6-5-2021

Adaptive Routing Behavior with Real-Time Information Under Multiple Travel Objectives

Ravi Venkatraman
Portland State University

Stephen D. Boyles
The University of Texas at Austin, sboyles@austin.utexas.edu

Rachel James
Federal Highways Administration

Avinash Unnikrishnan
Portland State University, uavinash@pdx.edu

Priyadarshan N. Patil
Portland State University

Citation Details

Venkatraman, R., Boyles, S. D., James, R., Unnikrishnan, A., & Patil, P. N. (2021). Adaptive routing behavior with real-time information under multiple travel objectives. Transportation Research Interdisciplinary Perspectives, 10, 100395.

Transportation Research Interdisciplinary Perspectives 10 (2021) 100395

Contents lists available at ScienceDirect

Transportation Research Interdisciplinary Perspectives

journal homepage: www.elsevier.com/locate/trip

Adaptive routing behavior with real-time information under multiple travel objectives

Ravi Venkatraman, Stephen D. Boyles, Rachel James, Avinash Unnikrishnan, Priyadarshan N. Patil *

[a] *Wayfair, United States*
[b] *The University of Texas at Austin, United States*
[c] *Federal Highways Administration, United States*
[d] *Portland State University, United States*
[e] *The University of Texas at Austin, United States*

ARTICLE INFO

Keywords:
Online shortest path
Adaptive routing
Real-time information
Routing games
Driver behavior

ABSTRACT

Real-time information about traffic conditions is becoming widely available through various media and connected-vehicle technology. In such conditions, travelers have better knowledge about the system and adapt as the system evolves dynamically during their travel. Drivers may change routes during their travel in order to optimize their own objective of travel. Various travel objectives are captured in mathematical models via disutility functions. The focus of this research was to study the behavior of travelers with multiple trip objectives when they are provided real-time information, and assess their ability to determine "optimal" routing policies, compared to exact solutions based on the online shortest path problem. A web-based experiment was carried out to simulate a traffic network with limited information provision. The decision strategies of participants were analyzed and compared to a variety of decision policies established in the literature – optimal, greedy, and *a priori* – and a general model to describe the observed travelers' decision strategies was calibrated from over 40,000 decision points extracted from the collected data. Apart from trip objective, other factors such as relative position in the network and experience gained were found to influence user decisions. This research is a step towards calibrating equilibrium models for adaptive behavior with multiple user classes.

1. Introduction

Uncertainty is inherent in every transportation network in the form of variable congestion levels, traffic incidents, and network closures due to bad weather conditions or road work. This stochasticity is a hindrance both to the users – who try to choose the best route within a transportation network – and to the planners – who try to develop an efficient transportation system. Recently, much effort has been devoted toward mitigating the effects of such uncertainties by providing users with real-time information about the network. This information may be provided through various media including, but not limited to, variable message signs (VMS), smartphones, GPS devices, or the radio.

Travelers respond to uncertainty and travel time information in different ways. Some drivers may stick with their initial, *a priori* route selection, while others may change their route online. This behavior partially depends on risk attitudes and trip purpose. One way of incorporating such considerations into route choice is through the use of disutility functions expressing arrival-time preferences. We specifically focus on disutility functions capturing different trip purposes. In this study, users' reactions were observed under uncertain network conditions and real-time information. This class of problems, where information is revealed *en route* and the user can make a series of adaptive decisions, is called online shortest path (OSP) problems.

The specific objectives of this study were twofold. The first objective was to study the behavior of individual travelers in a simulated environment with real-time local congestion information under multiple travel objectives. The travel objectives included a shopping trip (no target arrival time with lateness disincentive), work trip (fixed target time with early/late disincentive), social trip (fixed arrival time with less harsh disincentives), and airport trip (fixed target time with harsh lateness disincentive). The second objective was to validate, through comparisons, the use of adaptive routing policies in practice. Given that the behavioral routing policies executed by participants were observed to not match those detailed in the literature, this research developed a *"hybrid"* policy from over 40,000 collected decision points that is more consistent with user behavior.

This study compared several disutility functions, corresponding to decision rules, that have embedded non-linear preferences of travelers. Driver routing behavior is known to be complex in the face of

* Corresponding author at: The University of Texas at Austin, 204 E. Dean Keeton St., C2200, Austin, TX 78712-1172, United States.
E-mail address: Priyadarshan@utexas.edu (P.N. Patil).

https://doi.org/10.1016/j.trip.2021.100395

uncertain travel times, and disutility functions are a flexible way to capture a range of behaviors that have been proposed in the literature, including the scenarios described in the previous paragraph. We describe several other proposed behaviors in the literature review.

This study also compared observed user behavior strategies to routing policies detailed in the literature for different travel conditions. Multinomial logit models were developed to understand user decision making as a function of network characteristics, current location, trip type, and uncertainty characteristics. Identifying appropriate specifications for disutility functions under uncertainty and information provision is critical as a better understanding of traveler decision making will help academics, planners, and engineers to develop better planning models that can more accurately predict flows and travel times. As such, the results of this study are a step towards calibrating equilibrium models with adaptive routing and heterogeneous routing policies.

Our study complements existing experiments involving route choice behavior. In particular, while there have been several recent studies using field data from driving observations, from driving simulators, or with multiple participants interacting simultaneously, we chose to conduct our experiments in a simpler, asynchronous web-based simulation. While the former setups are likely more realistic, our choice of experimental setting (1) allowed the collection of a very large data set, involving over 40,000 observations of route choice decisions; and (2) allowed us to specify a utility function directly (through the use of a numerical scoring mechanism and monetary incentives for performance). This allows us to measure drivers' ability to route adaptively under different utility functions without having to first estimate a utility function from field observations (which introduces confounding factors). Further, since our focus is on the ability of drivers to assimilate real-time information, and not on the emergence of equilibrium in multi-user settings, we do not believe an asynchronous experiment limits our analysis.

The remainder of this paper is organized as follows. Section 2 details literature important to the development of this experiment. Section 3 details the theoretical framework that is being used to model users' response to information. The theoretical framework is divided into two pieces: the traveler scenarios—which are the disutility functions used to encourage certain travel behaviors throughout the network through the use of a competitive incentive structure—and the policies—which are the mechanisms through which the results are being interpreted. Section 4 details the experimental framework designed to test this framework through the development of an interactive game called *SmartDrive*. Section 5 provides the results of the *SmartDrive* study segmented by trip purpose, user experience level, and policy. Section 6 details the comparison of the results of the study with existing routing policies and the creation of two new multinomial logit models that best match the behavior observed during *SmartDrive* play. Section 7 discusses limitations of our approach, and Section 8 provides conclusions and recommendations for future research.

2. Literature review

Ng et al. (1995) were among the earliest to emphasize the importance of travel information dissemination and analyzed the impact of various factors such as accuracy, cost, and type of information for the success of such systems. Srinivasan and Mahmassani (2000) modeled the route choice behavior under information provision as either complying with the active traffic information system (ATIS) information or using the same path as before. From their experiments, they found that tendency to comply with ATIS information increases with congestion levels and travel time savings. Mahmassani et al. (2003) analyze the factors impacting en-route switching to alternate destinations and routes for shopping trips using a stated-preference internet-based survey and develop a framework to account for the hierarchical nature of observed choices. They observe and corroborate route and destination switching phenomena as well as analyze the difference in behavior between providing full information as opposed to partial information. Peeta and Yu (2004) and Peeta and Yu (2005) propose a hybrid probabilistic-possibilistic model framework to incorporate the day-to-day and within-day dynamics of driver route choice given real-time information, as well as another hybrid route choice model using quantitative and fuzzy modeling. Both models are compared to traditional multinomial logit models and demonstrate better predictive power, showing potential for practical application. These studies propose advanced models to analyze this behavior, but are limited by the lack of generalization to different networks and in some cases, difficulty of intuitive interpretation. The standard MNL model is therefore chosen for its applicability and ease of interpretation.

ATIS systems aim to manage traffic better, enhance driving operations, and improve traveler safety (Adler and Blue, 1998). They have great potential for influencing user route choice (Abdel-Aty et al., 1997; Levinson, 2003). However, Avineri and Prashker (2006) noted that the provision of information does not always result in lower expected travel time, possibly because the nature of users' choices tends to be more heterogeneous when provided with information. The users either adopted a strategy to minimize expected travel time or became risk-averse by choosing the most reliable route. This corroborates the claims of travel time reliability for route choice made by Abdel-Aty et al. (1997). ATIS accuracy reduction has also been shown to shift driver choices towards reliable travel time routes (Ben-Elia et al., 2013). Dia and Panwai (2007) collected data from users and developed neural network models to ascertain the types of information provided in VMSs that are most influential. Prescriptive information provides the largest impact compared to descriptive and experiential information (Ben-Elia et al., 2013). Sawik et al. (2017) and Sawik et al. (2017) present two multi-criteria problem instances set in Spain and observe tradeoffs between various objectives. These experiments provide more insight on equilibrium and reaction to information systems than individual behavior.

Algorithms have been developed to describe routing behavior under uncertain travel times, incorporating different assumptions about risk preferences and how drivers choose routes in stochastic environments. As a few examples, researchers have proposed that drivers aim to minimize the probability of late arrival (Fan et al., 2005), a function of expected arrival time before and after a target (Gao, 2005a), exponential or quadratic utility functions (Eiger et al., 1985; Murthy and Sarkar, 1996), and linear combinations of mean and standard deviation (Khani and Boyles, 2015; Shahabi et al., 2015). Disutility functions can represent all of these behaviors exactly or nearly-exactly (the standard deviation can be approximated using the deviance metric we describe in the following section), allowing us to compare multiple proposed behaviors in a consistent framework.

Congestion experiments dealing with route choices under the provision of information have also been well documented in the literature (Chen and Mahmassani, 1993; Rapoport et al., 2014; Tang et al., 2017; Lu et al., 2011). Many of these studies have been carried out with the objective of assessing multiplayer interaction and convergence to equilibrium to assess the influence of information on driving behavior. Ramadurai and Ukkusuri (2007) conducted an online multiplayer network game to check the conversion to a steady state in a dynamic network with a single bottleneck. They also studied the impact of online information on users' payoff. The decisions in the experiment were to choose departure time to arrive at the destination to maximize payoff. In addition to observing no convergence to equilibrium, they also observed paradoxical behavior where providing information yielded smaller overall payoffs. However, the small sample size, in terms of players and number of rounds played, warrants further analysis. Morgan et al. (2009) tested for change in traffic flows in a multiplayer setting when changes were made to a network; they observed that flows did shift on changing network conditions, but more towards user equilibrium rather than system equilibrium.

The details of ATIS information reception and usage are not fully understood (Chorus et al., 2013; Ben-Elia and Avineri, 2015; Zhao et al., 2019), prompting theoretical and experimental studies. Liu et al. (2020) provide a comparative review of recent studies focusing on experimental studies about travelers' day-to-day route choices. They also explore the effect of ATIS penetration, and results show that flow patterns tend to converge to user equilibrium despite certain fluctuations. Additionally, they note that their system instance was most stable at about 75% ATIS penetration, implying that full ATIS availability might not best for system performance. Earlier studies have tried to analyze the impact of ATIS in the form of variable message signs (VMSs) and have found that around one-fifth of the drivers change their routes based on VMS information, as well as speed dips to compensate for information processing (Erke et al., 2007; Lee and Abdel-Aty, 2008; Harms et al., 2019).

Multiple recent studies have focused on experimental analysis of ATIS impact on various facets of traffic assignment. Ringhand and Vollrath (2018) study the impact of ATIS on moving the system towards a system optimal configuration and conclude that around one-tenth of the drivers accept the system optimal alternatives. Wijayaratna et al. (2017) investigated the empirical presence of the online information paradox in an experimental computerized setting and concluded that provision of online information can indeed deteriorate travel conditions for all users. Additionally, they conclude that information availability trades off system performance for a reduction in travel time volatility. Wijayaratna and Dixit (2016) test user risk preferences in the presence of online information using expected utility theory and indicate that information availability leads to a reduction of risk aversion.as well as Fechner error (behavioral error).

Delle Site (2018) study route choice in a multi-class setting where classes are differentiated by access to ATIS and experiential data. They observe the higher travel times experienced by ATIS users, as well as higher total system travel time savings for greater ATIS market penetration. Ramos et al. (2018) use reinforcement learning with regret-minimization for the route choice problem and show significant improvements over Q-learning approach as well as regret-minimization without ATIS. More importantly, the system is shown to converge to approximate user equilibrium under the influence of ATIS. Ma and Di Pace (2017) aim to model travelers' day-to-day route choice for ATIS using four strategy learning models, viz., reinforcement learning, extended reinforcement learning, joint strategy fictitious play, and Bayesian learning. The joint strategy fictitious play model outperforms the other proposed models for intermediate and high accuracy information scenarios, with low accuracy information reading to random behavior. However, this study does not incorporate traveler risk attitudes.

Ramos et al. (2020) present recursive route choice models based on both static and dynamic travel time representations and a case study based in the Netherlands with multiple information sources. They conclude that the parameter interpretations are drastically different depending on non-additive utility inclusion, and correlation inclusion, as well as dynamic representation, provides better predictive power. Lastly, Jiang et al. (2020) present a rational inattention model for the stochastic route choice problem. Inspired from their prior work showing the lack of real-time information usage, this model assumes that information is costly and users acquire limited information before choosing an alternative, incorporating driver information strategy endogenously. The single-user model shows the reduction in alternatives under consideration with information increase, and the model with multiple heterogeneous users has a unique equilibrium. However, the model is path-based, raising tractability questions, as well as requiring empirical information for calibration.

In summary, multiple routing behaviors have been proposed in the past literature. We hypothesize that a MNL model can represent the "choice" drivers (subconsciously) make as to what type of routing behavior they follow in given situations, and explore this further in the remainder of the paper.

2.1. Current study

This research looked at a more fundamental behavior at the individual level. Route choice behavior was analyzed assuming complete access to real-time information with non-linear preferences. This study analyzed multiple trip purposes (multiple payoff functions). Identifying the policy followed by users to make *en route* decisions was the ultimate goal of this effort; thus, multiplayer interaction was not considered. The focus of this project was more on individual behavior, not the presence, or lack thereof, of equilibrium.

The study was carried out using a web application that simulated the route choice process while navigating through a network with stochastic costs. Over 40,000 decision points were obtained and analyzed after repetitive simulations by users. Initial trial scenarios and provision of historical data on the network, described in the experimental design section, served as tools for learning network conditions. The number of data points and learning process enabled us to ensure the simulation resembles a practical scenario and can therefore draw statistically sound conclusions. Comparisons were made to the optimal, greedy, and *a priori* routing policies and multinomial logit (MNL) models were developed to capture discrepancies between the documented policies and the observed user behavior. To the authors' knowledge, no study specifically considers the decision strategy of users in a stochastic network with non-linear preferences. We, therefore, aim to provide useful insights about individual behavior. These results may also be a step towards calibrating equilibrium models with adaptive routing and heterogeneous routing policies.

3. Theoretical framework

Two critical pieces of this work were the use of various traveler scenarios – which allow the study of non-linear user preferences – and routing policies – which allow the comparative study of the decision strategies detailed in literature against those observed though *SmartDrive* gameplay. This section details the disutility functions used to capture users' routing decisions for various scenarios and the policies in the literature that are currently used to model travel decisions. The disutility functions are used to influence user behavior to mimic specific behavior patterns. Put together, this section describes the theoretical framework used to model response to information. The next section, Section 4 describes the experimental framework designed to test this theoretical framework.

3.1. Traveler scenarios

Traditional travel objectives of minimizing travel time do not capture non-linear user behavior (e.g. user preference for a target arrival time at the destination or the risk preferences of a user when faced with uncertain travel times). The disutility function, $f^q(t)$, of a specific user class q describes the 'cost' of arriving at the destination at time t. Possible disutility functions, described by Boyles (2009) (linear, deviance, quadratic, etc.), were modified to represent different scenarios.

The benefits and improved accuracy of incorporating non-linear preferences for traveler choices have been well documented in the literature for both the third and fourth steps of the planning process (Mandel et al., 1994; De Lapparent et al., 2002; Small et al., 2005; Palma and Picard, 2005; Pinjari and Bhat, 2006; Fosgerau and Karlström, 2010). Non-linear preferences have also been incorporated in routing decisions in the context of reliability and preference for robust paths (Gao, 2005b; Boyles and Waller, 2007). The following list provides a brief description of the possible disutility functions:

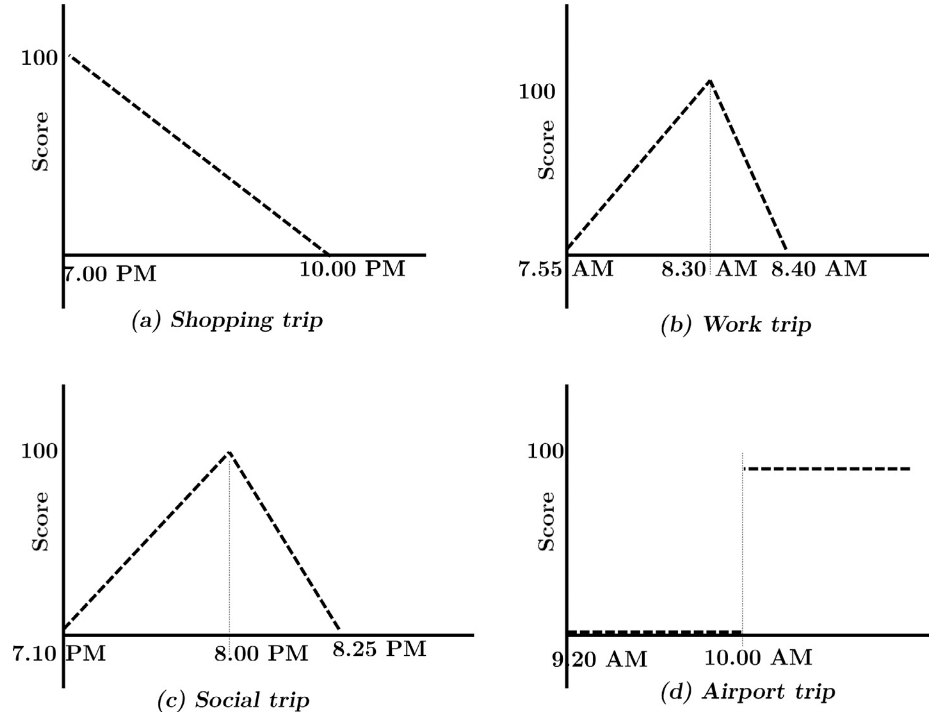

Fig. 1. Game scoring Rules.

- **Linear:** Linear disutility functions describe the standard shortest path objective, that is, we wish to minimize $E[t]$. Arriving at the destination as soon as possible is the primary concern, and the traveler is risk-neutral.
- **Deviance:** The deviance disutility is defined as $f(t) = (t - t^*)^2$, where t^* is the desired arrival time. The optimal policy now minimizes $E[(t - t^*)^2]$. In this case, the 'penalty' for an early or late arrival is the same i.e. there is no difference between arriving early by Δt units or late by Δt units. In some cases, late arrival may be a heavier burden than an earlier arrival, or vice versa. A simple modification to account for such a case would be a disutility function given below: $f(t) = \begin{cases} (t - t^*)^2, & \text{if } t \leqslant t^* \\ b(t - t^*)^2, & \text{if } t \geqslant t^* \end{cases}$ where $b > 1$ if the penalty for late arrival is greater than early arrival, $0 \leqslant b < 1$ if the penalty for early arrival is greater.
- **Quadratic:** A quadratic disutility function captures the risk-taking characteristics of a user class. Boyles and Waller (2007) parameterize this behavior using a single parameter k, representing the change in derivative of f between a range of possible arrival times. For the same value of disutility, a convex function allows for a later arrival time than a concave function. A user class with such behavior will be less prone to taking risks to arrive earlier. $k > 0$ yields a convex function, and represents a risk-averse behavior, whereas $k < 0$ yields a concave function representing a risk-prone behavior.
- **Arrival on time:** Nie and Fan (2006) and Nie and Wu (2009) address the problem of maximizing the probability of arriving at or before a specified time. The arrival on time disutility is used to represent the scenario of a traveler wishing to arrive at the destination no later than a threshold time t^*. The function is represented by an indicator function: $f(t) = \begin{cases} 1, & \text{if } t \leqslant t^* \\ 0, & \text{if } t > t^* \end{cases}$ since the expectation of this function is exactly the probability of late arrival (which, as a disutility, should be minimized).

The common travel objectives used in the *SmartDrive* web application were shopping, work, social, and airport trips. The corresponding disutility functions are shown in Fig. 1. The following bullets detail why each disutility function was selected for its respective travel scenario. Scores in the game were taken to be the negative of disutility (so higher scores correspond to lower disutility).

Shopping Trip : The disutility of users traveling for a shopping trip is modeled via a linear function with the behavioral assumption that the user wants to arrive at the destination as soon as possible with no particular target time.

Work Trip : Work trips are used to describe a scenario with a target arrival time. The user certainly does not wish to arrive late to the workplace because there is a high penalty associated with tardiness. Furthermore, one may argue that it is not beneficial to arrive early to work either (e.g. arriving early for collaborative meetings has no incentives and can be seen as unproductive time for fixed working hours). Note that the slope after the target arrival time is *steeper* than the slope before the arrival time. This scenario is an example of "schedule delay", which is a measure of the difference between a target arrival time and actual arrival time. This trip purpose corresponds to a deviance disutility function.

Social Trip : The social trip scenario assumed the user was traveling to a party, event, game, or something of the like. The user may not want to arrive too early for a social occasion, nor may he/she want to arrive late and miss the event. The target arrival times and penalties for late or early arrival are more relaxed than those for the work trip. This trip purpose also corresponds to a deviance disutility function.

Airport Trip : A user traveling to the airport to catch a flight needs to be on time at the airport. There is no extra incentive if he/she arrives before the target time. However, if the arrival time is later than the target time by any amount, the flight is missed. This is represented by the step function in the web application, with no incentive beyond the target arrival time. Specifically, the airport trip corresponds to the 'arrival on time' disutility function.

The users were encouraged to follow the presented travel scenario—shopping trip, work trip, social trip, or airport trip—through the use of the scoring rules in Fig. 1 and a competitive incentive structure. The top 100 scoring players all received a monetary award for participating in the experiment; the exact value of that award was tiered such that higher scoring players received more than lower scoring players.

3.2. Routing policies

The disutility function of user class affects the routing policy that the user belonging to a particular class follows. A policy is defined as a decision-making rule followed by a user, which may depend on factors like the current location of node i, current time t, the current state of the node θ, and target travel objective, which is represented in the disutility function. Such a policy may not lead to a fixed *a priori* path; instead, they will describe a *hyperpath*, or a path based on decision rules dependent on the current location, time, and information. Part of this research's contribution to the literature is the comparison of observed route choice behavior to policies well detailed in the literature. The following sections offer a brief summary of optimal, greedy, and *a priori* policies for routing decisions.

3.2.1. Online shortest path review

Networks with stochastic costs have been of interest to researchers for a long time. One of the first attempts at finding paths with minimum expected travel time in a stochastic and time-dependent network was by Hall (1986), who discussed the need for an adaptive decision strategy rather than an *a priori* path. More recently, considerable emphasis has been placed on the provision of real-time information through ITS technologies, leading to research on OSP problems (Boyles and Rambha, 2016; Cheung, 1998; Provan, 2003). Waller and Ziliaskopoulos (2002) studied two versions of the OSP problems, one with spatial dependence – where downstream arc travel time probabilities were conditioned on the travel times of upstream arcs – and temporal dependence – where the cost of the arc was learned upon arrival at the tail of the arc. OSP problems can also be cast as Markov decision processes (Boyles and Rambha, 2016). Unnikrishnan and Waller (2009) developed a convex programming formulation for user equilibrium when users follow a adaptive routing strategy. Boyles and Waller (2011) built on this and found the optimal locations for providing information by constructing contracted networks. While there have been many advances in this area, a few notable works are cited here for a general overview (Szeto et al., 2011; Wu, 2015; Khani and Boyles, 2015; Khani, 2019).

3.2.2. Optimal policy

The optimal strategy minimizes the expected disutility of the user class. With local information on the downstream arcs, the problem is similar to the one step temporal dependence in online shortest path (TD-OSP) problems considered by Waller and Ziliaskopoulos (2002). The TD-OSP algorithm suggested is based on the label correcting algorithm (Ahuja et al., 1993). Essentially, the algorithm starts at the destination node and works backward until the optimal labels are known for all nodes and time periods. This work uses an adaptive policy algorithm, suggested by Boyles (2009), that computes the optimal labels $L(i,t)$, and policy $\pi(i,t,\theta)$ for the network. For the exact algorithmic specification used in this work, see Algorithm 1.

Algorithm 1: AdaptivePolicy(**t**, **f**, $\mathcal{V}$)

$\boldsymbol{t}$ contains the travel times for each link in every state. $\boldsymbol{f}$ and v represent the disutility function and destination, respectively.
for $i \in N$ **do**
  for $t \in T_v$ **do**
    $L(i,t) \leftarrow \infty$
    for $\theta \in \Theta_i$ **do**
      $\pi(i,t,\theta) \leftarrow \phi$
for $t \in T_v$ **do**
  $L(v,t) \leftarrow f(t)$
while $t \geq 0$ **do**
  for $i \in N$ **do**
    $temp_L \leftarrow 0$
    for $\theta \in \Theta_i$ **do**
      $temp^\theta \leftarrow L(i,t)$
      for $j \in \Gamma(i)$ **do**
        $temp_j \leftarrow 0$
        if $t + t^theta_{ij} > T$ **then**
          $temp_j \leftarrow \infty$
        else
          $temp_j \leftarrow L(j, t + t_{ij}^{\theta})$
        if $temp_j < temp_\theta$ **then**
          $temp^\theta \leftarrow temp_j$
          $\pi(i,t,\theta) \leftarrow j$
      $temp_L \leftarrow temp_L + temp^\theta * Pr(\theta)$
    $L(i,t) \leftarrow temp_L$
  $t \leftarrow t - 1$

The algorithm AdaptivePolicy runs in $O((n+m)T|\Theta|)$ time, where $|\Theta|$ is the maximum number of node states examined for a node. Note that $|\Theta|$ can be $O(S^m)$. However, a reduction proposed by Waller and Ziliaskopoulos (2002) reduces the number of states scanned to $O(SM)$. For a node with A downstream arcs with S states each, we can reduce the node states from S^A to SA. While this reduction is not used for our implementation due to relatively small computational expense, it is useful for experiments on larger networks.

3.2.3. A priori policy

Another policy that might be observed is following an *a priori* path that yields the least expected disutility. Miller-Hooks (2001) proposed an algorithm to compute least expected travel time (LET) paths in uncertain networks with links whose probability distributions vary with time. A user with this strategy does not make use of local information but instead relies on past experiences.

The LET algorithm computes $\min_{p \in P} E[t(p)]$, where P is the set of all paths from the origin to the destination, and $t(p)$ denotes the travel time of path p. Irrespective of correlation between link travel times and distributions, $E[t(p)] = \sum_{(i,j) \in p} E[t_{ij}]$. In this case, the probability distributions of a link do not vary with time, hence $E[t_{ij}]$ does not depend on arrival time. Thus, the same adaptive policy algorithm stated above is used to calculate least expected disutility paths by replacing the multiple states of each arc with a single state with travel time $E[t_{ij}]$. Note that $|\Theta| = 1$ and the algorithm runs in $O(nT)$ time. The termination and correctness of this algorithm are guaranteed as this is a special case of the optimal policy algorithm.

3.2.4. Greedy policy

The greedy policy, similar to the traditional definition, represents a myopic user behavior. Though the local information readily available is used, the broader objective of minimizing expected disutility through route decisions is not taken into consideration.

Algorithm 2: GreedyPolicy(**t**, $\mathcal{V}$)

$\boldsymbol{t}$ contains the travel times for each link in every state and v represents destination.
for $i \in N$ **do**
    $RestrictedNodes(i) \leftarrow \phi$
    $updateRestrictedNodes(i)$
    for $\theta \in \Theta_i$ **do**
        $\pi(i, \theta) \leftarrow \phi$
for $i \in N$ **do**
    for $\theta \in \Theta_i$ **do**
        $temp^{\theta} \leftarrow \infty$
        for $j \in \Gamma(i)$ **do**
            $temp_j \leftarrow 0$
            if $temp_j > t_{ij}^{\theta} AND j \notin RestrictedNodes(i)$ **then**
                $temp_j \leftarrow t_{ij}^{\theta}$
                $\pi(i, \theta) \leftarrow j$

The greedy strategy involves choosing a link adjacent to the current node with the least congestion level to create a route to the destination; additionally, the algorithm implements two behavioral constraints: the user does not traverse in a direction opposite to the destination and spiral paths through the network are avoided. An example network is shown in Fig. 2. Assume a user traveling using the greedy strategy from node 1 to node 8. The double-sided arrows indicate the presence of two separate links in each direction.

- The user does not move away from the destination. For example, if the user is currently at node 5, the user can travel to nodes 4,6 and 8, but not to node 2, which is a direction opposite to the destination. This also includes directions (or vectors) that have a component in the opposite direction. For example, the user cannot travel from node 5 to nodes 3 or 1 (if links 5-3, 5-1 existed), since the vector 5-3 (or 5-1) has a component along 5-2. Dial (1971) considers 'reasonable paths' for logit based traffic assignment, and bases the criteria on shortest path distance of a node to (and from) the destination (source). Our criterion is based on the geographical direction of travel with respect to the destination, while allowing shortest path distances to the destination to increase on traversal. As a result, the user still moves in the general direction of the destination, but in a myopic manner not considering the remainder of the path.
- Spiral paths are avoided. Consider a hypothetical line joining the origin and destination, node 1 and 8 in 3.6. A spiral is a path that crosses this line more than once. For example, a path 1-2-3-6-5-4-7-8 is a spiral since it crosses the hypothetical line twice, once at 1-2 and again at 5-4.

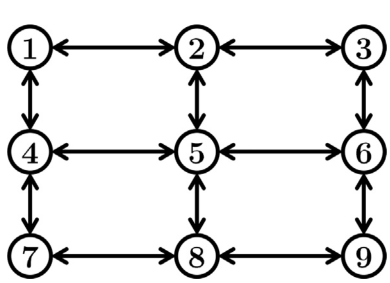

Fig. 2. Greedy policy example network.

To avoid this behavior, we define a set of restricted nodes from each node *i*, which consist of nodes that should not be traversed to from the current location. This set can be initialized *a priori*, but additions may be made during traversal. For example, to avoid a spiral path, once the path 1-2-3-6 has been traversed, nodes 4 and 7 may be added to the set of restricted nodes from node 5, if they were not already present in the set. This strategy does not consider nonlinear user behavior, the objective is to reach the destination minimizing the travel time.

The exact algorithmic specification of the greedy algorithm can be found in Algorithm 2. GreedyPolicy runs in $O(m|\Theta|)$ time. The termination and correctness of this algorithm are trivial.

3.2.5. Policy simulation

Simulation of all the policies is carried out on the network used for the application. The objective of carrying out simulations are threefold:

1. Verify the optimal and *a priori* labels
2. Compare route choices in each of the above policies
3. Use results to develop other policies that match user behavior closely and compare these with other definite policies.

We iteratively simulate each of the policies for each specified disutility function, as shown in Algorithm 3.

Algorithm 3: SimulatePolicy(**t, f**, *startTime, s, V*)

$\boldsymbol{t}$ contains the travel times for each link in ever state, $\boldsymbol{f}$ contains the disutility function, $startTime$, s, v represent the starting time, origin and destination, respectively.

$i \leftarrow s$

$t \leftarrow StartTime$

Sample θ from given probability distribution at node i

while $i \neq v AND \pi(i,t,\theta) \neq \phi$ **do**

  $updateRestrictedNodes(i)$

  $i \leftarrow \pi(i,t,\theta)$

  $t \leftarrow t + t^{s_\theta}_{i,\pi(i,t,\theta)}$

  Sample θ from given probability distribution at node i

$disutility = f(t)$

3.2.6. Example

A simple example is demonstrated to compare and contrast the three policies stated previously. Consider the familiar 4-node Braess network. The network, associated states and their costs are shown in Fig. 3 and Table 1 respectively. The probabilities are hypothetical for this example but can be gathered from historical data for real world networks, with states representing different roadway conditions. Consider a user class with a linear disutility function traveling from node 1 to node 4, i.e. the objective is to minimize expected travel time. We evaluate different policies and demonstrate an instance using the same.

The *a priori* path can be obtained by simply enumerating the three possible paths and choosing the one with least expected travel time. The *a priori* path is path 1-2-4 with expected travel time of 8. The optimal policy is a set of decisions to minimize expected travel time, using the observed downstream information. The labels with least expected values of travel time are shown in Fig. 4, and the corresponding policy is constructed in Table 2.

We construct an instance of this problem and demonstrate the progression using the three policies. The simulation, as outlined in Algorithm 3, is shown in Fig. 5. The highlighted nodes for each stage represent the current location. At node 1, information regarding travel times of downstream links to nodes 2 and 3 are obtained, as 4 and 3, respectively. The *a priori* strategy follows the least expected path and chooses node 2. The greedy strategy chooses node 3 since the immediate travel time node 3 is less than node 4, without any information regarding the travel time distributions of arcs downstream of 2 and 3. The optimal strategy is to observe the current information, and choose the node which minimizes the expected travel time from node 2 or 3 i.e. choosing node 2 would yield an expected travel time of 9 (= 5 + 4), whereas choosing node 3 would yield an expected travel time of 7 (= 3 + 4).

At the next stage, the *a priori* strategy continues from node 2 to the destination 4, yielding a travel time of 7 (less than the expected travel time of 8). At node 3, the deterministic link to 2 with travel time of 2 is observed, and the link to 4 is observed with travel time 3. Now, the optimal policy chooses the minimum of paths 3-2-4 (with expected travel time 6) and 3-4 (with deterministic travel time 3), and chooses to reach the destination. The total travel time in this case is 6 (less than the initial expected travel time of 10). However, the greedy policy chooses to reach node 2, since the travel time on link 3-2 is less than that of 3-4. At the next stage, the observed travel time on link 2-4 is 6, and the same is followed. The total travel time for this strategy is 11. In this specific instance, we have $T_{optimal} < T_{apriori} < T_{greedy}$.

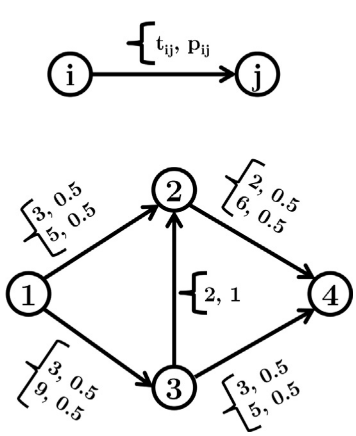

Fig. 3. Policy illustration: Braess' network.

4. Experiment design

A web application was created to assess the route choice decisions of users under the presence of local information and multiple travel objectives. In order to make the experience of using the interface enjoyable, the study was designed as a game with the objective of maximizing a score; based on their relative scores, the top 80 performers were eligible for an incentive ranging from $10–$50. These incentives were used to motivate people to use the application with a clear goal of achieving the stated objectives. The study was publicized through the web, university mailing lists, social networks, and fliers. The details of game design and the process of conducting the experiments are explained further in this section.

Table 1
Travel time distribution: Braess' network.

Arc	State 1		State 1	
	Cost	Probability	Cost	Probability
1–2	3	0.5	5	0.5
1–3	3	0.5	9	0.5
2–4	2	0.5	6	0.5
3–2	2	1.0	–	0.0
3–4	3	0.5	5	0.5

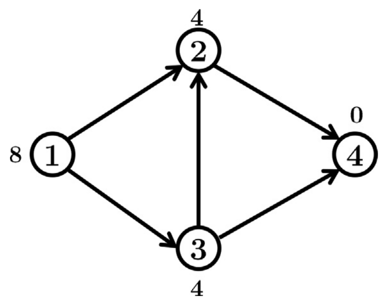

Fig. 4. Expected value labels.

Table 2
Optimal policy: Braess' network.

Node	Downstream Arc	State/Cost	Optimal Policy
1	1-2	3	2
	1-3	3	
	1-2	3	2
	1-3	9	
	1-2	5	3
	1-3	3	
	1-2	5	2
	1-3	9	
2	2-4	2	4
	2-4	6	4
3	3-2	2	4
	3-4	3	
	3-2	2	4
	3-4	5	

4.1. SmartDrive scenarios

Users repeated the game many times, with different travel objectives. These objectives are defined by the disutility functions listed in Fig. 1. In order to make it more comprehensible to the users, the disutility minimization objective was converted to a score maximization objective. Hence, the scoring scheme was defined as $score(t) = -f(t) + score_{max}$, where $score_{max}$ is the maximum achievable score and was set to 100. The revised scoring graphs are shown in Fig. 1, with t^* being the target time of arrival, if any, and t_M being the latest allowable time of arrival, beyond which there is no loss of incentive to arrive at the destination. The average score for a user is the average score obtained over all the scenarios. Section 5.2 presents analysis of average user score as a function of experience with the game.

4.2. SmartDrive navigation

The network used for the application was the Sioux Falls network with 24 nodes and 76 links. The origin and destination for each scenario were held constant. Fig. 6 shows an instance in the game in the 'airport' scenario.

Navigating the network was straightforward. The white circle indicated the user's current position; the objective was to get to the destination, marked by 'X', while maximizing the score. The user could choose his/her route at every intersection by clicking on the respective downstream arc. The current time, target arrival time, and the score were also displayed on the screen.

As the user proceeded through the network, real-time traffic information was provided. Congestion levels on all the downstream links

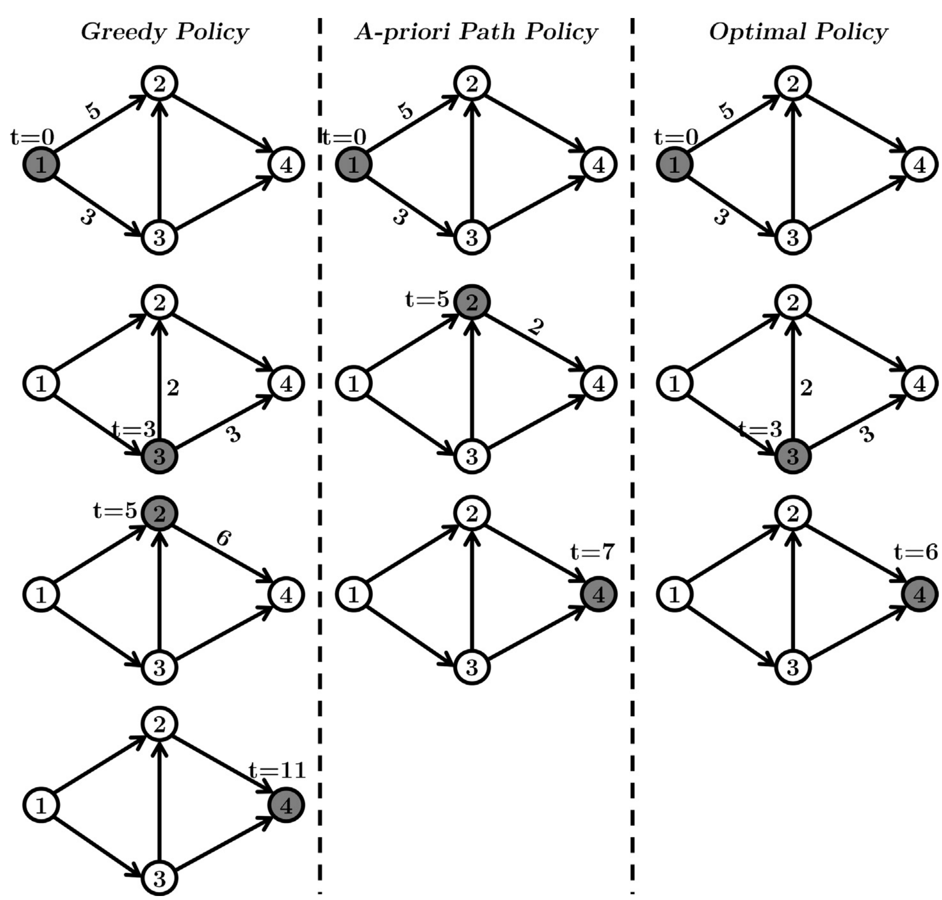

Fig. 5. Policy comparisons.

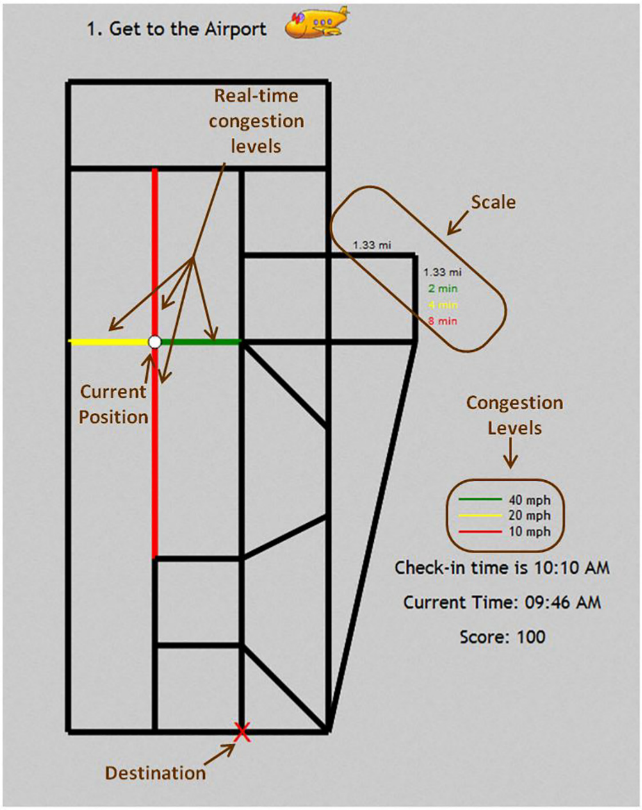

Fig. 6. Game: Network Navigation.

was shown through color codes, which were an indication of the average speed on that road. The distance scale was also displayed, highlighted in Fig. 6. Thus, users could use the scale and the immediate downstream congestion level to approximate travel time. When the destination was reached, users were given the option to end the trip or continue driving, which may have been the more optimal decision in cases with a target arrival time. If the user chose to end the trip, the path they followed was highlighted temporarily before beginning the next round in order to give them a recap of their trip and reassess the travel conditions.

To get familiarized with the navigation procedures, users were given three trial runs before beginning the simulations. Live information was only available on all the immediate downstream links. Using a drop-down menu, users were allowed to see the daily congestion levels on the entire network for the past 30 days. Each day's congestion level was formed by sampling the link travel time distribution, with each link's state drawn independently of any other state. This tool was designed to help users navigate through the network and enhance the process of learning the network conditions, simulating their experience on a real network they might travel daily. The next section discusses the score findings of the web application and its implications.

5. Analysis

The application was circulated through the web for a period of 2 months from November 2012 until December 2012. Users were eligible to compete for incentives if they played a minimum of 10 rounds, but were allowed to play up to 100. Upon conclusion, the game had 131 registered users who played a total of 5203 scenarios and provided over 40,000 individual decision points. Of the 131 respondents to the game, 67% were male. 57% were aged 18–25, 26% were aged 26–32, and the rest (17%) reported ages older than 32. 45% of the respondents were students and 40% were working in full-time positions. The numbers show that this sample does not represent the general US population; because the application was web-based and propagation was through web referrals, our control over the sample composition was minimal. Our intention was to study the decision-making process of users in a stochastic network with access to real-time information, subject to parameters specific to the transportation network (trip purpose, information observed, learning trends, etc.) and not trends dependent on demographics of the population. Demographic and socio-economic characteristics may certainly impact the decision-making process, but the main motive of this research focuses on a different set of parameters, thus eliminating the need for a representative sample to obtain meaningful results.

5.1. User performance segmented by trip purpose

Aggregated across scenario, the scores for all 5203 scenarios roughly followed a peaked distribution, with the majority (80%) of the respondents scoring an average of 60–80, and close to 10% of respondents each on the higher end – greater than 80 – and lower end – less than 60 – of the spectrum. Interestingly, the score distributions disaggregated by scenario are not peaked, as seen in Fig. 7. By absolute value of score, the 'social' trip and 'work' trip are the best and worst among the four, respectively. These two have a similar piece-wise disutility function, but the social trip has a smaller penalty for late arrival. The shopping trip, with a linear disutility, has a more consistent performance with close to 90% of the users scoring between 70–80. The airport trip, with a threshold arrival time and all or nothing score, has a wider distribution, with a higher composition (63%) of respondents scoring below 70. However, a decent percentage (22.90%) of respondents have scored above 80 in this scenario, which is equivalent to stating that they make the trip on time 8 out of 10 times.

5.2. Average score as a function of experience with game

We also studied the 'learning' process of users. Fig. 8 shows the learning curve with the number of rounds played, along with the expected score from the optimal policies. (Note: at one point the observed average score rises above the optimal value; this is because the optimal value in the Figure is an expected value, and for specific realizations, higher scores are possible with skillful play.) The scores of each user were calculated as a moving average with a fixed time period of 10 rounds. Let $S_i(n)$ be the score of user i in the n^{th} round of the specific trip, and let $S_i^A(n)$ denote the average score of user i

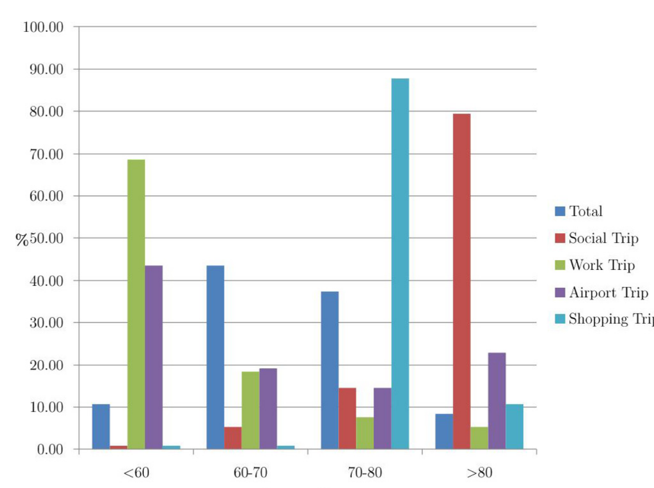

Fig. 7. Score distribution segmented by trip purpose.

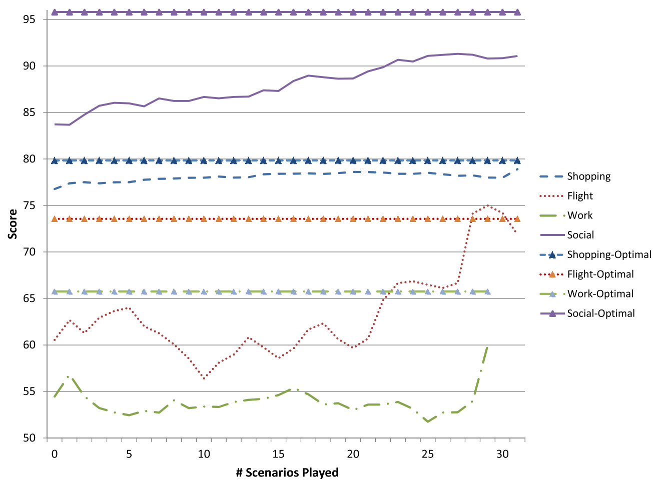

Fig. 8. Average score with #Rounds played.

in the n^{th} round. The moving average over a fixed $T = (5)$ periods for each user was computed as

$$S_i^A(n) = \sum_{t=n-T}^{t=n} S_i(t)/\min\{t, T\} \tag{1}$$

The average score over all users, $AvgScore(n)$, was computed as

$$AvgScore(n) = \sum_{i_n} S_{i_n}(n)/|i_n| \tag{2}$$

where i_n is the index of a user in the n^{th} round, and $|i_n|$ is the number of people who have played n rounds of the specific scenario. Fig. 8 shows the plot of $AvgScore(n)$ with the number of rounds played.

As seen in Fig. 8, there was a steep increase in average scores with the number of rounds for the social trip and an indication of a learning trend for the shopping trip. There was an increasing trend for the airport trip as well, with 'noise' in the curve. This could be attributed to the all-or-nothing scoring pattern of this trip, which leads to drastic fluctuations in the average scores. There was no apparent learning trend demonstrated in the work trip. This might be due to the strict time constraint and a heavy penalty for tardiness in this scenario. Additionally, it is possible that users were inhibited from trying out different strategies in this scenario. However, one must be cautious while studying these trends, particularly for numbers of rounds played greater than 25. As the number of rounds increased, the number of users that participated decreased. Hence, the average scores towards the end of the curve were based on fewer participants.

5.3. Score results segmented by policy

The three policies discussed in the policies subsection – the optimal, greedy, and *a priori* path policies – were used for comparison with users' observed path decisions. For each of the policies, 10000 Monte Carlo simulations were run for each of the four disutility functions to compute an average score for each policy. These are documented in Table 3. The labels from the optimal policy ('Expected Disutility') were compared with the simulation results and the scores were found to coincide with each other. Interestingly, the average score from the greedy strategy was better than following a fixed *a priori* path for the social and shopping trips. However, the greedy strategy performed very poorly for work and airport trips. In these two scenarios, late arrival was penalized more heavily than in the other two scenarios; the greedy strategy was observed to overshoot the target arrival time more often as there is no foresight while choosing the next node. From this analysis, it is evident that when embarking on trips with strict time constraints using a purely myopic decision rule is detrimental.

The users' average scores were both interesting and encouraging. For all trips, the user scores were less than the optimal values, as expected, but higher than the average scores from the other two strategies. This implies users followed a strategy that is not completely myopic nor did they follow a fixed path. Instead, the users' strategy may be a combination of the three strategies and may depend on various other parameters. This possibility is explored in the next section with the creation of a hybrid policy.

6. Development of new route choice models

One of the research objectives of this study was to determine what heuristic travelers use when making routing decisions and whether or not this heuristic is similar to an existing routing policy in the literature. The three policies discussed earlier – optimal, greedy, and *a priori* path policies – were used for comparison with observed user decisions. Preliminary investigations showed that users tend to follow a decision strategy that depends on the disutility function. The objective of this project was to try and capture this decision-making process into a new mathematical model. Further examination of the data revealed trends with respect to other parameters such as the distance from the destination and number of nodes to choose from (outdegree of the current node). However, these trends did not justify a single deterministic strategy that captures the decisions of all users or even the same user at different stages. Hence, a random utility-based discrete choice model was developed to determine the policy adopted by the users.

6.1. Methodology

Nine different multinomial logit alternatives were created for the model (see Table 4). We considered the alternatives as distinct policies themselves; hence, there were three immediate alternatives policies: optimal, greedy, or *a priori* path policy. Another factor of interest was to see how much these three policies overlapped with each other. 35–40% of the decisions overlapped for each pair of policies for each scenario. Out of those, there were significant cases for which all three overlapped simultaneously. To incorporate the fact that policies were observed to overlap with each other and that this may increase the probability of that node being chosen, overlapping policies were introduced as separate alternatives. There were also many instances where users made decisions that did not coincide with any of the aforementioned policies. Thus, two alternatives were introduced; they corresponded to choices that do not fall under any of the policies under consideration. These alternatives are chosen as the nodes j whose

Table 3
Score Comparison: Users vs. Policies.

Scenario	Social	Work	Airport	Shopping
User Average	83.72	54.46	60.54	76.77
Expected Score (from Optimal Policy)	95.78	65.76	73.56	79.85
Average Score (from Optimal Policy)	95.58	64.41	71.26	79.55
Average Score (from Greedy Policy)	79.45	20.79	21.03	73.46
Average Score (from A *priori* Path Policy)	70.38	61.53	72.20	60.19

Table 4
MNL alternatives – Notations.

Notation	Policy
O ≡ G ≡ A	optimal, greedy and *a priori* path policies overlap
O ≡ G	optimal and greedy policies overlap, *a priori* path policy is different
O ≡ A	optimal and *a priori* path policies overlap, greedy policy is different
O	optimal policy is different from greedy and *a priori* path policies
G	greedy policy is different from optimal and *a priori* path policies
G ≡ A	greedy and *a priori* path policies overlap, optimal policy is different
A	*a priori* path policy is different from optimal and greedy policies
Oth1	Neither optimal, greedy or *a priori* path policy, and $SP_{j^*} = \min_j SP_j, j$ is an adjacent node
Oth2	Neither optimal, greedy or *a priori* path policy, and $SP_{j^{**}} = \min_{j \neq j^*} SP_j, j$ is an adjacent node

shortest length to the destination (SP_j) is least and do not fall under any of the policies (refer Table 4). The nine alternatives were referred to by notations indicated in Table 4.

The following independent variables are hypothesized to explain the data:

- Outdegree (*OutDegree*): Number of alternative links from which to choose.
- Distance from destination (*distanceDest*): Length of the shortest path to the destination (with respect to arc lengths).
- Experience/Network familiarity (*NumScenarios*): Number of scenarios played by the user prior to that point.
- Relative arc state (*relativeArcState*): Parameter which quantifies the relative congestion level. Each arc (i,j) was denoted by their state, $ArcState_{ij}$, which takes values $1, 2$, and 3, with 3 being the highest congestion level. For any arc (i,j), the relative congestion level was defined by

$$relativeArcState_{ij} = \sum_{k \in \Gamma(i)} (ArcState_{ik} - ArcState_{ij}) \tag{3}$$

  Thus, the higher the relative arc state, the lower the congestion level on the arc relative to the other downstream arcs.
- Probability distribution of arc states (*arcReliability*): Arcs were categorized into three divisions based on their probability distribution, with category 1 representing a deterministic arc and 3 representing high variability and equal probability of all states.

One of the research objectives was to assess the decision-making process under multiple trip objectives or disutility functions. Hence, a model was developed by segmenting the data by trip scenario. The results of the model are described in the next section. Recall that this research effort did not consider user characteristics. In a sense, the individual decision maker is not the user of the game; rather, it is the current “state” in the network (e.g. location, experience, number of immediate alternatives)

6.2. Results

The specification of the partially segmented model, including network specific parameters and respective coefficients, is tabulated in Table 5. The table does not include the constants.

It is interesting to see that many parameters have different and sometimes opposite impacts on different trip types. Recall the following features of the disutility functions. The work and airport trips have strict time constraints and high penalties for violating the target arrival time. The shopping trip has a time constraint in the sense that the disutility function is continuously increasing with time and there is no incentive for a user to delay the arrival time. Relative to other trips, the social trip has a more lenient time constraint and the penalty for late arrival is smaller. With this in mind, the impact of each of the parameters on the decision-making process is detailed below.

Table 5
MNL results – partially segmented model.

Variables	Coeff.	t-stat
O ≡ G ≡ A		
Link Parameters		
relativeArcState (Work, Airport)	0.11	3.45
relativeArcState (Shopping)	0.27	4.99
O ≡ G		
Current Location		
OutDegree (Work)	0.24	1.88
OutDegree (Airport, Shopping)	0.54	5.25
O ≡ A		
Current Location		
distanceDest (Social)	−2.18	−5.93
distanceDest (Work)	0.60	1.97
distanceDest (Airport)	1.24	1.99
O		
Learning		
NumScenarios (Social)	0.13	2.31
NumScenarios (Work, Airport, Shopping	0.09	2.48
Link Parameters		
relativeArcState (Work, Airport)	0.11	3.45
relativeArcState (Shopping)	0.27	4.99
G		
Current Location		
OutDegree (Social)	−0.23	−2.14
OutDegree (Work)	−2.14	−6.73
OutDegree (Airport, Shopping)	−1.72	−8.61
Link Parameters		
relativeArcState (Work, Airport)	0.11	3.45
relativeArcState (Shopping)	0.27	4.99
arcReliability (Social, Shopping)	0.35	5.49
arcReliability (Work)	1.38	10.46
arcReliability (Airport)	0.91	7.68
G ≡ A		
Current Location		
distanceDest (Social, Shopping)	−0.72	−3.63
distanceDest (Work, Airport)	−1.12	−5.04
Link Parameters		
arcReliability (Social, Shopping)	0.35	5.49
arcReliability (Work)	1.38	10.46
arcReliability (Airport)	0.91	7.68
A		
Current Location		
distanceDest	−1.39	−10.37
Oth1		
Current Location		
distanceDest (Work, Airport)	0.31	4.65
distanceDest (Shopping)	0.84	6.32
Log-likelihood (constants only)	−10105.08	
Final log-likelihood	−4882.57	
ρ^2	0.52	

Out-degree

The out-degree represents the number of alternatives to choose from, as complete information is available at all nodes in the experiment. Interestingly, a higher number of alternatives negatively impacted the preference to choose a greedy policy relative to the case when all three policies overlap. Intuitively, one would expect an increase in the number of choices to result in users following a more myopic strategy. In the case of a higher number of alternatives, users look at strategies other than a greedy strategy and work towards that. Moreover, the impact is highest in airport, work, and shopping trips. A plausible reason for this is that the three scenarios have higher constraints on time (strict time constraints or continuously increasing disutility), therefore users look to optimize their route more effectively. The social trip is more relaxed, which leads users to choose myopic strategies to explore other routes. When the optimal and greedy policies coincide, the likelihood of the alternative being chosen

Table 6
Model Validation.

Policy	Data	Constants only model	Full specification		
		Estimate	%age difference	Estimate	%age difference
O ≡ G ≡ A	11144	8694	21.98	9743	12.58
O ≡ G	12905	11531	10.65	13761	6.63
O ≡ A	1051	235	77.62	829	21.17
O	8405	7150	14.93	7587	9.73
G	2266	1611	28.9	1908	15.79
G ≡ A	947	393	58.54	1353	42.86
A	725	910	15.48	852	17.55
Oth1	4801	11455	138.59	6018	25.34
Oth2	717	983	37.03	911	27.02

increases with the number of alternatives. This may be because this alternative is usually the more 'easy' or 'obvious' choice to minimize disutility.

Distance from destination

One can see that the distance from the destination has a negative impact on the greedy and *a priori* strategies and a positive impact on the optimal policy alternative when it overlaps with the *a priori* strategy, relative to the case where all three policies overlap. It is interesting to note the highly positive impact of distance on choosing a policy (*Oth1*) that is neither of the policies documented in the literature. One plausible conclusion is that users are trying to optimize their routes closer to the origin when they start navigating. As they get closer to the destination, users prefer greedy strategies as they get them nearer to the destination in less time. In the process of optimizing their route when they start, they are neither following a purely myopic strategy or an optimal strategy, resulting in a choice that is not defined by any of the three policies. Further, the impact is higher for the work and airport trips for the greedy strategy. These trips, with strict time constraints and heavier penalties, discourage the user from following a purely myopic or *a priori* decision strategy; instead, they try to use information and experience to optimize their route. Given this outcome, one would expect the tendency to use an alternative policy in the work and airport trips to be higher than that for other alternative policies. However, the results show that distance from destination has less of an impact on these scenarios than a shopping trip while choosing other policies, *Oth*1. Interestingly, the location in the network affects the likelihood to choose an optimal decision (overlapping with an *a priori* strategy) in opposite ways depending on the scenario. For social trips, users are more likely to make optimal decisions closer to the destination while the trend is opposite for airport and work trips. One possible explanation is that users tend to follow optimal strategies with respect to the time-constrained work and airport scenarios (which may be similar) and follow the same path for social trips as well, even though the optimal strategy for the social trip may be to take a different route.

Experience/network familiarity

A learning process was demonstrated when the optimal policy was selected even though the optimal decision was different from other policies. Further, knowledge about the optimal policy was gained more in the social trip than other trips. This may be due to a phenomenon discussed previously; it may be that users follow optimal decisions with respect to the more time-constrained trips initially and later learn optimal decisions with respect to social trips.

Relative arc state

The relative arc state was found to influence the decision of four alternatives, which can primarily be classified as an optimal strategy and a greedy strategy. The higher the relative state of the arc – or the lower its relative congestion level – the higher the likelihood of it being preferred. The relative state in shopping trips has a greater influence than in the 'stricter' work and airport trips.

Probability distribution of arc states

A higher value for this parameter indicates more variability in the arc states. Thus, the parameter positively impacted the greedy strategy. Less arc reliability increases the possibility of a greedy strategy being adopted. That is, given an option between two greedy strategies at different instances, users preferred a greedy strategy when the arc reliability was lower. This may be due to a behavioral instinct to explore a new route when it has a lower travel time than typically observed, but an intuitive explanation is not clear. Further, the influence was higher on the work and airport trips than on social and shopping trips. This result indicates that it may be more insightful to study this simultaneously with the experience gained in a network (number of scenarios played) and to examine reliability at the path level rather than at the link level.

6.3. Validation

The model was developed using a random sample of 9,000 data points drawn from the dataset collected (close to 43,000). The model was validated using the remaining data points. For each observation i, the probability of choosing policy $\pi, P_i(\pi)$, was calculated as $e^{V_\pi}/\sum_j e^{V_j} {}^{*av_j}$, where av_j indicated if alternative j was available. V_j was the utility of the alternative computed using the specification mentioned in Table 5. The total estimates over all observations, $\sum_i P_i(\pi)$ are reported in Table 6 for the developed model and the model with constants only; note, the model with constants only was omitted for brevity and can be found in Venkatraman (2013). The model with constants only gives estimates of policy preferences when they overlap with each other. The model with specifications further builds on these preferences with network parameters. The results show that the model estimated trends in the entire dataset fairly accurately, given the number of uncertainties in user decisions. Specifically, all the estimates involving the optimal policy have estimates within 13% (except O ≡ A). In the estimate of the alternative being chosen when G ≡ A, the model with specifications overestimated the preference – a high transfer of alternative preferences from the underestimated count in the model with constants only. Drastic improvement in forecasts in the policy *Oth*1 and O ≡ A alternatives are also observed. Note that the three alternatives discussed now have *distanceDest* in their specification. Traveler's relative location along their path seems to be a factor that strongly influenced their driving behavior. Overall, the model gives a good description of the drivers' route choices.

7. Limitations

While our experimental setup yielded a large number of observations, allowing us to estimate the discrete choice models above with a high level of significance, it nevertheless also has significant limitations. This section discusses these issues.

First, and most significantly, web-based experiments can only capture parts of actual driver decision making process. There are major questions about the extent to which routing behavior in a “web game” format corresponds to routing behavior in actual driving environments, and our ability to control for external factors, such as distractions while participants are engaging in the experiment. It would be valuable to see whether the trends observed in our experiments can be replicated in more immersive settings, such as virtual reality experiments (Bowman and McMahan, 2007) or driving simulators (Blaauw, 1982; Nilsson, 1993; Godley et al., 2002).

A second major limitation is our assumption of independent link states. It is well-known that travel times exhibit spatial and temporal correlation. At the same time, many different models have been proposed for representing such correlations (for the purposes of routing algorithms and sampling), and the choice of correlation model significantly affects the ability to obtain benchmark policies. For example, with limited spatial dependency, for any constant “history” length the problem can still be solved in polynomial time (Waller and Ziliaskopoulos, 2002). By contrast, if revisited links are constrained to have the same cost as when first visited (“no reset”), the problem is either NP-complete or NP-hard, depending on additional choices of correlation structure (Provan, 2003). Repeating our experiments with correlations could be a highly valuable exercise, but choosing a particular correlation structure is nontrivial, from the standpoint of identifying a “realistic” structure, from the standpoint of implementation, and from the standpoint of interpretation (distinguishing the effect of correlation from other effects). We therefore chose to leave such investigations to future work as well.

An additional limitation is the use of our disutility functions and monetary incentives to simulate risk preferences under different scenarios. This experimental choice allowed us to represent scenarios with different levels of urgency and risk in a standardized way for all users of the web game, but participants’ risk preferences regarding the monetary incentive are unknown. Presuming risk neutrality regarding the monetary incentive, optimal behavior in the game corresponds exactly to maximizing expected score with the given disutility functions, which was our intent. Essentially, we assumed risk neutrality towards the monetary incentive in order to standardize and evaluate risk preferences regarding routing in the experiment. We believe this to be reasonable, given our experimental setup, but nevertheless an implicit assumption in our analysis.

8. Conclusions

This research effort studied individual travel behavior in response to real-time information under multiple travel objectives. The class of problems studied for this purpose was the online shortest path problems, where one can observe and adapt to the information gained *en route*. Different functional forms were used to represent disutility for the travel objectives. A web-based application was developed to study the behavior of travelers in such conditions via the simulation of an environment with multiple travel objectives and real-time information. Responses from 131 participants—from which over 40,000 decision points were extracted—were analyzed in detail. In order to compare user decision strategies to well defined mathematical policies (or decision rules), the observed user behavior was compared to the optimal, greedy, and *a priori* path policies.

The decisions have a common trend with respect to location in the network, specific information on downstream congestion levels gained, and familiarity gained with the network. In order to incorporate all such possibilities and uncertainty in strategy followed, an MNL model was developed to determine the preference for each policy. The results show that users’ decision strategies vary with travel objectives. Previously proposed policies were not observed in our experimental data, therefore this study proposes a random utility-based discrete choice model to determine the policy adopted by the users. For each policy and objective, we observe the significant variables from the following set: outdegree, distance to destination, network familiarity, congestion levels, and probability distribution of arc states. We observe that these factors also affect the chosen policy dependant on the trip objective. This new model demonstrates higher predictive power and provides better understanding of driver decision making.

Hence, the outcome of this research is an insight into the traveler decision strategy in response to real-time information under multiple travel objectives. With connected and autonomous vehicles (CAVs) on the horizon, optimal policy adoption and processing for large networks can become a reality. However, till human drivers remain active, the insights into human choice making under the influence of real-time information are valuable for network planning.

Future research includes extending these individual routing strategies towards equilibrium while accounting for dynamic flow evolution of the system by representing the extent of temporal and spatial dependence of arcs. Another potential extension involves risk profiling of the users, then separate analysis of learning patterns for risk averse and risk seeking individuals. Link correlations can also be considered in the optimal policy formulation, requiring solution for OSP with link correlations sub-problem. While a harder sub-problem, solution methods have been developed in the literature allowing optimal policy formulation for the correlations case.

CRediT authorship contribution statement

Ravi Venkatraman: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing - original draft, Visualization. **Stephen D. Boyles:** Conceptualization, Methodology, Formal analysis, Investigation, Resources, Writing - original draft, Writing - review & editing, Project administration, Funding acquisition. **Rachel James:** Validation, Data curation, Visualization, Writing - original draft. **Avinash Unnikrishnan:** Conceptualization, Methodology, Writing - original draft, Project administration, Funding acquisition. **Priyadarshan N. Patil:** Writing - original draft, Writing - review & editing, Visualization, Validation.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors would like to acknowledge that this work is based on the work done by Mr. Ravi Venkatraman for his Master’s Thesis and has not been published elsewhere. The authors would like to acknowledge the National Science Foundation for funding this work through Award Number 1069141/ 1157294 “Collaborative Research: Stochastic and Dynamic Hyperpath Equilibrium Models” and Award Number 1562109/ 1562291 “Collaborative Research: Non-additive Network Routing and Assignment Models.” This research was also partially supported by the U.S. Department of Transportation through the Data Supported Transportation Operations and Planning (D-STOP) Tier 1 University Transportation Center. Any error, mistake, or omission related to this research paper is the sole responsibility of the authors.

References

Abdel-Aty, M.A., Kitamura, R., Jovanis, P.P., 1997. Using stated preference data for studying the effect of advanced traffic information on drivers’ route choice. Transp. Res. Part C 5 (1), 39–50.
Adler, J.L., Blue, V.J., 1998. Toward the design of intelligent traveler information systems. Transp. Res. Part C 6 (3), 157–172.
Ahuja, R.K., Magnanti, T.L., Orlin, J.B., 1993. Network Flows: Theory, Algorithms, and Applications. Prentice-Hall Inc.

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

Ben-Elia, E., Avineri, E., 2015. Response to travel information: A behavioural review. Transp. Rev. 35 (3), 352–377.

Ben-Elia, E., Di Pace, R., Bifulco, G.N., Shiftan, Y., 2013. The impact of travel information's accuracy on route-choice. Transp. Res. Part C 26, 146–159.

Blaauw, G.J., 1982. Driving experience and task demands in simulator and instrumented car: a validation study. Hum. Factors 24 (4), 473–486.

Bowman, D.A., McMahan, R.P., 2007. Virtual reality: how much immersion is enough? Computer 40 (7), 36–43.

Boyles, S.D., 2009. Operational, supply-side uncertainty in transportation networks: causes, effects, and mitigation strategies (Ph.D. dissertation). The University of Texas at Austin.

Boyles, S.D., Rambha, T., 2016. A note on detecting unbounded instances of the online shortest path problem. Networks 67 (4), 270–276.

Boyles, S.D., Waller, S.T., 2007. In: The impact of nonlinear objective functions in online shortest paths presented at the 86th Annual Meeting of the Transportation Research Board, Washington, D.C.

Boyles, S.D., Waller, S.T., 2011. Optimal information location for adaptive routing. Networks Spatial Econ. 11 (2), 233–254.

Chen, P.S.-T., Mahmassani, H.S., 1993. Dynamic interactive simulator for studying commuter behavior under real-time traffic information supply strategies. Transp. Res. Rec. 1413, 12–21.

Cheung, R.K., 1998. Iterative methods for dynamic stochastic shortest path problems. Naval Res. Logist. 45 (8), 769–789.

Chorus, C.G., Walker, J.L., Ben-Akiva, M., 2013. A joint model of travel information acquisition and response to received messages. Transp. Res. Part C 26, 61–77.

De Lapparent, M., De Palma, A., Fontan, C., 2002. Non-linearities in the valuation of time estimates. Publication of: Association for European Transport.

Delle Site, P., 2018. A mixed-behaviour equilibrium model under predictive and static advanced traveller information systems (atis) and state-dependent route choice. Transp. Res. Part C 86, 549–562.

Dia, H., Panwai, S., 2007. Modelling drivers' compliance and route choice behaviour in response to travel information. Nonlinear Dyn. 49 (4), 493–509.

Dial, R.B., 1971. A probabilistic multipath traffic assignment model which obviates path enumeration. Transp. Res. 5 (2), 83–111.

Eiger, A., Mirchandani, P.B., Soroush, H., 1985. Path preferences and optimal paths in probabilistic networks. Transp. Sci. 19 (1), 75–84.

Erke, A., Sagberg, F., Hagman, R., 2007. Effects of route guidance variable message signs (VMS) on driver behaviour. Transp. Res. Part F 10 (6), 447–457.

Fan, Y., Kalaba, R., Moore, J., 2005. Arriving on time. J. Optim. Theory Appl. 127 (3), 497–513.

Fosgerau, M., Karlström, A., 2010. The value of reliability. Transp. Res. Part B 44 (1), 38–49.

Gao, S., 2005a. Optimal adaptive routing and traffic assignment in stochastic time-dependent networks Unpublished doctoral dissertation. Massachusetts Institute of Technology, Cambridge, MA.

Gao, S., 2005b. Optimal adaptive routing and traffic assignment in stochastic time-dependent networks (Ph.D. dissertation). Massachussets Institute of Technology.

Godley, S.T., Triggs, T.J., Fildes, B.N., 2002. Driving simulator validation for speed research. Accident Anal. Prevent. 34 (5), 589–600.

Hall, R.W., 1986. The fastest path through a network with random time-dependent travel times. Transp. Sci. 20 (3), 182–188.

Harms, I.M., Dijksterhuis, C., Jelijs, B., de Waard, D., Brookhuis, K.A., 2019. Don't shoot the messenger: traffic-irrelevant messages on variable message signs (VMSs) might not interfere with traffic management. Transp. Res. Part F 65, 564–575.

Jiang, G., Fosgerau, M., Lo, H.K., 2020. Route choice, travel time variability, and rational inattention. Transp. Res. Part B 132, 188–207.

Khani, A., 2019. An online shortest path algorithm for reliable routing in schedule-based transit networks considering transfer failure probability. Transp. Res. Part B 126, 549–564.

Khani, A., Boyles, S.D., 2015. An exact algorithm for the mean–standard deviation shortest path problem. Transp. Res. Part B 81, 252–266.

Khani, A., Boyles, S.D., 2015. An exact algorithm for the mean-standard deviation shortest path problem. Transp. Res. Part B 81, 252–266.

Lee, C., Abdel-Aty, M., 2008. Testing effects of warning messages and variable speed limits on driver behavior using driving simulator. Transp. Res. Rec. 2069 (1), 55–64.

Levinson, D., 2003. The value of advanced traveler information systems for route choice. Transp. Res. Part C 11 (1), 75–87.

Liu, S., Guo, L.-D., Easa, S., Yan, H., Wei, H., & Tang, Y.-N., 2020. Experimental study of day-to-day route-choice behavior: Evaluating effect of ATIS market penetration. J. Adv. Transp.

Lu, X., Gao, S., Ben-Elia, E., 2011. Information impacts on route choice and learning behavior in a congested network: experimental approach. Transp. Res. Rec. 2243 (1), 89–98.

Ma, T.-Y., Di Pace, R., 2017. Comparing paradigms for strategy learning of route choice with traffic information under uncertainty. Expert Syst. Appl. 88, 352–367.

Mahmassani, H.S., Huynh, N.N., Srinivasan, K., Kraan, M., 2003. Tripmaker choice behavior for shopping trips under real-time information: model formulation and results of stated-preference internet-based interactive experiments. J. Retailing Consumer Services 10 (6), 311–321.

Mandel, B., Gaudry, M., Rothengatter, W., 1994. Linear or nonlinear utility functions in logit models? The impact on German high-speed rail demand forecasts. Transp. Res. Part B 28 (2), 91–101.

Miller-Hooks, E., 2001. Adaptive least-expected time paths in stochastic, time-varying transportation and data networks. Networks 37 (1), 35–52.

de Moraes Ramos, G., Mai, T., Daamen, W., Frejinger, E., Hoogendoorn, S., 2020. Route choice behaviour and travel information in a congested network: Static and dynamic recursive models. Transp. Res. Part C 114, 681–693.

Morgan, J., Orzen, H., Sefton, M., 2009. Network architecture and traffic flows: Experiments on the pigou–knight–downs and braess paradoxes. Games Econ. Behav. 66 (1), 348–372.

Murthy, I., Sarkar, S., 1996. A relaxation-based pruning technique for a class of stochastic shortest path problems. Transp. Sci. 30 (3), 220–236.

Ng, L., Barfield, W., Mannering, F., 1995. A survey-based methodology to determine information requirements for advanced traveler information systems. Transp. Res. Part C 3 (2), 113–127.

Nie, Y., Fan, Y., 2006. Arriving-on-time problem: discrete algorithm that ensures convergence. Transp. Res. Rec. 1964 (1), 193–200.

Nie, Y.M., Wu, X., 2009. Shortest path problem considering on-time arrival probability. Transp. Res. Part B 43 (6), 597–613.

Nilsson, L., 1993. Behavioural research in an advanced driving simulator-experiences of the vti system. In Proceedings of the human factors and ergonomics society annual meeting (37, 612–616).

de Palma, A., Picard, N., 2005. Route choice decision under travel time uncertainty. Transp. Res. Part A 39 (4), 295–324.

Peeta, S., Yu, J.W., 2004. Adaptability of a hybrid route choice model to incorporating driver behavior dynamics under information provision. IEEE Trans. Syst. Man Cybern.-Part A 34 (2), 243–256.

Peeta, S., Yu, J.W., 2005. A hybrid model for driver route choice incorporating en-route attributes and real-time information effects. Networks and Spatial Economics 5 (1), 21–40.

Pinjari, A.R., Bhat, C.R., 2006. Nonlinearity of response to level-of-service variables in travel mode choice models. Transp. Res. Rec. 1977, 67–74.

Provan, J.S., 2003. A polynomial-time algorithm to find shortest paths with recourse. Networks 41 (2), 115–125.

Ramadurai, G., Ukkusuri, S.V., 2007. Dynamic traffic equilibrium: theoretical and experimental network game results in single-bottleneck model. Transp. Res. Rec. 2029, 1–13.

Ramos, G.d.O., Bazzan, A.L., da Silva, B.C., 2018. Analysing the impact of travel information for minimising the regret of route choice. Transp. Res. Part C 88, 257–271.

Rapoport, A., Gisches, E.J., Daniel, T., Lindsey, R., 2014. Pre-trip information and route-choice decisions with stochastic travel conditions: experiment. Transp. Res. Part B 68, 154–172.

Ringhand, M., Vollrath, M., 2018. Make this detour and be unselfish! Influencing urban route choice by explaining traffic management. Transp. Res. Part F 53, 99–116.

Sawik, B., Faulin, J., Pérez-Bernabeu, E., 2017. Multi-criteria optimization for fleet size with environmental aspects. Transp. Res. Proc. 27, 61–68.

Sawik, B., Faulin, J., Pérez-Bernabeu, E., 2017. A multicriteria analysis for the green vrp: a case discussion for the distribution problem of a spanish retailer. Transp. Res. Proc. 22, 305–313.

Shahabi, M., Unnikrishnan, A., Boyles, S.D., 2015. Robust optimization strategy for the shortest path problem under uncertain link travel cost distribution. Comput.-Aided Civil Infrastruct. Eng. 30, 433–448.

Small, K., Winston, C., Yan, J., 2005. Uncovering the distribution of motorists' preferences for travel time and reliability. Econometrica 73, 1367–1382.

Srinivasan, K.K., Mahmassani, H.S., 2000. Modeling inertia and compliance mechanisms in route choice behavior under real-time information. Transp. Res. Rec. 1725 (1), 45–53.

Szeto, W.Y., Solayappan, M., Jiang, Y., 2011. Reliability-based transit assignment for congested stochastic transit networks. Comput.-Aided Civil Infrastruct. Eng. 26 (4), 311–326.

Tang, Y., Gao, S., Ben-Elia, E., 2017. An exploratory study of instance-based learning for route choice with random travel times. J. Choice Model. 24, 22–35.

Unnikrishnan, A., Waller, S.T., 2009. User equilibrium with recourse. Netw. Spatial Econ. 9, 575–593.

Venkatraman, R., 2013. Adaptive routing behavior with real time information under multiple travel objectives (Master's thesis). The University of Texas at Austin, TX, USA.

Waller, S.T., Ziliaskopoulos, A.K., 2002. On the online shortest path problem with limited arc cost dependencies. Networks 40 (4), 216–227.

Wijayaratna, K.P., Dixit, V.V., 2016. Impact of information on risk attitudes: implications on valuation of reliability and information. J. Choice Model. 20, 16–34.

Wijayaratna, K.P., Dixit, V.V., Denant-Boemont, L., Waller, S.T., 2017. An experimental study of the online information paradox: does en-route information improve road network performance?. Plos One 12, (9) e0184191.

Wu, X., 2015. Study on mean-standard deviation shortest path problem in stochastic and time-dependent networks: a stochastic dominance based approach. Transp. Res. Part B 80, 275–290.

Zhao, W., Quddus, M., Huang, H., Lee, J., Ma, Z., 2019. Analyzing drivers' preferences and choices for the content and format of variable message signs (vms). Transp. Res. Part C 100, 1–14.