in (3) gives the Expected Utilities for the two modes. Based on these Expected Utilities, choice probabilities are derived using the binary logit model presented in section 2.1. Upon choosing one of the travel modes, based on these anticipations of quality observations during trip t and resulting messages to be received when planning tip $t+1$, a noisy observation is made of the quality of the chosen mode, leading to updated beliefs regarding the chosen mode's quality.

4.2.   Inertia in the presence of multimodal travel information (simulation)
Because the derivation of choice probabilities among forward-looking travelers in the presence of information involves the evaluation of several integrals without a closed form solution, we study inertia by mans of numerical simulation. Assume the following settings: like in section 3.2, $\hat{x}_{car}^{t} = \hat{x}_{train}^{t} = 0$, $VAR_{car}^{t} = VAR_{train}^{t} = 5$, $VAR_{obs} = 2.5$ and $\beta_{x}=1$. Risk aversion parameter $\beta_{VAR} = 0.25$. Given these settings, the following results are obtained (note that a series of sensitivity analyses show that these results are robust[7] with respect to varying one or more of the above parameter settings):

Result 4a: Inertia strength is an increasing function of information unreliability (VAR_{I}). For large magnitudes of VAR_{I}, higher values of γ imply lower levels of inertia strength.

Result 4b: For small values of VAR_{I} (implying reliable information), the effect of γ on inertia strength vanishes.

Figure 2 illustrates this result. VAR_{I} is varied from 0 to 5 (this latter extreme being the value of VAR_{car}^{t} and VAR_{train}^{t} which implies that the information is anticipated to be equally unreliable as the traveler's own initial knowledge. Simultaneously, γ is varied from 0 to 1. The dependent variable, inertia strength, is measured (in line with definition 1) in terms of the probability that car is chosen during trip $t+1$ (given a choice for the car mode during trip

---

[7] An exception is the level of risk aversion: as was established in result 3b, high levels of risk aversion imply that the effect of γ becomes negligible, irrespective of the value of VAR_{I}.

t and an associated observation of car quality that exactly matches a priori expectations) minus the probability of choosing the car during trip *t* (which equals 50%). Integration over the density functions that represent what observations the traveler expects to make when choosing a particular mode is performed by means of Monte Carlo simulation (100 pseudo-random draws for each density). Messages are also simulated by making 100 pseudo-random draws from the pdfs presented in (6). Crucially, to reflect that the traveler anticipates that received messages are conditional on his (potentially updated) perceptions of quality after having made a trip (see (6)), messages are conditioned on anticipated quality levels. That is, for each possible quality level drawn, 100 messages are drawn, which results in 10,000 messages being drawn in total. Sensitivity analyses showed that these numbers were sufficiently high. Figure 2 shows simulated levels of inertia strength (note that the scale of the Z-axis differs from that in Figure 1).
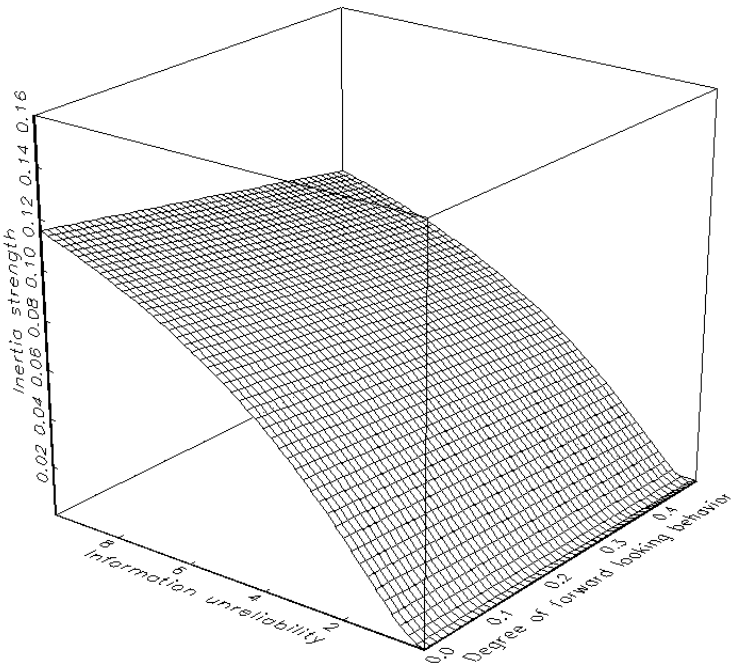

Figure 2: Inertia in the presence of multimodal travel information

The positive effect of information unreliability on inertia strength is clearly visible. The intuition behind this result is straightforward: as discussed in section 2.2, inertia arises from a lock-in effect based on travelers' ability to learn from observing a chosen mode's quality level. However, these observations becomes less important when information becomes more and more reliable, and as a result the cognitive lock-in effect causing inertia becomes less pronounced. Take the extreme situation where a traveler believes that he will receive fully reliable information at the start of each trip: in that case, paying attention to quality levels observed during trips made is of no use as it does not add to the learning process. As a result, the cognitive lock-in effect based on learning from observations vanishes.

When information is unreliable, the negative effect of γ on inertia strength is clearly visible as well. This effect is in line with result 3a as illustrated in Figure 1 (note that to the extent that information is anticipated to be unreliable, the 'with information' case

becomes equivalent to the 'without information' case): forward-looking travelers are relatively prone to explore new alternatives, because they know that observing their quality levels helps them achieve higher levels of expected quality in future choices. To the extent that information is anticipated to be reliable, the effect of γ on inertia strength rapidly diminishes because reliable information not only diminishes the usefulness of observing quality during the current trip, but also the usefulness of anticipated observations during the next trip (the traveler knows that information will also be available during the next trip).

Two final remarks are in place here. First, it is obvious that any impact of multimodal travel information on car inertia is conditional on car-drivers' (passive or active) acquisition of this information. In this light, it should be noted that information acquisition is not costless: even when the information is provided for free, empirical research shows that the non-monetary costs associated with acquiring the information (in terms of effort, attention, time) are generally perceived as high, leading to relatively low usage levels (e.g. Chorus et al., In Press). As shown by Aarts et al. (1997), this holds especially for the situation where travelers are provided with information about another than their usual alternative. In sum, these studies show that the findings presented in this section provide an upper bound for the actual impact of the provision of multimodal information on car inertia.

Second, as pointed out by a referee, it has been shown empirically that the provision of travel information to travelers may impact their level of risk aversion. Depending on the nature of the information (whether or not it is dynamic, and depending on whether it informs travelers about one or more of the moments of the distribution of quality (or travel times), travelers may become more or less risk averse (Avineri & Prashker, 2006; Ben-Elia et al., 2008; Bogers, 2009). This cause-effect relationship is not studied in this paper, but may be addressed in further (empirical) research.

5. A case study: Impact of a forced highway closure on car inertia

Whereas the previous three sections focused on the reasons underlying inertia among rational travelers, this section uses the developed model of rational inertia to study the impact – in terms of reducing car inertia – of a forced highway closure. The aim of this analysis not to obtain new results (as is discussed below, the effects of highway-closures on car inertia have been investigated in-depth using empirical data), but to show how these empirical results can be explained by our model of inertia emergence among rational travelers.

It is empirically well established, and in line with intuition, that a forced change of travel mode during a brief period of time has the potential to reduce inertia strength (Fujii et al., 2001; Fujii & Gärling, 2003; Fujii & Kitamura, 2003). Specifically, Fujii and coauthors found that car drivers, after being forced to abandon their usual car-route during an eight day highway closure exhibited less inertia with respect to the car mode than they did before the closure. Public transport use after the closure (immediately after as well as one year after the closure) was highest among those car-drivers that either used the car mode relatively infrequently beforehand (Fujii et al., 2001) or were offered free bus tickets (Fujii & Kitamura, 2003). The cited studies have in common that they interpret inertia as resulting from habitual behavior ("goal-directed automaticity"), among car-drivers that are "unlikely

to think of public transport as an alternative". In contrast, our model of traveler behavior postulates that travelers – irrespective of their inertia level – always consider both the car and the train mode when making a travel choice. In an attempt to study whether our model of rational inertia can explain the findings of Fujii and coauthors, we perform a numerical simulation. We assume the absence of pre-trip multimodal quality-related information.

Assume the following settings: $\hat{x}_{car}^t = \hat{x}_{train}^t = 0$, that is: expected quality of both the car and the train mode, as anticipated when planning trip t, equals zero. Furthermore, $VAR_{train}^t = 5$. Quality is evaluated in terms of one util per unit ($\beta_x = 1$), as is quality uncertainty ($\beta_{VAR} = 1$). Finally, $VAR_{obs} = 2.5$ and $\gamma = 0$.[8] The dependent variable is the probability that car is chosen during trip $t+1$, given a (forced) choice for the train during trip t. We simultaneously vary anticipated uncertainty concerning car quality under normal conditions (VAR_{car}^t, between 2.5 and 5) and the observed train quality during the forced execution of the train alternative during day t ($\breve{x}_{train}^t$, between -10 and 10). Note that values of VAR_{car}^t close to 2.5 represent the situation where repeated car use in the past has led to a relatively large reduction in uncertainty. In contrast, values of VAR_{car}^t close to 5 (= VAR_{train}^t) imply that no learning has taken place yet. In sum, low (high) values of VAR_{car}^t reflect the presence (absence) of car inertia before occurrence of the highway closure. Higher values of $\breve{x}_{train}^t$ represent situations similar to the situation described in Fujii & Kitamura (2003) where car-drivers were offered free bus tickets. Note that because $\hat{x}_{car}^t = \hat{x}_{train}^t = 0$, any preference for the car mode before the highway closure results solely from the relatively small uncertainty surrounding its quality. Figure 3 shows how the probability of choosing the car mode after the forced one-day train choice depends on car uncertainty before the highway closure (being a proxy for inertia strength) and on train quality observed during the forced train choice[9].

Shown relations are fully in line with intuition and the empirical results obtained by Fujii and co-workers: a) lower levels of car uncertainty (being a result from inert behavior in the past) are associated with higher probabilities of returning to the car mode after the forced change; b) higher levels of observed train quality lead to lower probabilities of returning to the car mode after the forced change. As a further illustration, Figure 4 plots the difference between the probability of choosing the car mode before and after the highway closure, again as a function of car uncertainty before the highway closure and train quality observed during the highway closure. It becomes clear that when the observed train quality is much lower than expected, the probability of choosing the car mode after the highway closure will in fact be higher than the corresponding probability before the closure (this is reflected in negative values of the dependent variable[10]). This is especially the case among travelers who were relatively uncertain about car quality before the highway-closure.

---

[8] Note that sensitivity analysis regarding the value of γ showed that variation between zero and one did not affect the overall result obtained from the simulation – which is in line with result 3b.

[9] Again, integration over the density functions is performed by means of Monte Carlo simulation (involving 100 pseudo-random draws for each density function).

[10] Note that we do not go into the precise magnitude of this difference, since it depends on simulation settings.
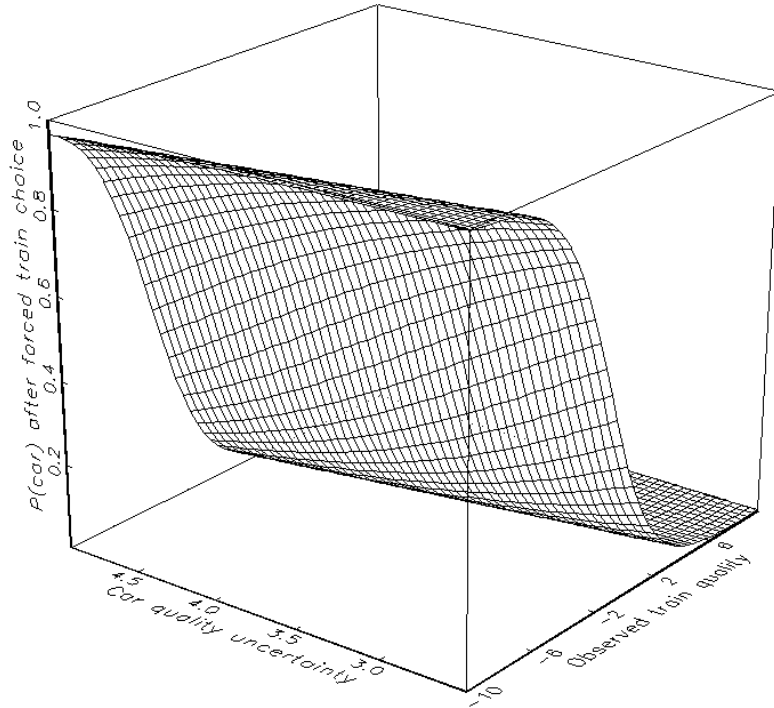
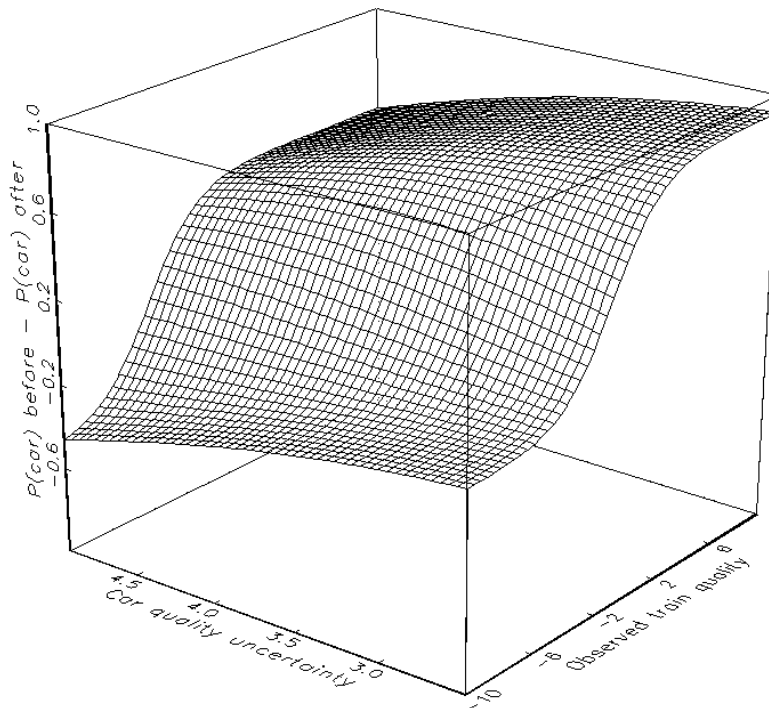**Figure 3: Car probabilities after a one-day highway closure**

Figure 4: Difference in car probabilities before and after a one-day highway closure

Again, these results are fully in line with intuition as well as empirical literature – for example, it is well established that when drivers divert from their usual route or mode,

based on positive information (expectations) about an alternative, a subsequent negative experience heavily affects their propensity to divert again in the future (e.g. Bonsall et al., 2004; Chen et al., 1999; Srinivasan and Mahmassani, 2000; Jou et al. 2005).

6.　　Conclusions and discussion

This paper shows how travelers that make rational choices exhibit inertia when faced with risky mode choice-situations, due to a learning-based lock-in effect. We also show how the interplay between forward-looking behavior and risk aversion determines the level of inertia. In addition, we show how the costless provision of multimodal travel information may slow down inertia emergence, but only to the extent that the information is considered reliable. As a case study, we show how our model predicts that a forced highway closure may help break inertia to the extent that an alternative train alternative performs satisfactory and car inertia strength is limited – a result that is fully in line with empirical literature.

The main message that this paper tries to convey is that the postulates of bounded rationality (sticking to a good enough alternative that has been chosen before arises from the wish to save effort, time and attention) are not necessary to explain inertia in a travel mode choice context. We show that inertia emerges rapidly and forcefully among travelers that live up to high standards of 'unbounded' rationality: they consider all alternatives available in their choice set, evaluate by means of Expected Utility maximization and learn from past observations and information in a Bayesian manner. Of course, there is no reason to conclude from the theoretical analysis presented here that inertia has little to do with bounded rationality – on the contrary: we acknowledge the intuitive and empirically well established notion of inertia arising from travelers' wish to economize on cognitive resources. In fact, we believe that actual patterns of traveler inertia are the result of an interplay between boundedly and 'unboundedly' rational behavioral phenomena.

From a theoretical perspective, we believe that a particularly fruitful direction for further research would be to incorporate the notion that quality uncertainty may not only result from a lack of experience, but also from day-to-day variability. While the assumption that quality uncertainty solely arises from a lack of experience is computationally efficient (see for example Kaysi (1991), Jha et al. (1998) and Chen & Mahmassani (2004) for recent examples), it actually implies that making repeated observations will in the long run eliminate uncertainty. However, when day-to-day variability is present this implication is not realistic, as for example Jha et al. (1998) acknowledge. Instead, in that case repeated observations enable a traveler to learn actual day-to-day variability in quality, which should be reflected in an additional quality-variance term that gradually approaches actual quality-variance. We consider the treatment of day-to-day variability impacts on inertia emergence as an important avenue for further research, although the increased behavioral realism associated with incorporating day-to-day quality variability is likely to come at the price of substantial increases in model intractability.

Notwithstanding that the presented model leaves room for theoretical model extensions like the one discussed directly above, we consider empirical testing to be the paramount direction for further research. In terms of data collection, basically two alternative courses of action are possible: one is to observe choices in a hypothetical choice experiment, preferably using a carefully designed incentive structure to control preferences

– see Denant-Boèmont & Petiot (2003) and Han et al. (2007) for recent examples in a travel behavior context. The other one is to collect panel data sets of choice behavior observed in real life, ideally enriched with data concerning travelers' anticipations of uncertainty, information reliability and so on. The (dis-)advantages of both approaches are obvious, and a trade-off between efficiency (in terms of the needed data-collection effort) and validity must be made. In terms of model estimation, the identification of inertia from repeated choices is not trivial, given endogeneity issues and the fact that it is by definition difficult to distinguish whether a repeated choice for a particular alternative follows from its perceived superiority or from some form of inertia. However, recent developments in discrete choice theory will help enable the needed analyses (Cantillo et al., 2007). Hopefully, such empirical analyses will in time highlight the role of bounded and unbounded rationality in explaining and predicting travelers' repeated choice-behavior.

To conclude, we believe that the results and overall conclusion we present here have an important practical implication: they suggest that inertia might be even more difficult to 'break' then we thought. Whereas transport policy-makers often implicitly assume that helping make travelers choose in a more 'rational' way (e.g. by providing travel information; see Chorus et al. (2006a) for an overview of such attempts) will reduce inertia, our results suggest that such an approach will not be very helpful. On the other hand, our analyses do highlight the potential of other measures that have been proposed in the past, such as the promotion of alternative (sustainable) travel modes when travelers are forced to abandon their usual mode. Our model of rational behavioral inertia suggests that travelers may be quite good in quickly learning 'good' habits.

References

Aarts, H., Verplanken, B., van Knippenberg, A. , 1997. Habit and information use in travel mode choices. *Acta Psychologica*, 96, pp. 1-14

Allen, W., B., Hamhmoud, M.M., McNeil, D., 1985. The importance of time in transit and reliability of transit time for shippers, receivers and carriers. *Transportation Research Part B*, 19(5), pp. 447-456

Arentze, T.A., Timmermans, H.J.P., 2005. Information gain, novelty seeking and travel: A model of dynamic activity-travel behavior under conditions of uncertainty. *Transportation Research Part A*, 39(2-3), pp. 125-145

Avineri, E., Prashker, J.N., 2006. The impact of travel information on travelers' learning under uncertainty. *Transportation*, 33(4), pp. 393-408

Ben-Elia, E., Erev, I., Shiftan, Y., 2008. The combined effect of information and experience on drivers' route-choice behavior. *Transportation*, 35(2), pp. 165-177

Bogers, E.A.I., Viti, F., Hoogendoorn, S.P., 2005. Joint modelling of ATIS, habit and learning impacts on route choice by laboratory simulator experiments. *Transportation Research Records*, 1926, pp. 189-197

Bogers, E.A.I., 2009. *Traffic information and learning in day-to-day route choice*. TRAIL PhD-Thesis Series, Delft, The Netherlands

Bonsall, P., Firmin, P., Beale, J., 2004. Perception of modal attributes: how accurate and how susceptible to change? Paper presented at the *83rd meeting of the Transportation Research Board*, Washington, D.C.

Cantillo, V., Ortuzar, J. de D., Williams, H.C.W.L., 2007. Modeling discrete choices in the presence of inertia and serial correlation. *Transportation Science*, 41(2), pp. 195-205

Chancelier, J-P., De Lara, M., de Palma, A., 2007. Risk aversion, road choice, and the one-armed bandit problem. *Transportation Science*, 41(1), pp. 1-14

Chen, P.S-T., Srinivasan, K.K., Mahmassani, H.S., 1999. Effect of information quality on compliance behavior on commuters under real-time traffic information. *Transportation Research Record*, 1676, pp. 53-60

Chen, R.B., Mahmassani, H.S., 2004. Travel time perception and learning mechanisms in traffic networks. *Transportation Research Record*, 1894, pp. 209-221

Chorus, C.G., Molin, E.J.E., van Wee, G.P., 2006a. Travel information as an instrument to change car-drivers' travel choices: A literature review. *European Journal of Transport and Infrastructure Research*, 6(4), 335-364

Chorus, C.G., Arentze, T.A., Molin, E.J.E., Timmermans, H.J.P., van Wee, G.P., 2006b. The value of travel information: Decision-strategy specific conceptualizations and numerical examples. *Transportation Research Part B*, 40(6), pp. 504-519

Chorus, C.G., Molin, E.J.E., Arentze, T.A., Timmermans, H.J.P., van Wee, G.P., 2007. Travelers' need for information in traffic and transit: Results from a web-survey. *Journal of Intelligent Transportation Systems*, 11(2), 1-11

Chorus, C.G., Timmermans, H.J.P., 2008. Revealing consumer preferences by observing information search. *Journal of Choice Modeling*, 1(1), 3-25

Chorus, C.G., Walker, J.L., Ben-Akiva, M.E., In Press. The value of travel information: A search-theoretic approach. *Journal of Intelligent Transportation Systems*

Commission of the European Communities, 2007. *Green Paper: Towards a new culture for urban mobility.* Brussels, Belgium

Denant-Boèmont, L., Petiot, R., 2003. Information value and sequential decision-making in a transport setting: an experimental study. *Transportation Research Part B*, 37, pp. 365-386

de Palma, A., Picard, N., 2006a. Equilibria and information provision in risky networks with risk-averse drivers. *Transportation Science*, 40(4), pp. 393-408

de Palma A., Picard, N., 2006b. Route choice behavior with risk averse users. In: *Spatial Evolution and Modelling*, P. Nijkamp and A. Reggiani (eds), Edward Elgar, pp. 139-178.

Department for Transport, 2004. *The future of transport: a network for 2030.* London, UK

Edwards, W., Lindman, H., Savage, L.J., 1963. Bayesian statistical research for psychological research. *Psychological Review*, 70, pp. 193-242

Erev, I., Barron, G., 2005. On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, 112(4), pp. 912-931

Eriksson, L., Garvill, J., Nordlund, A., 2008. Interrupting habitual car use: The importance of car habit strength and moral motivation for personal car use reduction. *Transportation Research Part F*, 11(1), pp. 10-23

Fujii, S., Gärling, T., Kitamura, R., 2001. Changes in drivers' perceptions and use of public transport during a freeway closure. *Environment and Behaviour*, 33, pp. 796-808

Fujii, S., Gärling, T., 2003. Development of script-based travel mode choice after forced change. *Transportation Research Part F*, 6, pp. 117-124

Fujii, S., Kitamura, R., 2003. What does a one-moth free bus ticket do to habitual drivers? *Transportation*, 30, pp. 81-95

Gabaix, X., Laibson, D., Moloche, G., Weinberg, S., 2006. Costly information acquisition : Experimental analysis of a boundedly rational model. *American Economic Review*, 96, 4, pp. 1043-1068

Gardner, B., 2009. Modelling motivation and habit in stable travel mode contexts. *Transportation Research Part F*, 12(1), pp. 68-76

Gärling, T., Eek, D., Loukopoulos, P., Fujii, S., Johansson-Stenman, O., Kitamura, R., Pendyala, R., Vilhelmson, B., 2002. A conceptual analysis of the impact of travel demand management on private car-use. *Transport Policy*, 9, pp. 59-70

Gärling, T., Axhausen, K.W., 2003. Introduction: Habitual travel choice. *Transportation*, 30, pp. 1-11

Han, Q., Dellaert, B.G.C., van Raaij, W.F., Timmermans, H.J.P., 2007. Modeling strategic behavior in anticipation of congestion. *Transportmetrica*, 3(2), pp. 119-138

Harker, P.T., Hong, S., 1994. Pricing of track time in railroad operations: an internal market approach. *Transportation Research Part B*, 28(3), pp. 197-212

Hey, J.D., 1995. Experimental investigations of errors in decision making under risk. *European Economic Review*, 39, pp. 633-640

Jha, M., Madanat, S., Peeta, S., 1998. Perception updating and day-to-day travel choice dynamics in traffic networks with information provision. *Transportation Research Part C*, 6, pp. 189-212

Jou, R., Lam, S., Liu, Y., Chen, K., 2005. Route switching behavior on freeways with the provision of different types of real-time traffic information. *Transportation Research Part A*, 39(5), 445-461

Kaysi, I., 1991. *Framework and models for provision of driver information system*. PhD-thesis, Massachusetts Institute of Technology, Cambridge, MA

Kenyon, S., Lyons, G., 2003. The value of integrated multimodal traveler information and its potential contribution to modal change. *Transportation Research Part F*, 6(1), pp. 1-21

Laibson, D., 1997. Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112, pp. 443-477

Lam, T., and Small, K., 2001. The value of time and reliability: measurement from a value pricing experiment. *Transportation Research Part E*, 37(4), pp. 297-320.

Liu, H.X., He, X., Recker, W., 2007. Estimation of the time-dependency of values of travel time and its reliability from loop detector data. *Transportation Research Part B*, 41(4), pp. 448-461

Mahmassani, H.S., Chang, G., 1987. On boundedly rational user equilibrium in transportation systems. *Transportation Science*, 22, 89-99

Manski, C.F., 1977. The structure of random utility models. *Theory and Decision*, 8, pp. 229-254

March, J.G., 1988. Bounded rationality, ambiguity and the engineering of choice. *Bell Journal of Economics*, 9(2), pp. 587-608

Ministry of Transport, Public Works and Water Management, 2002. *De markt voor multi-modaal personenvervoer.* Rijkswaterstaat, Rotterdam, The Netherlands (in Dutch)

Payne, J.W., Bettman, J.R., Johnson, E.J., 1993. *The adaptive decision maker.* Cambridge University Press, Cambridge

Shugan, S.M., 1980. The cost of thinking. *Journal of consumer Research*, 7, pp. 99-111

Simon, H.A., 1978. Rationality as process and as product of thought. *American Economic Review*, 68(2), pp. 1-16

Srinivasan, K.K., Mahmassani, H.S., 2000. Modeling inertia and compliance mechanisms in route choice behavior under real-time information. *Transportation Research Record*, 1725, pp. 45-53

Sun, Z., Arentze, T.A., Timmermans, H.J.P., 2005. Modeling the impact of travel information on activity-travel rescheduling decisions under conditions of travel time uncertainty. *Transportation Research Record*, 1926, pp. 79-87

Appendix: derivations of results 1 and 2

Derivation of result 1.
By definition, the rational traveler exhibits inertia when $\left[\left(E\tilde{U}_{car}^{t+1}\middle|y^t=car,\breve{x}_{car}^t=\hat{x}_{car}^t\right)-\left(E\tilde{U}_{train}^{t+1}\middle|y^t=car,\breve{x}_{car}^t=\hat{x}_{car}^t\right)\right] > \left[E\tilde{U}_{car}^t-E\tilde{U}_{train}^t\right]$. Because anticipated train quality and train quality uncertainty are unaffected by a traveler's choice for the car-mode, $E\tilde{U}_{train}^t$ equals $\left(E\tilde{U}_{train}^{t+1}\middle|y^t=car,\breve{x}_{car}^t=\hat{x}_{car}^t\right)$. As a result, the traveler exhibits inertia when $E\tilde{U}_{car}^{t+1}\middle|\left(y^t=car,\breve{x}_{car}^t=\hat{x}_{car}^t\right) > E\tilde{U}_{car}^t$. Defining Expected Utilities in terms of a mean-variance linearization gives the following rewrite: $\beta_x \cdot \hat{x}_{car}^t - \beta_{VAR} \cdot \dfrac{VAR_{car}^t \cdot VAR_{obs}}{VAR_{car}^t + VAR_{obs}} > \beta_x \cdot \hat{x}_{car}^t - \beta_{VAR} \cdot VAR_{car}^t$. This inequality can be rearranged as follows: $\beta_{VAR} \cdot \dfrac{\left(VAR_{car}^t\right)^2}{VAR_{car}^t + VAR_{obs}} > 0$. Under the condition that β_{VAR}, VAR_{car}^t and VAR_{obs} are all strictly positive, this condition holds, and inertia is established. By definition, inertia strength under the condition that quality is as expected, in terms of expected utilities, equals

$\left[\left(EU_{car}^{t+1} \middle| y^t = car, \breve{x}_{car}^t = \hat{x}_{car}^t \right) - EU_{train}^t \right] - \left[EU_{car}^t - EU_{train}^t \right]$. It is directly seen that this

difference equals $\beta_{VAR} \cdot \dfrac{\left(VAR_{car}^t \right)^2}{VAR_{car}^t + VAR_{obs}}$.

Derivation of result 2.
When the observed level of quality is Δ units lower than expected, and given iid errors, a choice for the car mode during day t implies an increased choice probability for the car during day $t+1$ when $\left[\left(E\tilde{U}_{car}^{t+1} \middle| y^t = car, \breve{x}_{car}^t = \hat{x}_{car}^t - \Delta \right) - \left(E\tilde{U}_{train}^{t+1} \middle| y^t = car, \breve{x}_{car}^t = \hat{x}_{car}^t - \Delta \right) \right] >$

$\left[E\tilde{U}_{car}^t - E\tilde{U}_{train}^t \right]$. Because anticipated train quality and train quality uncertainty are unaffected by a traveler's choice for the car-mode, $E\tilde{U}_{train}^t$ equals $\left(E\tilde{U}_{train}^{t+1} \middle| y^t = car, \breve{x}_{car}^t = \hat{x}_{car}^t - \Delta \right)$. As a result, the above condition reduces to $E\tilde{U}_{car}^{t+1} \middle| \left(y^t = car, \breve{x}_{car}^t = \hat{x}_{car}^t - \Delta \right) > E\tilde{U}_{car}^t$. Defining Expected Utilities in terms of a mean-

variance linearization presented gives: $\beta_x \cdot \dfrac{\left(VAR_{car}^t \right)^{-1} \cdot \hat{x}_{car}^t + \left(VAR_{obs} \right)^{-1} \cdot \left[\hat{x}_{car}^t - \Delta \right]}{\left(VAR_{car}^t \right)^{-1} + \left(VAR_{obs} \right)^{-1}}$

$- \beta_{VAR} \cdot \dfrac{VAR_{car}^t \cdot VAR_{obs}}{VAR_{car}^t + VAR_{obs}} - \left(\beta_x \cdot \hat{x}_{car}^t - \beta_{VAR} \cdot VAR_{car}^t \right) \quad > \quad 0.$ Manipulating gives:

$\beta_{VAR} > \beta_x \cdot \Delta \cdot \dfrac{\left(VAR_{obs} \right)^{-1}}{\left(VAR_{car}^t \right)^{-1} + \left(VAR_{obs} \right)^{-1}} \cdot \dfrac{VAR_{car}^t + VAR_{obs}}{\left(VAR_{car}^t \right)^2}$. After further manipulation, this in

turn reduces to: $\beta_{VAR} > \dfrac{\beta_x \cdot \Delta}{VAR_{car}^t}$.