Congestion Information and Efficiency: An Experiment*

Charles N. Noussair† Liang Qiao‡

July 24, 2026

Abstract

Users are now typically able to access historical data in traffic, restaurant, and retailing contexts. This availability of information can be viewed as a means to reduce congestion. However, scientific support for this notion is weak. We investigate, using a laboratory experiment, whether an intermediate level of information provision might reduce congestion. We also study how the effect of making more information available changes over time. We show that providing all users with information is counterproductive in the short run. In the long run, users' behavior adjusts so that providing some or all individuals with information leads to more efficient outcomes than when no information is made available. The incidence of group over-and under-reaction to congestion cost differences is only mildly responsive to the information level, a pattern consistent with Small Sample models, which do not assume sequential adjustment.

Keywords: Coordination Game, Information Use, Information Value, Learning, Experiment.

JEL: C72, C92, D83

*The authors would like to express their appreciation to Xiduo Chen, Yan Chen, John Duffy, Amanda Friedenberg, Antonio F. Galvao, David Gill, Xiao Han, Nobuyuki Hanaki, Yuge Hao, Haoran He, Shuya He, Jin Li, Tianyi Li, Hang Qi, Julian Romero, Zhewei Song, Hang Wu, Wenhao Wu, Wendan Zhang, Daniel Zizzo, three anonymous referees and the anonymous Associate Editor, for helpful comments. We also thank participants at seminars at the University of Arizona, Nankai University, Tianjin University, Harbin Institute of Technology, the 2021 North American Regional Economic Science Association (ESA) Meeting, the 2022 Virtual Asia Pacific ESA Meeting, the Chinese Economists Society 2022 Annual Conference, the 2023 Tsinghua Conference on Behavioral, Experimental and Theoretical Economics, and the UC Riverside Conference in honor of Amnon Rapoport for helpful comments and discussions. We would like to thank the Economic Science Laboratory at the University of Arizona for funding and the Southwest University of Finance and Economics for allowing us to use their research facilities.

†Department of Economics, University of Arizona, Tucson, USA. cnoussair@email.arizona.edu

‡Department of Economics, Nankai University, Tianjin, China. liangqiao@nankai.edu.cn

1 Introduction

Individuals often employ waiting time or congestion information gleaned from past experience when making decisions, such as where to shop or dine, or how to travel to a given destination. This has become easier and more widespread in recent years, as numerous apps whose purpose is to help individuals avoid delays and congestion have become available. These apps use historical information as inputs into their predictive models. For example, Advanced Traveler Information Systems (ATIS), such as Google Maps or Gaode Maps, provide drivers with predicted travel time information on different routes. When choosing a restaurant, diners can use apps that show how crowded a restaurant tends to be at each time of day to assist them in avoiding long wait times. When deciding which store to purchase a product from, apps are available that display information on how crowded the store typically is on different days and times.

Is having this data available efficiency-enhancing? At first glance, it would appear likely. Such information allows drivers to avoid routes and times that tend to be congested and to use more efficient alternatives. This behavior could reduce average travel times. A restaurant could make more efficient use of its seating capacity if more diners arrive when the restaurant tends to be less crowded. However, one can also make a plausible argument for why universal access to congestion information would not be desirable. If too many people respond simultaneously to congestion information by changing routes, they may merely move the congestion elsewhere[1]. Similarly, information regarding available restaurant capacity may lead to a rush of diners arriving at traditionally off-peak periods and may actually increase waiting times.[2]

Indeed, academic research using agent-based modeling has shown that the effect of traffic information on congestion is not necessarily beneficial. Cabannes et al. (2018a), in their simulations of traffic on I-210 in Los Angeles, find that travel time can be increasing in the percentage of drivers using traffic apps. The tension between self-interested behavior and socially optimal outcomes is laid bare in the work of Cabannes et al. (2018b), who find that app users spend less time traveling than non-app users, but that the average travel time for app users increases with the number of users. This

[1]For a similar argument about overreaction, see Ben-Akiva et al. (1991).

[2]Indeed, dissatisfaction with the traffic congestion on small routes running through residential areas caused by the use of traffic apps has become a widespread problem that is routinely discussed in the popular press (Moore, 2016; Madrigal, 2018; Macfarlane, 2019). This new traffic has caused cities to try to strategically distort the output of the apps' algorithms (Weise, 2017). Similarly, restaurants are prone to stop transmitting waiting time information when the times get too long, distorting the output of the apps.

suggests that information has a greater value when fewer other individuals have it. As we discuss later in this paper, well-known learning models also suggest that giving only some people the information is preferable to making it available to all users, and that information is more valuable to an individual when fewer other drivers also have it.

An interesting though small experimental literature has considered the value of historical information in facilitating more efficient route choice decisions. This literature has typically found that providing all individuals with historical information about congestion does not increase traffic efficiency (Selten et al., 2007; Rapoport et al., 2014; Wijayaratna et al., 2017), with one study (Liu et al., 2020) suggesting that it is preferable to provide only a fraction of users rather than none or everyone with information. Our study borrows design elements from a number of these studies, and this literature is discussed in Section 2.

We use a laboratory experiment to compare congestion patterns under different levels of information provision. Unlike the prior literature, the focus of our analyses is the learning dynamics at work and the transition to the long-run congestion level. That is, we are interested in short-run behavior as individuals begin to acquire experience as well as long-run behavior. We frame our experiment in a route choice context and adopt a simplified version of the environment studied by Selten et al. (2007). There are a number of drivers, each of whom can choose between two routes to reach their destination. One of the routes becomes congested more quickly than the other as traffic increases. Congestion reduces the payoff to drivers. The game is played repeatedly. Drivers who are informed know the number of drivers who have taken each route in the previous period before making their choice. Those who are uninformed do not have this knowledge. The experiment includes four treatments that vary the percentage of informed drivers.[3]

The stage game has many Nash equilibria, and the repeated game has a very large number of subgame perfect equilibria. The equilibria are all similar in that they all have, in expectation, the same number of drivers on each route. Although prior experiments with similar games tend to observe convergence to the equilibrium proportions, the

[3] The fact that there are only two routes and we have a discrete period structure means that in the traffic app context, the environment most closely corresponds to a situation in which there are two roads to reach a destination, and switching between them is impractical once one is underway, so the route choice decision is made once with the traffic information available at the beginning of one's journey. However, the correspondence to the traffic app context is imperfect since traffic apps use historical information along with other variables in predicting traffic on routes. Thus, the environment is not meant to simulate the manner in which traffic apps operate but rather is guided by the previous literature and the isolation of the effect of historical information.

process whereby this takes place is so poorly understood, that it has been termed as "magic".[4] This convergence presumably takes place through the learning dynamics at work, and our approach allows us to consider whether a model of learning can account for this behavior. Consideration of equilibria only does not yield any predictions about the dynamics at work or the role of information. Therefore, rather than basing our hypotheses on equilibrium analysis, our approach is to formulate hypotheses based on a learning model that explicitly models the process whereby individuals update the route they choose. Specifically, we employ *Experience Weighted Attraction-lite*(EWA-lite) (Ho et al., 2007), a simplified one-parameter version of EWA (Camerer and Ho, 1999).

Our experiment has two parts. The first part includes three treatments (No, Partial, and Full) that vary the number of drivers with traffic information. Under the No treatment, all players are uninformed about past traffic on each route. Under the Partial treatment, one-half of the players are informed about prior traffic, and the other half are uninformed. Finally, in the Full treatment, all players are informed. By comparing the congestion level among the three treatments, we are able to observe the impact of information levels on efficiency. An endogenous information condition characterizes the second part of the experiment: players decide whether to pay to be informed. Their valuation of past traffic information is elicited using a strategy method protocol. This part of the experiment allows us to determine the value individuals place on traffic information and how it depends on how many others also have access to it.

The experimental results show that whether having past traffic information available is beneficial for overall efficiency depends on the time period that is considered. In the short run, interpreted as the first ten periods of play in our experiment, it is best to have nobody or 50% of drivers with access to information. In the long run, defined as the last ten periods of a treatment, Full information adoption is significantly better than Partial information adoption, and both Partial and Full adoption are significantly better than No adoption. The cost from congestion decreases over time under Partial and Full, but not under No Information. The underperformance of Full in early periods

[4] Kahneman (1988) writes: "Observing the regularity of behavior in these markets was a bewildering experience; to a psychologist, it looked almost like magic. The bewilderment was not eased by the debriefing conversations in which we engaged participants after the experiment. They described a large variety of strategies and expectations as guiding their behavior. Most of the strategies were completely unfounded (as they must have been, since the equilibrium effectively precluded any successful strategy). Furthermore, in at least one case there was evidence that the strategy which the individual described did not in fact account for his choices. The equilibrium outcome (which would be generated by the optimal policies of rational players) was produced in this case by a group of excited and confused people, who simply did not seem to know what they were doing."

is accompanied by a relatively high rate of underreaction (defined here as a failure to change the choice of one's route to the one that would have been cheaper last period) to new information compared to Partial, a difference that disappears over time. The lower the number of other users who are informed, the higher the value individuals place on traffic information.

The EWA-Lite model has mixed success in predicting outcomes. It correctly predicts that Partial Information is the most efficient condition in early periods. However, it overpredicts the percentage of drivers choosing route a under No Information as well as the variance in route choice in the late periods. The presence of underreaction and its insensitivity to treatment suggests that the assumption of sequential adjustment, the immediate response to outcomes from the last period, within the EWA-lite model might be a source of inaccuracies. We consider whether these two inaccuracies can be explained with the application of *Small Sample (SSM) Models* (Erev and Roth, 2014; Yakobi et al., 2020). We consider two SSM models. The first is the *Sample of Five Model* (SSM5), under which an individual is assumed to randomly make five draws from prior experience to calculate the payoff from the relevant options. The second is the *Sample of Nine or Less Model* (SSM1-9), where the number of prior periods sampled varies uniformly between 1 and 9. We find that the SSM models do capture a number of patterns in the data that EWA-lite does not.

The remainder of the paper is organized as follows. Section 2 discusses some of the related literature. Section 3 introduces the model and reports simulation results. Section 4 describes the design and implementation of the experiment. In Section 5, we present and discuss the main experimental results. Finally, Section 6 concludes with a summary of the main findings and some thoughts about policy implications.

2 Related Literature

This paper is closely related to five strands of literature. For brevity, we summarize related previous studies in Table 1.

The first strand consists of previous experimental studies of transportation route choice. Early studies concerning exogenous uncertainty about conditions on different routes (Ben-Elia et al. (2013), Rapoport et al. (2014) and Wijayaratna et al. (2017)) tend to show that travel information helps individual decision makers cope with the uncertainty. Other studies focus on the congestion resulting from interacting travelers choosing routes (see, e.g., Rapoport et al. (2009), Lu et al. (2011), Mak et al. (2015)).

The most closely related experimental studies to ours are those of Selten et al. (2007), Liu et al. (2020), and Ashraf et al. (2023). These studies vary the amount of information available to drivers. In all cases, the information provided consists of the past history of congestion on each route, as in our study.[5]

Second, the present paper has a close relationship with market entry games, which typically exhibit a large multiplicity of equilibria (Erev and Rapoport, 1998). The two types of games share the property that all pure-strategy equilibria involve the same number of entrants (interpreted in our environment as the number of cars taking a given road). The unique symmetric mixed strategy equilibrium leads to a slightly higher average cost than in pure-strategy equilibria. Experimental studies have shown that in entry games with this structure, behavior converges to the equilibrium proportion over time.

The third literature related to our study is the experimental literature on learning. A variety of learning models have been proposed, and we apply two such models that have received empirical support in other environments. One of the models is a simplified version of the Experience-Weighted Attraction (EWA) model, introduced by Camerer and Ho (1999), which hybridizes belief-based models (see Fudenberg and Levine (1995), Cheung and Friedman (1997)) and choice-reinforcement models (e.g., Roth and Erev (1995), Erev and Roth (1998)) and includes versions of each of them as special cases. To address the criticism that the EWA model has too many free parameters, Ho et al. (2007) introduce a one-parameter learning theory, EWA-lite (also known as self-tuning EWA). We extend EWA-lite to an environment with limited feedback and use it to make predictions. The second model that we consider is the Small Sample Model (Erev and Roth, 2014). Agents are assumed to take a small sample of past plays of the task they face, compute the average payoff for that sample, and choose the option with the highest average. Small sample models have been shown to predict well in several environments (Erev and Barron, 2005; Yakobi et al., 2020; Plonsky et al., 2015; Erev et al., 2017, 2023). Here, we employ SSMs ex-post to explain some features of the experimental data.

Fourth, this paper is related to the study of Kremer et al. (2014), which examines a social planner's problem of computing and implementing an optimal information dis-

[5] Zhang and Verhoef (2006) present a theoretical model, based on Emmerink et al. (1996), that combines a transport market and a monopolistic market for exogenous traffic information. Denant-Boèmont and Petiot (2003), in a study focused on departure time and traffic mode rather than route choice, allow users to choose whether or not they purchase travel time information. However, to our knowledge, no previous experimental studies explore the dependence of the value of traffic information on how widely available the information is.

closure policy in a sequential game. Their findings share a similar flavor to ours in that disclosing all information can be sub-optimal.

Finally, this paper applies the strategy method (Selten, 1967) to elicit contingent strategies in a transportation game. We do so in a similar manner as Fischbacher et al. (2001) and Fischbacher and Gächter (2010). The central idea is that participants make two types of decisions, an "unconditional" and a vector of "conditional" decisions. This protocol ensures that both unconditional and conditional decisions are incentivized. We use this augmented strategy method to elicit a valuation for the traffic information and determine how this valuation depends on how many other drivers also acquire the information.

Table 1: Summary of related literature

Study	Main Features of Study	Key Findings	Relation to this Study
Selten et al. (2007)	Two (endogenous) information treatments (`No` and `Full`). Informed drivers learn travel time on the unselected route in the last period of a two-route choice scenario.	The average route choice is close to the symmetric mixed strategy equilibrium frequency, though with fewer individual route changes than in the equilibrium. Information reduces traffic fluctuations. A Payoff Sum Model explains their data well.	They have a transportation network structure similar to ours. We extend the design by introducing the `Partial` information treatment and consider the effect of information on payoffs. We also estimate learning models on the data.
Wijayaratna et al. (2017)	Two (exogenous) information treatments (`No` and `Full`) vary whether the state of one link at a node in a three-route Breass network (Braess, 1968) is public or not.	Providing public information at one node in the network serves to increase congestion, which is referred to as an Online Information Paradox. The data move toward analytic solutions over time.	The results have a similar flavor to ours, which is that more information is not necessarily better.
Rapoport et al. (2014)	Two (exogenous) information treatments (`No` and `Full`). Information is about travel conditions ("good" or "bad") on each route in a two-route structure.	No significant differences between the `No` and `Full` Information conditions; EWA describes the broad patterns well, but EWA predicts more switches than are actually observed.	They have a transportation network structure similar to ours. They did not compare the performance of EWA to any other model.
Liu et al. (2020)	Five (endogenous) information treatments that vary the fraction of individuals that have information about the travel time of each route in the last period in the three-route Braess network.	The total travel time in the system as a whole increases with an increase in information acquisition.	They include `Partial` information treatments. They conduct only one session under each treatment, so we view their study as being suggestive rather than conclusive.

Ashraf et al. (2023)	Two experiments vary who receives additional information about the unchosen road in the binary route choice game of Selten et al. (2007). The first varies the number of informed individuals based on the frequency of changing their route, while the second varies the delay before everyone receives information.	The different interventions have little effect on the aggregate payoff.	They have a transportation network structure similar to ours. The design includes a `Partial` information treatment.
Zhang and Verhoef (2006)	Presents a theoretical model, based on Emmerink et al. (1996), that combines a transport market and a monopolistic market for exogenous traffic information.	Increasing information penetration leads to a negative externality to existing informed users.	We explore the dependence of the value of traffic information on how widely available the information is by conducting an experiment.
Denant-Boèmont and Petiot (2003)	Focus on individual decisions of departure time and traffic mode. They endogenize the level of information by allowing users to choose whether or not to purchase more or less detailed travel time information.	The willingness to pay is greater for more detailed information. Players are more likely to buy information when the variance in payoffs is higher (the payoff depends on both the travel time and the difference between the actual and target arrival time).	Same as above.
Camerer and Ho (1999)	Introduces the EWA model, which hybridizes belief-based and choice-reinforcement models.	EWA predicts human behavior in a variety of games well.	EWA is the foundation of EWA-lite, which we used in the paper.
Ho et al. (2007)	Introduces a one-parameter learning theory, EWA-lite, to reduce the number of parameters in EWA by fixing the initial parameters at specific values and replacing estimatable parameters with self-tuning functions.	Although EWA-lite is econometrically simpler than EWA, it does as well as EWA in producing accurate predictions of behavior in a wide class of games.	We extend the EWA-lite model to an environment with limited feedback and use it for predictions.
Erev and Roth (2014)	Explores the conditions under which mechanisms that rest on the rationality assumption are likely to be successful. They discuss how to facilitate market design via promoting human learning.	When people quickly learn to maximize expected returns, Small Sample models predict choices well.	We use Small Sample models to try to explain some patterns in our experimental data.
Kremer et al. (2014)	Examines a social planner's problem of computing and implementing an optimal information disclosure policy. Does so in an environment where agents arrive sequentially and choose one action from a set of actions with unknown payoffs.	Disclosing all, as well as not disclosing any, information is not the optimal policy.	Our paper obtains results of a similar flavor in that we observe that revealing information to all can be suboptimal in a simultaneous repeated game in the early periods.

Fischbacher et al. (2001)	Design a protocol that has participants make one "unconditional" and a vector of "conditional" decisions in a public goods game based on the strategy method (Selten, 1967).	Approximately one-third of subjects can be classified as free riders, whereas 50% are conditional cooperators.	We use a similar augmented strategy method to elicit a valuation for the traffic information and determine how this valuation depends on how many other drivers also acquire the information.
Fischbacher and Gächter (2010)	Explore why voluntary contributions to public goods decline over time.	Cooperation declines due to preferences of conditional cooperators for contributing somewhat less than others.	Same as above.

3 The Model, Simulations, and Hypotheses

This section is organized as follows. Section 3.1 describes the coordination game we created for our experiment and its Nash equilibria. In Section 3.2, we briefly describe the EWA-lite learning model that we consider and report simulation results for the model in our environment. Testable hypotheses derived from the simulations are reported in Section 3.3.

3.1 The One-shot Transportation Game

Consider N risk-neutral players $i = 1, ..., N$, who move simultaneously in a one-shot game. Player i chooses her action $s_i \in \{0, 1\}$, which we think of as choosing a route to reach a destination. $s_i = 1$ if i chooses route a, and $s_i = 0$ if i chooses route b. A driver taking route a pays a cost of $C(a) = k \times n_a$ while a driver taking route b pays $C(b) = m \times n_b$, where k and m are constants, and n_a and n_b are the number of players choosing routes a and b, respectively. Reaching the destination yields a value v_i. Nature independently draws the travel value v_i from a uniform distribution $v_i \in [v_L, v_H]$ which is common to each player i.[6] v_i is not revealed to any player, including i herself, before they choose strategies.[7] In the experiment, $N = 6, v_L = 6, v_H = 16, k = 1$ and $m = 2$. Route b is twice

[6]In the repeated game we conduct in our experiment, v_i is independently drawn each period for each i. This means that learning v_i in one period provides no information about v_i in the next period.

[7]Since the travel value is random, a player is unable to infer perfectly the number of players choosing one specific route after observing the realized payoff. This is critical because our experimental design involves a No Information treatment, in which an individual cannot determine how many drivers have taken each route in the past. A random travel value corresponds to many scenarios in life. For example, you may not know how enjoyable a party will be until you get there, or you may not know how valuable

as costly to take as route a for a given number of commuters on that route, and can be thought of as a road with lower capacity. Note that $v_L \geq \min\{k \times n_a, m \times n_b\}, \forall n_a, n_b$, so that there will always exist a route to the destination which has a non-negative payoff. The payoff function of player i is the difference between the travel value and the traffic cost:

$$\pi_i(s_1, ..., s_N) = \begin{cases} v_i - \sum_{k=1}^{N} s_k \text{ if } s_i = 1, \\ v_i - 2\sum_{k=1}^{N}(1 - s_k) \text{ if } s_i = 0. \end{cases}$$

There are two types of players: informed (Inf) and uninformed ($UInf$). After the one-shot game, both types of players learn their individual payoffs. Only an informed player, however, can observe how many players have chosen each route at the end of each play of the game. Because the information arrives after all strategies have been chosen, whether a player i is informed of v_i or of the number of players selecting each route at the end of the game does not affect the Nash equilibria.

A number of pure strategy Nash equilibria exist. Given others' strategies, player i's payoff is $\pi_i(1, s_{-i}) = E(v_i) - 1 - \sum_{k \neq i} s_k$ if i chooses $s_i = 1$; player i's payoff is $\pi_i(0, s_-i) = E(v_i) - 2 - 2\sum_{k \neq i}(1 - s_k)$ if i chooses $s_i = 0$. The best response for player i is therefore $s_i = 1$ if and only if $\pi_i(1, s_{-i}) \geq \pi_i(0, s_{-i})$. In other words, if at least two others select the strategy $s = 0$, i is better off picking $s_i = 1$; otherwise, player i is better off choosing $s_i = 0$. A pure strategy NE consequently implies that four players choose $s = 1$ and two players select $s = 0$. Since the players are symmetric, there are $6!/(4! \times 2!) = 15$ pure-strategy Nash equilibria, each a distinct mapping from the set of players to the set of routes that results in the same route split. All players in any of these equilibria have the same cost of 4, which is the lowest possible average cost. The outcome is therefore socially optimal.

There is also a unique symmetric mixed strategy Nash equilibrium and many asymmetric mixed strategy equilibria. We omit the discussion of asymmetric equilibrium since coordinating on an asymmetric mixed strategy equilibrium in the absence of direct communication is very unlikely. Nevertheless, the symmetric mixed strategy NE is a plausible benchmark for an experiment. Suppose each player i chooses $s_i = 1$ with probability α. Given others' strategies, player i's payoff will be $\pi_i(1, \alpha) = E(v_i) - 1 - 5\alpha$ if $s_i = 1$ is selected; his payoff will be $\pi_i(1, \alpha) = E(v_i) - 2 - 2 \times 5(1 - \alpha)$ if $s_i = 0$ is chosen. In order for player i to mix, the expected payoff of choosing each route must be

a class session is until you attend.

equal, which requires $\alpha = 11/15$.[8] Thus, each player independently chooses $s = 1$ with probability 11/15 in the symmetric mixed strategy NE. In this equilibrium, the expected individual cost is 14/3, higher than in the social optimum or in the asymmetric pure strategy equilibria discussed earlier.

3.2 The EWA-lite Learning Model

In our experiment, players can learn from each play of the stage game from the information they receive and update their strategies accordingly. The EWA-lite learning model (Ho et al., 2007) is well-suited to our environment, and our approach is to employ the model as our source of hypotheses.

The EWA-lite model updates two sets of variables for each player after each period: the *observation-equivalents* of experience and the *attraction* of each action. Following Ho et al. (2007), we initialize experience at one observation equivalent and each action's attraction at one. Additionally, we specify two self-tuning functions that govern experience accumulation and strategy reinforcement to update experience and attractions after each period. The only free parameter is the response sensitivity, λ, which regulates choice probabilities. It measures how sensitive a person's choice behavior is to differences between the attractions of the available actions. $\lambda = 0$ corresponds to randomly choosing routes with equal probability. As $\lambda \to \infty$ increases, the behavior converges to error-free decision-making, where players always select the action with the highest attraction.

We add two assumptions to fit EWA-lite into our setting. First, the action of players other than i is proxied by the number of other players choosing route a. Second, the uninformed players have a uniform belief over all possible outcomes of others' actions that could lead to their observed payoffs. See Appendix A for details.

To formulate point predictions, we set the sensitivity parameter at $\lambda = 5.87$, the estimate reported in Table 4 of Ho et al. (2007), who estimated λ by pooling data across seven games.[9] We assume all initial attractions are equal to one, which leads to an equal probability of choosing each route in period 1.

Figure 1 displays the simulation results over 40 periods. The three panels in the figure display the average cost, as well as the average and the standard deviation of the number of players choosing route a, respectively. The precise average values and

[8]The probability of selecting route a in the symmetric mixed strategy equilibrium for N players is $2/3 + 1/[3(N-1)]$, and converges to $2/3$ as $N \to \infty$.

[9]Appendix A.3 contains simulations of the EWA-lite model for different levels of the sensitivity parameter, λ.

standard deviations in the first 10 and last 10 periods are given in Table 2.

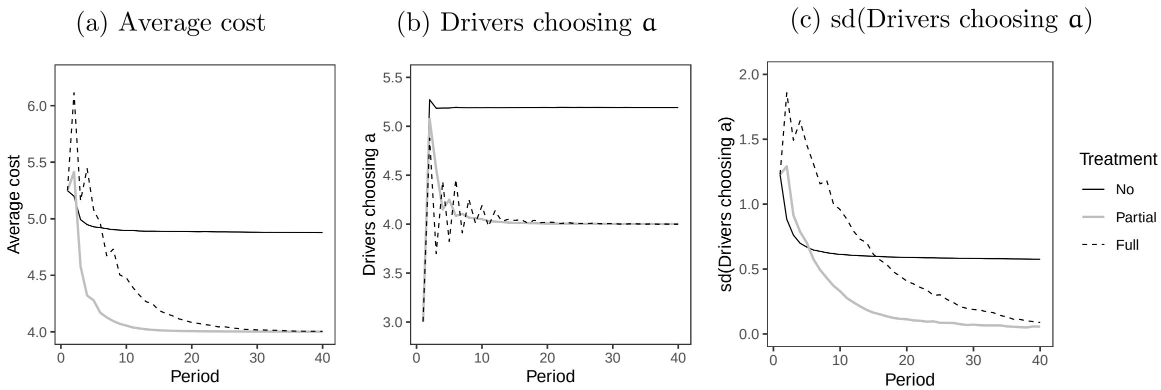

Figure 1: Dynamics of cost and traffic patterns under the EWA-lite model. $\lambda = 5.87$ is assumed. 10,000 simulations of a six-player group are conducted under each of the No, Partial and Full treatments. (a) plots the average individual cost in the simulations over a 40-period horizon. The minimum possible cost is 4, and the average cost converges downward toward this level over time in Partial and Full, with Partial converging faster; (b) shows the average number of players choosing route a, the higher capacity route, in a group. The efficient level is 4 and Partial and Full converge to an average of 4 over time, while too many drivers choose route a under No; (c) presents the standard deviation of the number of players choosing route a. This measure of dispersion converges to 0 asymptotically under Partial and Full, while converging to a positive level of approximately 0.6 under No.

Figure 1a shows that the cost is lowest under the Partial treatment, with the advantage particularly pronounced from periods 5 to 15. The average costs over the 40 simulated periods under No, Partial, and Full are 4.91 (sd=0.70), 4.11 (sd=0.09), and 4.32 (sd=0.38), respectively. The corresponding average costs during the first ten periods are lowest in Partial, while they are comparable in the other two conditions. During periods 31 to 40, the average costs under Partial and Full Information are almost indistinguishable, but considerably lower than the No condition.

Figure 1b illustrates the dynamics of the number of players choosing the larger route a. The number is higher under the No treatment than in the other two conditions. Specifically, the average number of players choosing route a over the 40 periods under No, Partial, and Full are 5.14 (sd=0.55), 4.04 (sd=0.05), and 4.03 (sd=0.04), respectively. In early periods, there is considerable oscillation in the number of drivers taking route a under Full, reflecting some overreaction to the relative costs of the two roads in the prior periods. Figure 1c reveals that in the later periods, the standard deviation of drivers taking route a remains considerable under No, but becomes very small under Partial and Full.

We conducted additional simulations, described in Appendix A.3.2, to compare the expected payoffs of individuals with or without information, conditional on how many other drivers have information. The expected payoffs can be compared to yield a willingness-to-pay for information. Under conditions where zero to five others are informed players, the implied information values are 1.79 (sd=1.46), 0.50 (sd=0.48), -0.03 (sd=0.13), -0.10 (sd=0.22), -0.16 (sd=0.29), and -0.23 (sd=0.36), respectively. This means that drivers should only be willing to pay for information if at most one other person is informed as well.

3.3 Hypotheses

Table 2: EWA-lite simulation results

Model	Info.	Avg. Cost Periods 1-10	Avg. Cost Periods 31-40	Num. Routa A Periods 1-10	Num. Routa A Periods 31-40
EWA-lite $\lambda = 5.87$	No	4.99 (0.71)	4.88 (0.72)	4.98 (0.51)	5.19 (0.58)
	$\mathsf{Partial}$	4.44 (0.35)	4.00 (0.02)	4.14 (0.15)	4.00 (0.03)
	Full	5.04 (0.97)	4.01 (0.12)	4.07 (0.09)	4.00 (0.02)

Notes: The average values of cost and number of members choosing route a across the first ten and last ten periods in 10,000 simulations are reported. Standard deviations are given in parentheses. A standard deviation is calculated by averaging the cost of the members or counting the number of members choosing route a of each group in each period, averaging over the 10 (periods) for each group, and then computing the between-group standard deviation of these averages.

Our hypotheses are derived from predictions of the EWA model. Support for the hypotheses, which would not be anticipated by an equilibrium analysis, would be evidence for the notion that the data can be described by behavioral models of learning. We treat two simulation sample means as different in formulating our hypotheses if their difference is greater than 0.5 times the pooled sample standard deviation. 0.5 standard deviations corresponds to a medium-sized treatment effect (see Cohen (1988)).

We first consider patterns in average cost among treatments and over time predicted by EWA-lite. Let $C(T)$ denote the average cost under treatment T. In the early periods 1 - 10, EWA-lite predicts that $C(\mathsf{Partial}) < C(\mathsf{Full}) = C(\mathsf{No})$. In the late periods 31 - 40, EWA-lite predicts $C(\mathsf{Full}) = C(\mathsf{Partial}) < C(\mathsf{No})$. Let the subscripts E and L indicate early (1-10) and late (31-40) periods, respectively. Comparing Periods 1 - 10 with Periods 31 - 40, EWA-lite predicts that $C_L(\mathsf{No}) = C_E(\mathsf{No})$, $C_L(\mathsf{Partial}) < C_E(\mathsf{Partial})$, and $C_L(\mathsf{Full}) < C_E(\mathsf{Full})$. These average cost relationships form the basis for our first hypothesis.

Hypothesis 1: a) In the early periods, providing information to one-half of drivers leads to the lowest average congestion cost. b) The congestion cost decreases over time when some or all drivers have information, but not when no drivers have it. c) In the late periods, providing some or all drivers with information results in lower costs than when none have it.

We denote the number of drivers taking route a under treatment T as $\mathsf{A(T)}$. In periods 1 - 10, EWA-lite predicts that $\mathsf{A(No)} > \mathsf{A(Partial)} > \mathsf{A(Full)}$. In periods 31 - 40, EWA-lite predicts $\mathsf{A(No)} > \mathsf{A(Partial)} = \mathsf{A(Full)}$. As to comparisons between early and late periods, EWA-lite predicts that $\mathsf{A_L(No)} > \mathsf{A_E(No)}$, $\mathsf{A_L(Partial)} < \mathsf{A_E(Partial)}$, and $\mathsf{A_L(Full)} < \mathsf{A_E(Full)}$. The predicted route choice patterns under EWA-lite are summarized as Hypothesis 2.

Hypothesis 2: a) In the early periods, more drivers choose route a *under* No *than in the other two treatments, followed by* $\mathsf{Partial}$*, then by* Full*. b) In the late periods, the number of drivers on route* a *is largest under* No*, followed by the other two treatments. c) The number choosing* a *under the* No *treatment increases from the early to the late periods, while it decreases under the* $\mathsf{Partial}$ *and* Full *treatments.*

We advance a third Hypothesis that applies to Part 2 of our experiment, in which we measure the willingness-to-pay for information as a function of how many others also have it. EWA-lite predicts a positive willingness-to-pay for the information if and only if zero or only one other driver has it. See Section 4.2 for details. This is the basis of our third hypothesis.

Hypothesis 3: The willingness-to-pay for information is strictly positive when either zero or one other driver has the information, and is zero otherwise.

4 The Experiment

The experiment consisted of 24 sessions, conducted at ..... Laboratory at ..... University, USA, and ..... University, China. Both are large public universities. There were 288 participants, 72 from ..... University, USA and 216 from ..... University, China. All procedures received prior approval from the Institutional Review Board at ..... University, USA. The experiment was computerized and used the zTree platform (Fischbacher, 2007).[10]

Each experimental session had 6 (4 sessions), 12 (16 sessions), or 18 (4 sessions) participants taking part, all in the role of drivers. Participants remained in the same group of six for the duration of the experiment. A session lasted, on average, for about 100 minutes. All participants were undergraduate students registered in the local subject pool. 33.3% were male.[11]

Each session of the experiment consisted of four 40-period segments, followed by a short questionnaire. The four treatments were implemented in a within-subject design so that each subject took part in each treatment. The same group of six participants went through the four treatments together in a session. In Segments 1 - 3, which constituted Part 1 of the session, the *No*, *Partial*, and *Full* treatments were conducted in random order, with each treatment appearing only once. Specifically, zero, three, and six participants in a group were informed about past traffic congestion under the No, Partial, and Full treatments, respectively. The rest remained uninformed for the entire 40-period segment. Part 2, which consisted of the *Endogenous* treatment, allowed individuals to choose whether or not to acquire information.

To help participants become more comfortable with the task, we included 20 practice periods preceding Part 1, played under the same conditions as Segment 1. Before Part 2, there was also a practice period that allowed participants to get used to the Part 2 task. During the practice periods, each participant was grouped with five computerized bots.[12] Interacting with bots enabled participants to become familiar with the experiment while minimizing contamination of the data from observing the behavior of other participants.

4.1 Part 1: The No, Partial, and Full Treatments

The No, Partial, and Full treatments constituted Part 1 of the experiment. Each treatment was in effect for 40 periods (not including the 20 practice periods). The traffic task in each period can be described as follows. As shown in Figure 2, there are

[10]IRB approval was not required at ..... University, China. Our procedures followed the applicable Chinese laws on conducting research with human subjects.

[11]The gender composition of participants differed in the two locations. 61.1% of participants were male at ..... University in the U.S. and 24.1% at ..... University in China, primarily reflecting the differences in the gender composition of students studying economics and business.

[12]The bots selected each route with probability 0.5 independently in each practice period of Part 1. They chose to acquire information with a probability of 0.5 in the practice period of Part 2. Participants were not informed about the bots' strategy in these practice periods.

two routes that can be taken between the origin and the destination, called routes $\mathtt{a}$ and $\mathtt{b}$. Drivers decide simultaneously which route to take. However, congestion is increasing in the number of drivers taking a route and this congestion imposes a cost on drivers. Specifically, the congestion cost a driver incurs by taking route $\mathtt{a}$ equals the number of drivers on route $\mathtt{a}$. If the driver had chosen route $\mathtt{b}$, the cost would equal two times the number of drivers on route $\mathtt{b}$. One can think of route $\mathtt{a}$ as having a greater capacity than route $\mathtt{b}$.

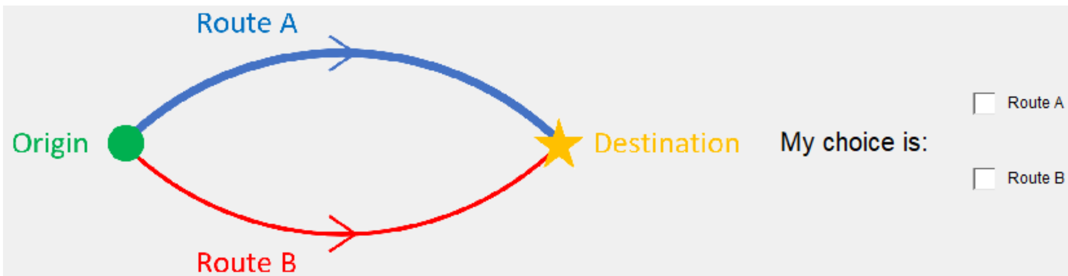

Figure 2: The traffic task in the experiment. Participants decide simultaneously which route, $\mathtt{a}$ or $\mathtt{b}$, to take in each period. Route $\mathtt{a}$ is the lower cost-higher capacity route, while route $\mathtt{b}$ is the higher cost-lower capacity route.

Reaching the destination, regardless of the route taken, earned a driver an unknown amount of money, called the travel value. The travel value was drawn randomly from a uniform distribution on $[6, 16]$ in each period. The draws were independent across drivers and periods. The travel value was not revealed to any player before players made route decisions. A driver's earnings equaled the difference between her travel value and the cost of the route the driver chose in each period.

At the beginning of each period, drivers decide simultaneously which route to take. Once all participants in a group decided, every participant would receive some feedback. The three treatments differed in the feedback about traffic congestion that was provided. Informed and uninformed drivers differed regarding whether or not they had access to traffic congestion information, i.e., the number of drivers that took each route in the period just concluded. Under the $\mathtt{No}$ treatment, all six drivers were uninformed, so at the end of a period, they learned their stage game payoff only. The $\mathtt{Partial}$ treatment randomly assigned three drivers as informed and the three others as uninformed. Informed drivers observed their payoff and the number of drivers on each route at the end of the period. Which drivers were informed remained the same for the duration of the treatment. Under the $\mathtt{Full}$ treatment, all six drivers were informed.

To control for any order effect, the sequence of the three treatments was counterbalanced. Seven or nine groups of participants received each of the 3! = 6 possible sequencings of the three treatments.

4.2 Part 2: The Endogenous Treatment

A driver's type in the three Part 1 treatments was assigned exogenously. In contrast, the Endogenous treatment in Part 2 allowed participants to decide for themselves whether they wished to be informed or not. We elicited valuations for the past congestion information by asking participants to complete the decision matrix shown in Table 3.

Table 3: The decision matrix for the Endogenous treatment of the experiment.

(1) Decision row	(2) Access bonus/fee per period	(3) Will you accept the offer? (Unconditional)	(4) to (9) If 0/1/2/3/4/5 other(s) accept(s) the offer, will you accept it? (Conditional)
1	Bonus: receive 1 point	Yes/No	Yes/No ... Yes/No
2	Bonus: receive 0.75 points	Yes/No	Yes/No ... Yes/No
3	Bonus: receive 0.5 points	Yes/No	Yes/No ... Yes/No
4	Bonus: receive 0.25 points	Yes/No	Yes/No ... Yes/No
5	Receive/Pay 0 points	Yes/No	Yes/No ... Yes/No
6	Fee: pay 0.25 points	Yes/No	Yes/No ... Yes/No
7	Fee: pay 0.5 points	Yes/No	Yes/No ... Yes/No
8	Fee: pay 0.75 points	Yes/No	Yes/No ... Yes/No
9	Fee: pay 1 point	Yes/No	Yes/No ... Yes/No
10	Fee: pay 1.25 points	Yes/No	Yes/No ... Yes/No
11	Fee: pay 1.5 points	Yes/No	Yes/No ... Yes/No
12	Fee: pay 1.75 points	Yes/No	Yes/No ... Yes/No
13	Fee: pay 2 points	Yes/No	Yes/No ... Yes/No

Note: For each of the prices of the traffic information specified in column (2), each participant is asked to enter their decisions in columns (3) to (9).

The decision matrix combines a price list and the use of the strategy method (Selten, 1967). The first column of Table 3 shows the decision row number, with each row corresponding to the price of information specified in Column (2). Both positive (Fees) and negative (Bonuses) prices are included. Providing negative prices allows us to capture the possibility that participants place a negative value on traffic information. In Column (3), individuals enter their *unconditional* decisions. For example, the first decision of Column (3) asks "Will you accept the offer that provides you an access bonus such that you can receive 1 point per period if you decide to be informed?" We call the maximum valuation implied by the decisions in Column (3) the *Unconditional Valuation*. The

Conditional Valuations were obtained by adding conditions to the price list using the strategy method as shown in Columns (4) to (9). In particular, a column of conditional decisions asks participants to decide whether they are willing to accept an offer given that 0 to 5 other player(s) decided to accept the same price to obtain traffic information.

To ensure that all decisions were incentivized, we used the same technique as Fischbacher et al. (2001) and Fischbacher and Gächter (2010), which study behavior in the public good game. Specifically, a decision row is randomly chosen to be implemented and one participant per group is randomly selected. The conditional decisions of this individual for the selected row become operative. For the other five drivers, their unconditional decision in the chosen row is implemented. This random mechanism enables all decisions in the matrix to be potentially payoff-relevant for a participant.

4.3 Questionnaire

A short questionnaire was displayed on participants' screens after Part 2. Some questions asked about demographic characteristics such as gender and major, and we included two quiz questions in the questionnaire to better understand participants' beliefs and intentions when they chose their actions.

The first quiz question asked participants to state what they thought was the lowest travel cost the six group members could have on average in a given period. This question is incentivized by employing a quadratic scoring rule to provide a bonus to participants who answered close to the correct answer, which is a cost of 4. This question helps us to understand whether participants thought that they were doing well in their attempts to minimize costs.

The second quiz question was not incentivized. It asked participants to select a better strategy for the game between two options. One option described a mixed strategy: "Try to change the route you choose frequently and keep people guessing about your route". The other option corresponded to a pure strategy: "Try to find a good assignment of players and routes, and try to stick with the same route." Answers to this question provide clues about the strategy a participant was likely to have been using in the experiment.

4.4 Timing of Events and Payment

Figure 3 illustrates the timeline of an experimental session. The instructions for the first segment were given at the beginning of the experiment. All instructions were read aloud

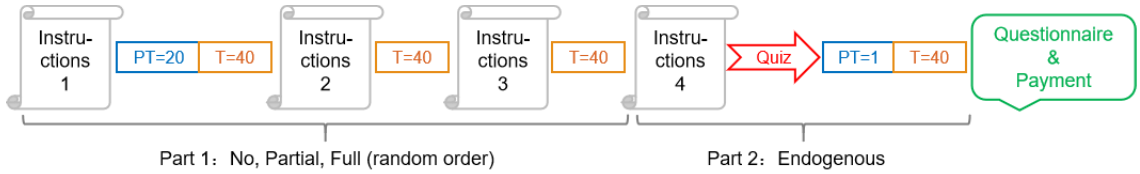

Figure 3: The timeline of an experimental session. In the three segments of Part 1, the No, Partial, and Full treatments were conducted sequentially in an order that differed by session. Part 2 consisted of the Endogenous treatment. The instructions (shown as parchment papers in the figure) for each treatment were read, and the corresponding 40-period task (orange rectangles) was implemented afterward. The practice period(s) (blue rectangles) was (were) played before each part. A quiz (the red arrow) followed the Part 2 instructions to check participants' understanding. At the end of the sessions, participants filled in the questionnaire and received their payment.

while participants followed on their own copies. After participants finished 20 practice periods and 40 remunerated periods of one of the treatments, the instructions for the second segment were provided. Participants played the traffic game of Segment 2 repeatedly for 40 periods. The instructions for the third segment followed, and participants played their third treatment for 40 periods. Afterward, we read the instructions for Part 2, in which the Endogenous treatment was introduced. The Endogenous treatment was always the last treatment in a session. A quiz, added after the first few sessions of the study, followed the instructions for Part 2 to check participants' understanding of the decision matrix. Answers and explanations to the quiz questions were revealed automatically after the responses were submitted. Afterward, each participant practiced filling in the decision matrix once. Their decisions were then played against five computerized bots behaving randomly and they observed the outcome. Then they filled out the matrix again, and this time their decisions determined their information level for the 40 periods of Part 2. At the end of the session, participants filled in the questionnaire described in the last subsection. They were paid their earnings electronically on the same day.[13] The instructions are available in Appendix J.

The final payment to ..... University in the U.S. (..... University in China) participants equaled the sum of (1) the $10 (15 RMB) show-up fee, (2) the total earnings from one of the four treatments in the session (randomly chosen at the end of the session) with an exchange rate of US$1=14 points (1 RMB=4.7 points), and (3) the earnings from the remunerated quiz question. The average payment was US$29.1 at ..... University in the U.S. and 71.8 RMB at ..... University in China.

[13]Due to COVID-19, we used electronic transfer apps, mainly Venmo, Alipay and WeChat, to make payments.

5 Results

This section is structured as follows. Section 5.1 compares the average costs and route choices for the three treatments of Part 1. Section 5.2 reports the tests of Hypotheses 1 and 2. Section 5.3 considers whether Small Sample Models can explain some features of the data. Section 5.4 analyzes how individuals adjust their decisions from one period to the next. Section 5.5 reports the data from Part 2. Section 5.6 presents the responses to our questionnaire.[14]

5.1 Part 1: Comparison among No, Partial and Full

Figure 4 displays the average cost and route choice for the three treatments of Part 1 for Segment 1 only and the complete data from Segments 1 - 3. Figures 4a and 4d present the dynamics of the average cost over time. The connected lines represent the five-period moving average of the average cost. The two figures indicate that Partial and No tend to be more efficient than Full before period 15 and Full tends to be the most efficient in later periods. Figures 4b and 4e present the dynamics of the average number of players choosing route a, averaged over all groups, by period. There is no obvious difference among the three treatments. The average number of players choosing route a in a group fluctuates around the socially optimal level of four. Figures 4c and 4f show the between-group standard deviation of the number of players choosing route a in a period for each treatment. Its time profile is very similar to that of the average cost.

[14]Our analysis begins by considering whether data from the two locations can be pooled. There is no particular reason to expect any cultural difference in the game studied here since it does not involve distributional norms or social preferences. The key variable of interest is the average cost that players incur. Over the 12 groups at ..... University in the U.S., the average cost across the No, Partial, and Full treatments is 4.51 (sd=0.17). Over the 36 groups at ..... University in China, the average cost is 4.50 (sd=0.12). We test the two-sided hypothesis that there is no difference in the cost distribution between the two sites. A Mann-Whitney U test yields a test statistic equal to 217 (p-value of 0.99). We also test whether the mean average costs between the two sites are equal. The bootstrapped two-sample t-statistic is 0.28 (p-value of 0.78). Additionally, we conduct a Mann-Whitney U test and a bootstrapped two-sample t-test of the hypothesis that the rank distribution of, and average, cost in both locations is the same for each treatment separately. The Mann-Whitney U test statistics for the No, Partial, and Full treatments between two locations are 200.5 (p-value=0.72), 252.5 (p-value=0.39), and 144.5 (p-value=0.09), respectively. The corresponding bootstrapped two-sample t-test statistics are 0.36 (p-value=0.72), -0.87 (p-value=0.38), and 1.37 (p-value=0.17). All six test statistics are insignificant at a significance level of 5%. In view of the lack of significant difference in the key performance measure between locations, we pool the data from two sites in our analysis. Appendix I splits the data by location and provides results for each subset. These findings in each location are very close to the aggregate results.

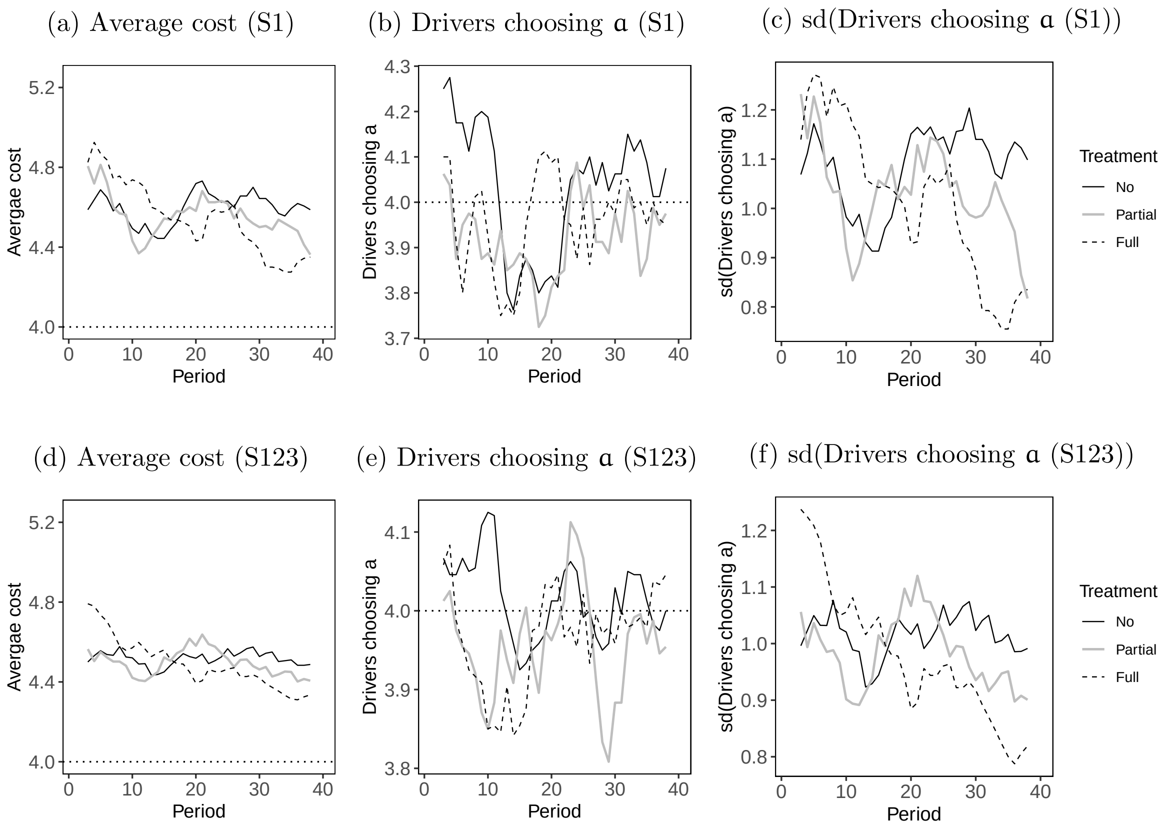

Figure 4: The realized average cost and the number of players choosing route $\mathfrak{a}$ in each treatment over time in Segment 1 only (S1) and in the full data of Segments 1 - 3 (S123). Panels (a) and (d) illustrate the realized average individual cost. Panels (b) and (e) show the average number of players who choose route $\mathfrak{a}$. Panels (c) and (f) present the standard deviation of players choosing route $\mathfrak{a}$. The connected lines indicate five-period moving averages. The horizontal dotted lines indicate the socially optimal levels.

Table 4 summarizes the average cost, the average congestion (the average number of drivers on route $\mathfrak{a}$), and the mean standard deviation of the number of drivers on route $\mathfrak{a}$ (MSDNA) for Segment 1 in panel (a) and for Segments 1 - 3 in panel (b). These values are provided for the first ten periods, for the last ten periods, and for all 40 periods, in each of the three treatments. [15] There are several broad patterns in the data. The first is that the average cost considerably exceeds the optimal level of 4. The second is that the average number of drivers on route $\mathfrak{a}$ is very close to 4, the equilibrium level. The

[15] Appendix E reports tests for carryover effects. These indicate no carryover effects from prior treatments, with the exception of the `No` and `Partial` treatments in Segment 3 of the experiment. In our analysis here, we report the data from Segment 1, which is free of carryover effects, and the full data from all three segments. The data for segments 1 and 2 only is presented in Appendix F, demonstrating that the results are robust to the choice of the data set.

third is that, while the overall average cost under the three treatments is comparable, Full has the highest average cost in the early periods and the lowest in the late periods.

Table 5 reports between-subject (within-subject) test statistics of both Mann-Whitney U and bootstrapped t-tests (Wilcoxon signed-rank and bootstrapped pairwise t-tests) between each pair of treatments for Segment 1 (Segments 1–3).[16] All tests are two-sided. The p-value of the bootstrapped pairwise t-test is computed by applying the bootstrap-based clustered error (Cameron et al., 2008), with clustering at the level of the group.

Table 4: Average cost and the number of drivers on route a

Periods	Average cost			Drivers on route a			sd(Drivers on route a)			Obs.
	First10	Last10	All40	First10	Last10	All40	First10	Last10	All40	
(a) Segment 1										
No	4.60	4.58	4.59	4.22	4.09	4.05	1.09	1.08	1.08	16
	(0.19)	(0.21)	(0.10)	(0.35)	(0.32)	(0.20)				
Partial	4.69	4.45	4.56	4.01	3.98	3.94	1.13	0.94	1.03	16
	(0.34)	(0.18)	(0.16)	(0.31)	(0.25)	(0.19)				
Full	4.79	4.32	4.55	4.06	3.97	3.98	1.19	0.81	1.01	16
	(0.39)	(0.23)	(0.23)	(0.25)	(0.17)	(0.12)				
(b) Segments 1 - 3										
No	4.54	4.49	4.51	4.06	4.02	4.02	1.04	1.00	1.01	48
	(0.23)	(0.25)	(0.17)	(0.39)	(0.35)	(0.21)				
Partial	4.53	4.42	4.50	3.97	3.97	3.96	1.02	0.92	0.99	48
	(0.28)	(0.22)	(0.17)	(0.28)	(0.25)	(0.16)				
Full	4.69	4.34	4.49	3.99	4.01	3.98	1.15	0.83	0.98	48
	(0.34)	(0.27)	(0.19)	(0.27)	(0.25)	(0.13)				

Notes: Panels (a) and (b) summarize results for data from Segment 1 only and Segments 1 - 3, respectively. The units of observation for the first three columns are the average cost for a participant for the first ten periods, for the last ten periods, and for all 40 periods in a given treatment. The next three columns show the average number of drivers on route a in a group for the first ten periods, for the last ten periods, and for all 40 periods in a given treatment. Standard deviations are in parentheses. A standard deviation is calculated by averaging the cost of the members (counting the number of drivers on route a) of each group in each period, averaging over the 10 or 40 periods for each group, and then computing the between-group standard deviation of these averages. Columns 7 to 9 display the mean standard deviation of the number of drivers on route a (MSDNA) for the first ten periods, for the last ten periods, and for all 40 periods in a given treatment. A mean is calculated by computing the between-group standard deviation of the number of drivers on route a in a given period, and then averaging the resulting values over the 10 or 40 periods.

Over all 40 periods, the average costs under No, Partial, and Full in Segment 1 are 4.59, 4.56, and 4.55. In the pooled data from Segments 1 - 3, the corresponding average costs are 4.51, 4.50, and 4.49. None of the pairwise comparisons of the average

[16]Tests are conducted for the average cost and the number of drivers on route a only. Since there is one observation of the standard deviation of the number of drivers choosing a in a given period of each treatment, we compare the value of MSDNA directly without conducting any statistical tests.

Table 5: Tests of treatment differences in average cost and the number of drivers on route a

Periods	Average cost			Drivers on route a			Obs.
	First10	Last10	All40	First10	Last10	All40	
(a) Segment 1							
No vs. Partial:							
Mann-Whitney U	121	172	134	170	147	163.5	32
	[0.81]	[0.10]	[0.84]	[0.12]	[0.48]	[0.19]	
Bootstrapped t	-0.87	1.86	0.69	1.77	1.16	1.57	32
	[0.40]	[0.06]	[0.51]	[0.09]	[0.25]	[0.13]	
No vs. Full:							
Mann-Whitney U	98	204	149	166.5	153	150	32
	[0.27]	[0.00]	[0.44]	[0.15]	[0.35]	[0.42]	
Bootstrapped t	-1.74	3.25	0.70	1.52	1.38	1.11	32
	[0.09]	[0.00]	[0.53]	[0.16]	[0.19]	[0.28]	
Partial vs. Full:							
Mann-Whitney U	109.5	175.5	133.5	123	140.5	112	32
	[0.50]	[0.08]	[0.85]	[0.86]	[0.65]	[0.56]	
Bootstrapped t	-0.80	1.75	0.16	-0.44	0.08	-0.79	32
	[0.45]	[0.12]	[0.88]	[0.68]	[0.94]	[0.46]	
(b) Segments 1 - 3							
No vs. Partial:							
Wilcoxon signed-rank	596	603.5	659.5	535	516.5	696.5	48
	[0.56]	[0.12]	[0.47]	[0.18]	[0.42]	[0.09]	
Bootstrapped pairwise t	0.10	1.98	0.54	1.57	0.81	1.84	48
	[0.93]	[0.04]	[0.59]	[0.10]	[0.42]	[0.07]	
No vs. Full:							
Wilcoxon signed-rank	302	773	620	571	536.5	574.5	48
	[0.02]	[0.00]	[0.39]	[0.24]	[0.83]	[0.52]	
Bootstrapped pairwise t	-2.75	3.56	0.67	1.20	0.14	1.38	48
	[0.01]	[0.00]	[0.50]	[0.21]	[0.89]	[0.17]	
Partial vs. Full:							
Wilcoxon signed-rank	312	763.5	589	434.5	370	428	48
	[0.01]	[0.07]	[0.80]	[0.84]	[0.44]	[0.44]	
Bootstrapped pairwise t	-2.42	1.57	0.12	-0.41	-0.85	-0.92	48
	[0.04]	[0.16]	[0.91]	[0.69]	[0.41]	[0.36]	

Notes: Panels (a) and (b) compare treatments using data from Segment 1 only and Segments 1 - 3, respectively. In Panel (a), between-subject pairwise comparisons are conducted with respect to the average cost and the number of players choosing route a among No, Partial, and Full reported in Table 4 with two statistical tests: Mann-Whitney U and the bootstrapped t-test. In Panel (b), within-subject pairwise comparisons are conducted with two statistical tests: the Wilcoxon signed-rank test and the bootstrapped pairwise t-test. All tests are two-sided. The p-values are in square brackets.

cost is significantly different between any two treatments in either data set. Pairwise comparisons between treatments also indicate no significant differences in the average number of drivers on route $\mathtt{a}$ over 40 periods. Moreover, the mean standard deviation of the number of drivers on route $\mathtt{a}$ during all 40 periods is very close between any two treatments. However, this lack of significant differences masks the fact that different patterns are observed in the early and late periods.

In the first ten periods, Full is the least efficient of the three treatments. For Segment 1, efficiency differences are not significant at the 5% level. However, using data from Segments 1 - 3, a Wilcoxon signed-rank test and a bootstrapped pairwise $\mathtt{t}$-test find that Partial is more efficient than Full at 1% and 5% significance levels, respectively. The two tests also conclude that No is more efficient than Full at significance levels of 5% and 1%. There is no significant difference in average cost between No and Partial. While the number of drivers on route $\mathtt{a}$ does not differ significantly across treatments at the 5% level in the first ten periods, the MSDNA is highest under Full. Overall, providing all individuals with past traffic information yields the least efficient traffic flow in the short run.

In contrast, during the last ten periods of a segment, Full Information is the most efficient arrangement, followed by Partial. For Segment 1, the Mann-Whitney $\mathtt{U}$ and bootstrapped $\mathtt{t}$-tests reveal that No is less efficient than Partial at a 10% significance level and less efficient than Full at a 1% significance level. Additionally, a Mann-Whitney $\mathtt{U}$ test shows that Full is more efficient than Partial at a 10% significance level. Analyzing data from Segments 1 - 3, the bootstrapped pairwise $\mathtt{t}$-test indicates that Partial is more efficient than No at a 5% significance level, while the Wilcoxon signed-rank test suggests that Partial is less efficient than Full at a 10% significance level. Both tests show that No is less efficient than Full at a 1% significance level. Pairwise comparisons between any two treatments indicate no significant differences in the average number of drivers taking route $\mathtt{a}$ in the long run. The MSDNA is highest under the No treatment and is lowest under the Full treatment.

Providing some or all individuals with past traffic information leads to more efficient outcomes as time passes by, whereas the average cost under No does not change significantly over time. In Segment 1, the average cost decreases by 0.03 (sd=0.30), 0.24 (sd=0.32), and 0.47 (sd=0.42) under No, Partial, and Full, respectively. For Segments 1 - 3, the corresponding decreases are 0.04 (sd=0.28), 0.11 (sd=0.27), and 0.34 (sd=0.39). Both the bootstrapped pairwise $\mathtt{t}$ and the Wilcoxon signed-rank test statistics suggest that the average cost under Partial and Full are lower in the long run compared to the short run at the 1% significance level under both data sets. No significant changes are observed in the average cost under No. [17]

5.2 Hypothesis Evaluation

This section evaluates our hypotheses, with a summary of the test results presented in Table 6. Non-parametric and bootstrapped t-tests are conducted using two data sets, resulting in four statistics for each relationship. To claim a relationship between two elements to be different, we require at least two out of four statistics (the t-test and non-parametric test for Segment 1 or the three segments) to be significantly different. Otherwise, the relationship is treated as equal.

Table 6: Hypothesis evaluation and comparison of EWA-lite

Time Period	Empirical Relationship of Average Cost	Consistent with H1 (.5)	Empirical Relationship of Drivers on Route a	Consistent with H2 (.5)
P1-P10	$C(N) = C(P)$	N	$A(N) = A(P)$	N
	$C(N) < C(F)$	N	$A(N) = A(F)$	N
	$C(P) < C(F)$	Y	$A(P) = A(F)$	N
P31-P40	$C(N) > C(P)$	Y	$A(N) = A(P)$	N
	$C(N) > C(F)$	Y	$A(N) = A(F)$	N
	$C(P) > C(F)$	N	$A(P) = A(F)$	Y
Early vs. Late	$C_L(N) = C_E(N)$	Y	$A_L(N) = A_E(N)$	N
	$C_L(P) < C_E(P)$	Y	$A_L(P) = A_E(P)$	N
	$C_L(F) < C_E(F)$	Y	$A_L(F) = A_E(F)$	N

Notes: Columns two and four of the table describe the empirical relationships of average cost (denoted as $C(\cdot)$) and the number of drivers selecting route a (denoted as $A(\cdot)$) across treatments No (N), $\mathsf{Partial}$ (P) and Full (F) during periods 1-10 (denoted as subscript E for Early) and periods 31–40 (denoted as subscript L for Late), along with the treatment-specific changes over time. Columns three and five indicate whether the empirical relationship is consistent with Hypotheses 1 and 2. The relationship of two elements in the hypotheses is considered different if the difference in their averages is at least 0.5 pooled-sample standard deviations in the simulations.

By the criterion, of the nine average cost relationships that Hypothesis 1 predicts, six are borne out in the data. $\mathsf{Partial}$ is accurately predicted to be the most and Full the least efficient treatment in the early periods. Full is correctly predicted to be the most and No the least efficient in the late periods. However, the performance in No is comparable to $\mathsf{Partial}$ in the early periods while $\mathsf{Partial}$ is not as efficient as Full

[17] With respect to the change in the number of drivers on route a, none of the null hypotheses of no change over time can be rejected. Thus, there are no significant trends over time in the number of drivers choosing route a for each treatment. The MSDNA decreases by 0.01, 0.19, and 0.38 under No, $\mathsf{Partial}$, and Full, respectively, from the early to the late periods in Segment 1. The corresponding reductions are 0.04, 0.10, and 0.32 for the data set of Segments 1–3.

in late periods, patterns unanticipated by EWA-lite. Hypothesis 2 does not fare as well, with only one of its nine predictions observed. Much of the failure of Hypothesis 2 to predict these empirical relationships is because the traffic on route $\mathtt{a}$ in the $\mathtt{No}$ treatment is less than predicted, especially in the late periods. We thus conclude that, overall, Hypothesis 1 is supported while Hypothesis 2 is not.

5.3 Can Small Sample Models Explain Some Patterns in the Data?

As we have seen in the last two subsections, the EWA-lite model has considerable success in predicting empirical patterns but exhibits a number of inaccuracies. We focus on three of these here. The first is that the $\mathtt{Partial}$ treatment fails to outperform $\mathtt{No}$ in the early periods. This is the main source of inaccuracies for Hypothesis 1. The second is that costs in the $\mathtt{Partial}$ and $\mathtt{Full}$ conditions remain considerably above the minimum level of 4 in the late periods, with $\mathtt{Partial}$ exhibiting higher costs than $\mathtt{Full}$. The reason for this is that there remains considerable variance in the number of drivers taking route $\mathtt{a}$ in the late periods. The third is that fewer individuals chose route $\mathtt{a}$ under the $\mathtt{No}$ condition than predicted, causing Hypothesis 2 to not be supported.

We consider here whether Small Sample Models explain some of these discrepancies between EWA-lite and the data. In these models, individuals learn the payoffs of their possible actions by sampling from their previous choices and choosing the action that yielded the higher average payoff within the sample. These models do not have sequential adjustment, the updating of actions based on the outcome of the preceding trial, because all prior periods have an equal chance of being sampled.

We employ two versions of the model, following Erev and Roth (2014). The first assumes that the individual samples exactly $\mathtt{q} = 5$ random draws with replacement from her prior history of play, drawing a new sample in each period. We refer to this as the Sample of 5 model (SSM5). The second assumes that in each period, an individual samples $\mathtt{q}_\mathtt{i}$ prior periods, where $\mathtt{i}$ indexes the individual. We assume that $\mathtt{q}_\mathtt{i}$ is drawn independently and uniformly from $\{1, 2, ..., \mathtt{q}\}$ and $\mathtt{q} = 9$ for each individual. We refer to this model as the Sample of 9 or Less model (SSM1-9). $\mathtt{q}$ is the sole free parameter of the two small sample models. It can be thought of as representing a measure of the willingness or the capacity for calculation.

The models assume random choice in the first trial. Starting in trial 2, an informed player selects the action with the highest average payoff in the small sample they have drawn. For an uninformed player in trial $t \geq 2$, if only one action has been selected in the past, the player will select an action randomly in the current trial. If both actions have been selected in the past, for each action, the uninformed player will draw a small sample of size 5 (or a number from 1 to 9), of past trials in which the action was chosen. Then, the uninformed player compares the average payoff of the two small samples and selects the action with the highest average payoff in its small sample. See Appendix B for details.

The simulation results of the two models are plotted in Figure 5. The top three panels in the figure illustrate the average cost, the average number of players choosing route a, and the standard deviation for the Sample of 5 model, respectively. The corresponding information for the Sample of 9 or Less model is presented in the bottom three panels. Detailed average values and standard deviations for the first and last ten periods can be found in Table 7. The two Small Sample models exhibit closely aligned patterns across all dimensions.

Both Panels (a) and (d) indicate that under the SSM models the average cost under Partial and No is lower than under Full over the entire 40-period horizon. Specifically, the average costs over the 40 periods in SSM5 are 4.39 (sd=0.16), 4.41 (sd=0.18), and 4.59 (sd=0.36) for No, Partial and Full, respectively. In SSM1-9, the corresponding values are 4.42 (sd=0.16), 4.43 (sd=0.17), and 4.56 (sd=0.34). Costs decrease over time substantially under all treatments. During the first ten periods, the cost is highest under Full and comparable between Partial and No. During the last ten periods, the cost is highest under Full and lowest under No. However, unlike under EWA-lite, the average cost remains considerably above the minimum level of 4 in all three treatments.

In Panels (b) and (e), the averages of the number of players choosing route a during the 40 periods under No, Partial and Full are 4.13 (sd=0.30), 4.14 (sd=0.13), and 4.09 (sd=0.08) in SSM5, and they are 4.12 (sd=0.30), 4.08 (sd=0.13), and 4.06 (sd=0.08) in SSM1-9, respectively. These numbers are comparable among treatments and consistent throughout the 40 periods. Panels (c) and (f) show that the standard deviation of drivers taking route a remains substantial even after 40 periods.

Unlike EWA-lite, the SSMs predict that average costs do not converge to the minimum of four under Partial and Full, even in the long run. There remains some variance in the number of individuals taking route a in the long run in these two treatments. Fewer drivers take route a in No under SSM than EWA-lite. Partial does not significantly outperform No in Periods 1 - 10. Thus, the SSMs capture some of the features of the data that EWA-lite misses. However, the SSMs also exhibit an important inac-

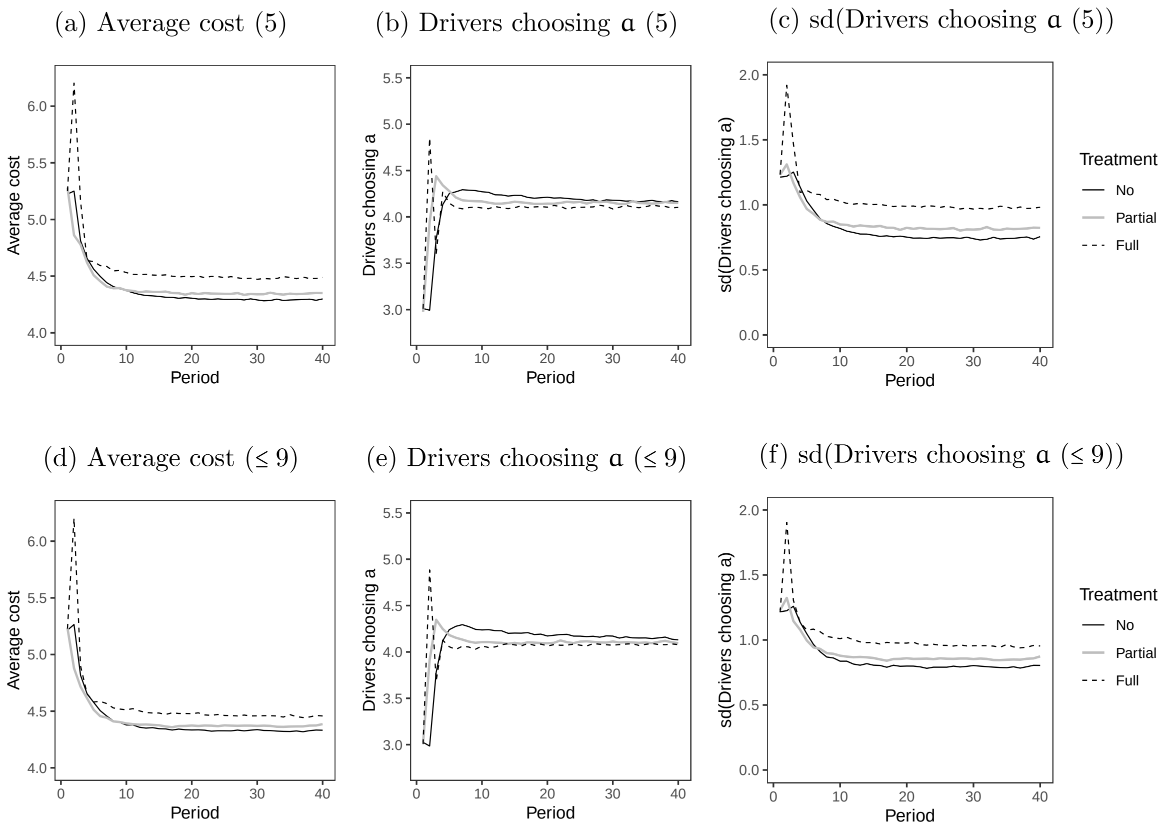

Figure 5: Dynamics of cost and traffic patterns for models "Sample of 5" (5) and "Sample of 9 or Less" (≤ 9) over a 40-period horizon. 10,000 simulations of a six-player group are conducted under each of the No, Partial, and Full treatments. Panels (a) and (d) plot the average individual cost. They show that Full exhibits the highest average cost throughout the 40 periods; Panels (b) and (e) show the average number of players choosing route a, the higher capacity route, in a group. This number converges to close to the optimal level of 4 in all treatments, though it does so most slowly under No; Panels (c) and (f) present the standard deviation of the number of players choosing route a. The standard deviation is highest under Full throughout the 40 periods.

curacy. They predict that Full would exhibit the highest average cost among the three treatments in both the early and the late periods, while the data show it does so only in the early periods.

Under the Small Sample models, in contrast to the EWA-lite model, the value of information is never significantly greater than zero regardless of how many other drivers also have the information. Appendix B.2.2 provides some details regarding the simulations that generate this prediction. The implied values of information under SSM5 are 0.10 (sd=0.39), -0.01 (sd=0.35), -0.13 (sd=0.32), -0.22 (sd=0.31), -0.30 (sd=0.34), and -0.28 (sd=0.48), when 0 to 5 other drivers also have information, respectively. Under SSM1-9,

Table 7: Small Sample models simulation results

Model	Info.	Avg. Cost Periods 1-10	Avg. Cost Periods 31-40	Num. Routa A Periods 1-10	Num. Routa A Periods 31-40
Sample of 5	No	4.66 (0.34)	4.29 (0.18)	3.96 (0.38)	4.17 (0.35)
	Partial	4.61 (0.32)	4.35 (0.21)	4.09 (0.21)	4.15 (0.25)
	Full	4.87 (0.58)	4.49 (0.35)	4.04 (0.14)	4.11 (0.25)
Sample of 9 or less	No	4.67 (0.34)	4.33 (0.19)	3.95 (0.38)	4.15 (0.35)
	Partial	4.61 (0.31)	4.37 (0.22)	4.03 (0.22)	4.10 (0.26)
	Full	4.82 (0.55)	4.46 (0.34)	4.00 (0.14)	4.08 (0.25)

Notes: The average values of cost and number of drivers choosing route a in the first ten and last ten periods in 10,000 simulations are reported. Standard deviations are given in parentheses. A standard deviation is calculated by averaging the cost of individual drivers or counting the number of drivers choosing route a of each group in each period, averaging over the 10 (periods) for each group, and then computing the between-group standard deviation of these averages.

the corresponding implied information values are 0.09 (sd=0.40), -0.01 (sd=0.34), -0.09 (sd=0.31), -0.16 (sd=0.30), -0.22 (sd=0.32), and -0.20 (sd=0.45). Nevertheless, both SSM models suggest, as does EWA-lite, that the value of getting informed decreases in the number of other drivers who are also informed.

5.4 Adjustment Behavior

In this section, we show that participants tend to adjust their behavior toward lower-cost routes. However, this adjustment is insufficient to move the group to the optimal traffic distribution.

We begin with an analysis of individual adjustment behavior. We follow Selten et al. (2007) and employ the Yule coefficient Q to study individual "response mode". The traffic cost to each driver in a pure-strategy equilibrium is 4. We classify a response of a driver as *direct* if she changes her route after incurring a cost greater than 4 (event 1) or not changed after a cost smaller than 4 (event 2). A response of a driver is classified as *contrary* if the route is changed after a cost smaller than 4 (event 3) or not changed after a cost greater than 4 (event 4). The Yule coefficient Q can then be calculated as follows.

$$Q = \frac{P(event\ 1) * P(event\ 2) - P(event\ 3) * P(event\ 4)}{P(event\ 1) * P(event\ 2) + P(event\ 3) * P(event\ 4)}$$

Q varies from -1 to 1. A higher Q represents more frequent direct adjustment behavior and vice versa. Drivers who never/always changed routes or whose cost was always equal to 4 are dropped as their Yule coefficients are not well defined.

Table 8: Yule coefficients Q by treatment: early, late and all periods

Treatment	First10			Last10			All39		
	mean	sd	num	mean	sd	num	mean	sd	num
(a) Segment 1									
No	0.21	0.92	62	0.22	0.87	57	0.26	0.63	90
Partial	0.21	0.91	64	0.19	0.92	51	0.29	0.62	88
UInf	0.12	0.90	32	-0.06	0.95	27	0.22	0.61	47
Inf	0.30	0.93	32	0.48	0.81	24	0.36	0.63	41
Full	0.15	0.83	62	0.09	0.94	35	0.25	0.65	88
(b) Segments 1 - 3									
No	0.28	0.86	164	0.19	0.88	140	0.29	0.60	244
Partial	0.27	0.87	166	0.25	0.88	143	0.31	0.63	254
UInf	0.16	0.86	80	0.08	0.93	68	0.21	0.64	129
Inf	0.37	0.87	86	0.40	0.82	75	0.41	0.61	125
Full	0.26	0.86	186	0.23	0.89	109	0.33	0.64	257

Notes: The Yule coefficients Q are calculated for individuals under No, Partial, and Full for $t \in \{1, ..., 39\}$. Observations are at the individual level. Panel (a) reports the mean Q values, standard deviations, and numbers of observations during the first ten, the last ten, and all 39 periods in Segment 1, while Panel (b) reports these values using data from Segments 1 - 3. The two types, UInf and Inf, under Partial are also reported separately.

The Yule coefficients Q for the three treatments for the first ten, the last ten, and all periods are presented in Table 8. The table indicates that mean Q values are always positive. Individuals are inclined to make direct responses under each treatment. Furthermore, the two data subsets both reveal that the informed under Partial have the highest average Q value, while the uninformed under Partial exhibit the lowest. Under Partial, informed individuals change over to less crowded routes more than the uninformed. In Segment 1, the average Q value for informed drivers under the Partial treatment surpasses that of uninformed drivers under the same treatment in the final ten periods, at a significance level of 5%. In Segments 1 - 3, the informed also have a significantly higher average Q than the uninformed, both over the last ten periods (p-value=0.03) and across all periods (p-value<0.01).

We now consider changes in route choice from one period to the next at the group level. We analyze the probability of overreaction and underreaction. *Overreaction* occurs when so many participants in a group choose the route that was less costly in the immediately preceding period to make it more costly than the other in the current period. In contrast, when so few participants select the previously lower-cost route, keeping it the less costly, *underreaction* occurs. Table 13 in Appendix C reports the observed transition probabilities under each treatment for data in Segment 1 and Segments 1 - 3.

The probability of overreaction between periods t and $t+1$ is calculated as:

$$\frac{P[(C^a_{t+1} > C^b_{t+1}) \text{ and } (C^a_t < C^b_t)] + P[(C^a_{t+1} < C^b_{t+1}) \text{ and } (C^a_t > C^b_t)]}{P[C^a_t \neq C^b_t]}$$

where $P[(C^j_t > C^k_t)]$ is the probability that the cost on route j is greater than that on route k in period t. Similarly, underreaction is defined as:

$$\frac{P[(C^a_{t+1} > C^b_{t+1}) \text{ and } (C^a_t > C^b_t)] + P[(C^a_{t+1} < C^b_{t+1}) \text{ and } (C^a_t < C^b_t)]}{P[C^a_t \neq C^b_t]}$$

During the first ten periods of Segments 1 - 3 (Segment 1), the observed probabilities of underreaction are 0.41 (0.41), 0.32 (0.35), and 0.44 (0.45) in the No, $\mathsf{Partial}$ and Full treatments, respectively. The EWA-lite model predicts the percentages of underreaction of 90%, 28%, and 15% for these treatments, while SSM5 (SSM1-9) predicts 48% (47%), 30% (28%), and 21% (21%). For overreaction during the same periods in Segments 1 - 3 (Segment 1), the probabilities are 0.28 (0.28), 0.25 (0.25), and 0.28 (0.30) in the No, $\mathsf{Partial}$ and Full treatments, respectively. EWA-lite predicts overreaction rates of 8%, 49%, and 75%, while SSM5 (SSM1-9) predicts 19% (20%), 33% (35%), and 55% (54%). While both models exhibit inaccuracies, the SSM models predict the percentages considerably more accurately.[18]

In periods 31 - 40, underreaction probabilities for Segments 1 - 3 (Segment 1) are 0.36 (0.40), 0.29 (0.27), and 0.30 (0.34) in the No, $\mathsf{Partial}$, and Full treatments, respectively. The EWA-lite model predicts underreaction rates of 100%, 37%, and 17%, while the SSM5 (SSM1-9) model, they are 37% (36%), 30% (29%), and 32% (31%). The long-run probabilities of overreaction for Segments 1 - 3 (Segment 1) are 0.27 (0.28), 0.30 (0.33), and 0.29 (0.26) in the No, $\mathsf{Partial}$, and Full treatments. EWA-lite predicts overreaction rates of 0%, 35%, and 68%, respectively, while SSM5 (SSM1-9) predicts 17% (20%), 26% (27%), and 35% (34%). The SSM models predict remarkably close to the actual frequencies.[19]

[18] A Wilcoxon signed-rank test reveals that the probability of underreaction in the early periods is significantly greater than that of overreaction under No and Full at a 5% significance level in both data sets. Comparing treatments, the probability of underreaction in Segments 1 - 3 is significantly greater under No than $\mathsf{Partial}$ at a 10% significance level and under Full than $\mathsf{Partial}$ at a 1% significance level.

[19] A Wilcoxon signed-rank test reveals that during the late periods, underreaction is more likely compared to overreaction under the No treatment at a significance level of 1% in Segments 1 - 3. Regarding treatment differences, the probability of underreaction is higher under No than $\mathsf{Partial}$ at

Compared to the early periods, the underreaction tendency under Full diminishes by 11% in Segment 1 and 14% in Segments 1 - 3. The tendency is significant at confidence levels of 5% and 1%, respectively. Additionally, the likelihood of overreaction in Full during the late periods decreases by 4% in Segment 1 and increases by 2% in Segments 1 - 3, but these changes are insignificantly different from zero. This evidence suggests that Full facilitates greater coordination success over time mainly by reducing underreaction without increasing overreaction.[20]

The relative success of the SSM models in capturing the extent of underreaction and overreaction appears to be due to their relative insensitivity to the different treatments. EWA-lite overpredicts (underpredicts) the incidence of underreaction (overreaction) in the No treatment while underpredicting (overpredicting) it in Full.

5.5 Part 2: The Valuation of Traffic Information

Part 2 of the experiment elicits individuals' valuations of the traffic information, conditional on how many other drivers are also informed. The heatmap in Figure 6 summarizes the observed choices of participants from the decision matrix in Table 3. This figure illustrates the proportion of participants choosing to receive traffic information at various prices. The left panel displays the unconditional decision column, UC, while the right panel depicts the conditional decision columns. We use C0 to C5 to denote the six conditional decision columns, representing scenarios where zero to five others are informed, respectively. Additionally, the price of the traffic information ranges from -1 to 2 in 0.25 increments. A negative price means giving participants who choose to access information a certain amount of additional points per period. Combining 13 price lines and seven decision columns forms a matrix containing 91 decisions.

We segment the data into five equal intervals, covering the range [0,1]. These intervals are visually represented by five distinct color shades in the heatmap. Darker colors correspond to a higher proportion of participants opting to buy the information. As prices increase, fewer participants are willing to accept the offer. The most significant shift occurs within the price range of -0.25 to 0.75.

The figure reveals that 22% to 28% of participants under C0 to C5 are not willing to accept the information even when it is free to take it. Notably, 5% to 11% decline

a 5% significance level, while the probability of overreaction is higher under Partial than under Full in Segment 1 at a 1% significance level.

[20] The overreaction rate under Partial increases by 8% in Segment 1 at a borderline significance level of 10%.

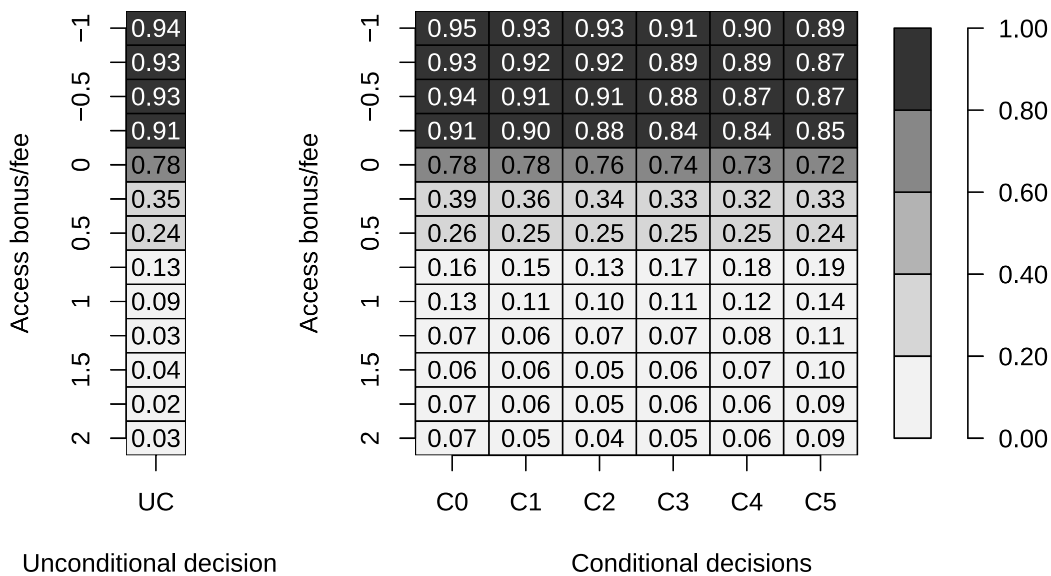

Figure 6: The heat map of the valuation of information for all 288 participants. Each decision cell shows the proportion of participants who chose to access traffic information. Negative values in the vertical axes correspond to access bonuses (or negative prices for information), while positive values represent access fees or positive prices. The horizontal axes specify both unconditional and conditional decision columns. UC in the left panel refers to the unconditional decision column. C0 to C5 in the right panel specify decisions conditional on how many other drivers are informed.

the offer even when offered a bonus of 1 point per period, consistent with a negative valuation for the information. The one-point bonus is equivalent to 15% of the average period profit and 22% of the average traffic cost. About one-third of participants are willing to pay 0.25 points per period for access to information. A small subset, 4% to 9%, values the traffic information at no less than 2 points, even though the average payoff difference in a period between informed and uninformed individuals is less than 0.05 points in Part 1. These individuals demonstrate a strong preference for being informed.

Among the 288 participants, 256 demonstrated consistency in their choices. Consistency is defined as having at most one switching point from accepting the offer to rejecting it as the price increases in each column of the decision matrix. Participants displaying inconsistency are excluded from the following analysis. For the consistent participants, the average unconditional willingness-to-pay (WTP) is 0.09 points. The average WTP conditional on no other person also buying is a price of 0.16 points. As the number of other informed players increases, the average WTP slightly decreases.

The average WTP for information is positive in all six conditions, ranging from 0.08 to 0.16 points. This is at odds with Hypothesis 3. However, a weaker prediction of the hypothesis is that the WTP under C0 and C1 would be higher than under the other conditions. A comparison of the mean WTP under C0–C1 and versus C2–C5 yields a t-statistic of 1.70, supporting the hypothesis that WTP is higher when no more than one other player is informed than otherwise at a significance level of 10% (two-sided).[21] Overall, there is modest evidence that, as the models predict, the willingness-to-pay for information is decreasing in the number of others who also have information.

5.6 Questionnaire: Beliefs about the decision environment

In the questionnaire after Part 2, there was an incentivized question that asked participants to state what they thought was the lowest travel cost the six group members could have on average in a given period. There were 282 respondents. From the analysis, we dropped 5 outliers, who thought that the lowest average travel costs were greater than 40. Among these, two reported 60, two reported 80, and one reported 160. Figure 7 is a histogram of the lowest average travel cost as stated by the remaining 277 participants. The figure shows that 112 participants (40%) provided the correct response, which is a cost of 4. However, 64 participants (23%) thought the cost was greater than 4, and 101 participants (36%) stated that the cost was less than 4. Participants who thought the cost was less than 4 account for more than one-third of the observations, indicating that even if a group reaches the lowest possible cost of 4, some members might change their actions in an attempt to do better, so optimal outcomes are unstable in a behavioral sense.

Another, unincentivized, quiz question asked participants the following question regarding the best type of strategy to employ in the route choice task:

Which of the following two options do you think is better?
A. Try to change the route you chose frequently and keep people guessing/uncertain about your route.
B. Try to find a good assignment of drivers and routes and stick with the same route in different periods.

Option A describes a mixed strategy, while option B refers to a pure strategy. Among the 282 participants who answered the question, 56 indicated that they thought a mixed

[21]The unconditional valuations of 97% of participants fall within the range of their conditional valuations, demonstrating a considerable degree of consistency between the two measures of valuation.

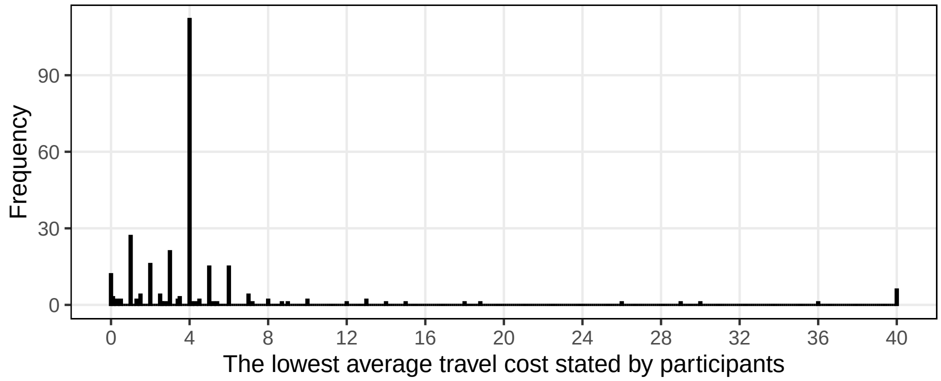

Figure 7: Histogram of beliefs regarding the lowest possible average travel cost. The average minimum travel cost as stated by 277 participants is 4.92 (sd=6.64). The median stated travel cost is 4, the correct response.

strategy was better than a pure strategy. Thus, while the majority responded as if they were seeking a stable traffic configuration, 20% thought that switching routes was a good strategy. This is problematic for efficiency because it means that even if an efficient traffic pattern is found, it is not likely to be stable, because a minority of 20% of mixers is sufficient to introduce great instability into the system.

6 Conclusion

This study considered the impact of historical information on congestion levels in a transportation network using a laboratory experiment. This type of information is used in commonly-employed apps reporting congestion on roads, in restaurants, and in shops. We employed the EWA-lite learning model as our source of hypotheses. Our predictions are therefore the outcomes of a learning process with feedback rather than of an equilibrium analysis. The hypotheses were tested in the experimental data.

The experimental results of Part 1 show that whether having past traffic information available is beneficial depends on the experience level of users. In the first ten periods of play in our experiment, there is greater traffic efficiency if nobody or 50% of users have the information compared to when all users have information. That is to say, everyone having the information harms efficiency in the beginning.

When participants have experience using the information in the last ten periods of a treatment, `Full` adoption is more efficient than `Partial` adoption, and both `Partial` and `Full` adoption are better than `No` adoption. This suggests that when users are sufficiently experienced, allowing all users to adjust to past information is beneficial for efficiency.

There are also a number of non-differences that we find interesting. Under No Information, there is no improvement in efficiency over time, as EWA-lite predicts. Partial is not more efficient than No in the early periods. The average number of drivers on route a is not different among treatments and does not change over time, contrary to the predictions of EWA-lite. These results indicate that the impact of information provision on traffic patterns is more modest than modeling might anticipate.

The EWA-lite model is a good predictor of the differences in congestion costs among treatments and over time, though a less accurate predictor of the amount of traffic on each route. Applying two versions of SSM models, we find that the SSM models do capture a number of patterns in the data that EWA-lite does not. It appears that its accuracy lies in the fact that it describes well the dynamics of adjustment of the overall traffic pattern from one period to the next and how these differ among information conditions. In other words, it predicts the extent of group over-and under-adjustment quite accurately.

The questionnaire data shed light on why the groups were typically unable to converge on an optimal configuration and remain there. When asked about the lowest possible cost they thought a group could average in a period, only two-fifths of respondents gave the correct answer of 4, and many indicated values of 3 or less. This suggests that even when groups were behaving close to optimally, they did not believe that they were, prompting them to continue to try new strategies. Changing strategies would tend to leave a group worse off if it were previously operating close to its optimum. In answer to the question that was trying to get at whether agents thought it was better to coordinate on a pure-strategy or a mixed-strategy equilibrium, a substantial minority indicated the mixed-strategy equilibrium. Because the mixed-strategy equilibrium leads to greater expected costs, and thus lower profit, than the pure-strategy equilibria, pursuing this strategy profile, even successfully, would fail to result in efficient outcomes. Even more adversely for efficiency, there was disagreement on whether a pure or a mixed strategy was preferable, guaranteeing problems with coordination.

The experimental results of Part 2 reveal some tendency for the value individuals place on traffic information to be higher, the fewer other users who are informed. This finding is consistent with the two learning models.

In our data, we observe that the average number of individuals choosing route a is close to the pure-strategy equilibrium number of 4 in all three treatments, though with considerable variability in behavior. As we have noted earlier, this phenomenon has been termed as “magic”. For a range of parameters, the EWA-lite models predict convergence to this level under `Partial` and `Full` Information, while the SSM models predict convergence to close to this level in all three treatments. These learning models, in our view, provide a good account for how this convergence to the equilibrium proportions can occur in the absence of equilibrium reasoning.

We believe that this paper provides some insights to inform a rollout strategy for congestion apps. If our results can be generalized, it is better in the short run to facilitate adoption of the apps for some rather than all users, so that initially, `Partial` information conditions are present. When users are experienced, `Full` adoption has no downside, and it is best for all users to have the information.

An important extension to the research reported here would be to add important features characteristic of applied settings. To better capture the way in which traffic apps operate, one could add a predictive algorithm to make the information reflect expected congestion, time of day effects so that overall congestion increases during peak hours, and shocks to the environment similar to those caused by sudden road closures or other unanticipated events.

References

Anderson, C. M. and C. F. Camerer (2000): “Experience-weighted attraction learning in sender-receiver signaling games,” *Economic Theory*, 16, 689–718.

Ashraf, S., A. L. Brown, M. W. Burris, and V. Vitaku (2023): “Aggregate and individual effects of information in a coordination (traffic) game,” *Economic Inquiry*, 61, 818–850.

Ben-Akiva, M., A. De Palma, and K. Isam (1991): “Dynamic network models and driver information systems,” *Transportation Research Part A: General*, 25, 251–266.

Ben-Elia, E., R. Di Pace, G. N. Bifulco, and Y. Shiftan (2013): “The impact of travel information’s accuracy on route-choice,” *Transportation Research Part C: Emerging Technologies*, 26, 146–159.

Braess, D. (1968): “Über ein Paradoxon aus der Verkehrsplanung,” *Unternehmensforschung*, 12, 258–268.

Cabannes, T., F. Shyu, E. Porter, S. Yao, Y. Wang, M. A. S. Vincentelli, S. Hinardi, M. Zhao, and A. M. Bayen (2018a): “Measuring regret in routing:

Assessing the impact of increased app usage," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, IEEE, 2589–2594.

Cabannes, T., M. A. S. Vincentelli, A. Sundt, H. Signargout, E. Porter, V. Fighiera, J. Ugirumurera, and A. M. Bayen (2018b): "The impact of GPS-enabled shortest path routing on mobility: A game theoretic approach," *Transportation Research Board 97th Annual Meeting*, 21.

Camerer, C. and T.-H. Ho (1999): "Experience-weighted attraction learning in normal form games," *Econometrica*, 67, 827–874.

Camerer, C. F., T.-H. Ho, and J.-K. Chong (2002): "Sophisticated experience-weighted attraction learning and strategic teaching in repeated games," *Journal of Economic Theory*, 104, 137–188.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008): "Bootstrap-based improvements for inference with clustered errors," *Review of Economics and Statistics*, 90, 414–427.

Chen, W., S.-Y. Liu, C.-H. Chen, and Y.-S. Lee (2011): "Bounded memory, inertia, sampling and weighting model for market entry games," *Games*, 2, 187–199.

Cheung, Y.-W. and D. Friedman (1997): "Individual learning in normal form games: Some laboratory results," *Games and Economic Behavior*, 19, 46–76.

Cohen, J. (1988): "Statistical power analysis for the behavioral sciences," *Lawrence Erlbaum Publishers.*

Denant-Boèmont, L. and R. Petiot (2003): "Information value and sequential decision-making in a transport setting: An experimental study," *Transportation Research Part B: Methodological*, 37, 365–386.

Emmerink, R. H., E. T. Verhoef, P. Nijkamp, and P. Rietveld (1996): "Information provision in road transport with elastic demand: A welfare economic approach," *Journal of Transport Economics and Policy*, 117–136.

Erev, I. and G. Barron (2005): "On adaptation, maximization, and reinforcement learning among cognitive strategies," *Psychological Review*, 112, 912.

Erev, I., E. Ert, O. Plonsky, D. Cohen, and O. Cohen (2017): "From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience," *Psychological Review*, 124, 369.

Erev, I., E. Ert, O. Plonsky, and Y. Roth (2023): "Contradictory deviations from maximization: Environment-specific biases, or reflections of basic properties of human learning?" *Psychological Review*, 130, 640.

Erev, I., E. Ert, and A. E. Roth (2010): "A choice prediction competition for market entry games: An introduction," *Games*, 1, 117–136.

Erev, I. and A. Rapoport (1998): "Coordination, "magic", and reinforcement learning in a market entry game," *Games and Economic Behavior*, 23, 146–175.

Erev, I. and A. E. Roth (1998): "Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria," *American Economic Review*, 848–881.

——— (2014): "Maximization, learning, and economic behavior," *Proceedings of the National Academy of Sciences*, 111, 10818–10825.

Estes, W. K. (1950): "Toward a statistical theory of learning," *Psychological Review*, 57, 94.

Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10, 171–178.

Fischbacher, U. and S. Gächter (2010): "Social preferences, beliefs, and the dynamics of free riding in public goods experiments," *American Economic Review*, 100, 541–56.

Fischbacher, U., S. Gächter, and E. Fehr (2001): "Are people conditionally cooperative? Evidence from a public goods experiment," *Economics Letters*, 71, 397–404.

Fudenberg, D. and D. K. Levine (1995): "Consistency and cautious fictitious play," *Journal of Economic Dynamics and Control*, 19, 1065–1089.

Ho, T. H., C. F. Camerer, and J.-K. Chong (2007): "Self-tuning experience weighted attraction learning in games," *Journal of Economic Theory*, 133, 177–198.

Ho, T. H., X. Wang, and C. F. Camerer (2008): "Individual differences in EWA learning with partial payoff information," *The Economic Journal*, 118, 37–59.

Ioannou, C. A. and J. Romero (2012): "Strategic learning with finite automata via the EWA-lite model," Purdue University, Department of Economics.

Kahneman, D. (1988): "Experimental economics: A psychological perspective," in *Bounded Rational Behavior in Experimental Games and Markets*, Springer, 11–18.

Kremer, I., Y. Mansour, and M. Perry (2014): "Implementing the "wisdom of the crowd"," *Journal of Political Economy*, 122, 988–1012.

Liu, S., L. Guo, S. M. Easa, H. Yan, H. Wei, and Y. Tang (2020): "Experimental study of day-to-day route-choice behavior: Evaluating the effect of ATIS market penetration," *Journal of Advanced Transportation*, 2020.

Lu, X., S. Gao, and E. Ben-Elia (2011): "Information impacts on route choice and learning behavior in a congested network: Experimental approach," *Transportation Research Record*, 2243, 89–98.

Macfarlane, J. (2019): "Your navigation app Is making traffic unmanageable," *IEEE Spectrum*, https://spectrum.ieee.org/your-navigation-app-is-making-traffic-unmanageable/.

Madrigal, A. C. (2018): "The perfect selfishness of mapping apps," *The Atlantic*, https://www.theatlantic.com/technology/archive/2018/03/mapping-apps-and-the-price-of-anarchy/555551/.

Mak, V., E. J. Gisches, and A. Rapoport (2015): "Route vs. segment: An experiment on real-time travel information in congestible networks," *Production and Operations Management*, 24, 947–960.

Moore, J. (2016): "Waze, please stop ruining Los Angeles," *GQ*, https://www.gq.com/story/waze-ruining-los-angels-side-streets/.

Plonsky, O., K. Teodorescu, and I. Erev (2015): "Reliance on small samples, the wavy recency effect, and similarity-based learning," *Psychological Review*, 122, 621.

Rapoport, A., E. J. Gisches, T. Daniel, and R. Lindsey (2014): "Pre-trip information and route-choice decisions with stochastic travel conditions: Experiment," *Transportation Research Part B: Methodological*, 68, 154–172.

Rapoport, A., T. Kugler, S. Dugar, and E. J. Gisches (2009): "Choice of routes in congested traffic networks: Experimental tests of the Braess Paradox," *Games and Economic Behavior*, 65, 538–571.

Romero, J. and Y. Rosokha (2019): "The evolution of cooperation: The role of costly strategy adjustments," *American Economic Journal: Microeconomics*, 11, 299–328.

Roth, A. E. and I. Erev (1995): "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term," *Games and Economic Behavior*, 8, 164–212.

Selten, R. (1967): "Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments," *Beiträge zur experimentellen Wirtschaftsforschung, Tübingen, JCB Mohr (Paul Siebeck)*, 136–168.

Selten, R., T. Chmura, T. Pitz, S. Kube, and M. Schreckenberg (2007): "Commuters route choice behaviour," *Games and Economic Behavior*, 58, 394–406.

Weise, E. (2017): "Waze and other traffic dodging apps prompt cities to game the algorithms," *USATODAY*, https://www.usatoday.com/story/tech/news/2017/03/06/mapping-software-routing-waze-google-traffic-calming-algorithmsi/98588980/.

Wijayaratna, K. P., V. V. Dixit, L. Denant-Boemont, and S. T. Waller (2017): "An experimental study of the Online Information Paradox: Does en-route information improve road network performance?" *PLoS One*, 12, e0184191.

Yakobi, O., D. Cohen, E. Naveh, and I. Erev (2020): "Reliance on small samples and the value of taxing reckless behaviors," *Judgment and Decision Making*, 15, 266–281.

Zhang, R. and E. T. Verhoef (2006): "A monopolistic market for advanced traveller information systems and road use efficiency," *Transportation Research Part A: Policy and Practice*, 40, 424–443.

Online Appendices

We include ten online appendices. Appendices A and B describe the EWA-lite learning model and the Small Sample models, respectively. Appendix C provides the observed probability transition matrix of adjustment to the relative cost of the two routes. Appendix D presents the estimation results of the two classes of models. In Appendix E, we discuss carryover effects from one treatment to the next. Appendix F shows the outcomes of a robustness check of our main results using different subsets of the experimental data. A comparison of outcomes between the informed and uninformed players under $\mathtt{Partial}$ is presented in Appendix G. Appendix H provides a more detailed analysis of the $\mathtt{Endogenous}$ treatment, Part 2 of our experiment. Appendix I considers the data from the two experimental laboratories in which the study was conducted separately. Finally, Appendix J consists of the instructions and quizzes used in the experiment.

A The EWA-lite Learning Model

A.1 The Model

The EWA-lite learning model was originally formulated by Ho et al. (2007). We follow their description of the model here. The model involves the updating of two state variables after each play of the game: (i) "observation-equivalents" of experience, and (ii) the attraction of each action. "Observation-equivalents" is a measure of experience after period t for player i and is denoted as $N_i(t)$.

$$N_i(t) = \phi_i(t) \cdot N_i(t-1) + 1,[22] \tag{1}$$

The "attraction" of the pure strategy $j \in \{0, 1\}$ after period t for player i is denoted as $A_i^j(t)$.

$$A_i^j(t) = \frac{\phi_i(t) \cdot N_i(t-1) \cdot A_i^j(t-1) + \delta_i^j(t) \cdot \hat{\pi}_i(s_i^j, s_{-i}(t))}{N_i(t)}. \tag{2}$$

Let $\hat{\pi}_i(s_i^j, s_{-i}(t))$ be player i's payoff from pure strategy s_i^j given that other players use strategy $s_{-i}(t)$ in period t. It may denote an actual, a foregone, or an expected foregone payoff, depending on whether s_i^j was actually chosen or not and the information that becomes available at the end of a period. The actual payoff received by player i after selecting pure strategy $s_i(t)$ is also written as $\pi_i(s_i(t), s_{-i}(t)) = \pi_i(t)$. The

[22] Ho et al. (2007) simplify the function by setting $\kappa = 0$ in the EWA-lite model.

foregone payoff for informed players corresponds to the payoff that the unselected pure strategy would have yielded, given others' strategies. The estimated foregone payoff for uninformed players is calculated under the assumption of uniform beliefs over the outcomes of the possible actions of other players that could have resulted in the observed payoffs. See Appendix A.2 for details.

The strategy of players other than i, $s_{-i}(t)$, is proxied by the number of other players selecting strategy $s = 1$ in period t. By making this simplification, we narrow the set of strategies of others that must be considered to $S_{-i} = \{0, 1, 2, 3, 4, 5\}$.[23]

The initial values of $N_i(t)$ and $A_i^j(t)$, $N(0)$ and $A^j(0)$, are pre-selected and identical across players. These initial values reflect pregame experience and introspection. Following Ho et al. (2007), we set the initial experience $N(0) = 1$, which is not critical after a few repetitions since its influence fades rapidly over time. Additionally, we assume that all initial attractions are equal and set the initial value $A^j(0) = 1, \forall j \in \{0, 1\}$.[24]

Two self-tuning functions govern experience accumulation and strategy reinforcement. The first self-tuning function for player i after period t is the change detector function, $\phi_i(t)$. It captures both forgetting and a player's perception of how quickly the learning environment is changing. When others' behavior tends to be consistent, a larger weight is put on experience in the more distant past. When others' behavior changes frequently and drastically, however, a player weights distant past experience less heavily.

$\phi_i(t) \in [0, 1)$ is based on the difference between the immediate history vector and the cumulative history vector of other players' actions. The immediate history of others' actions s_{-i}^k for an informed player is denoted as $r_i^{Inf,k}(t) = I^{Inf}(s_{-i}^k, s_{-i}(t)) \in \{0, 1\}$. $I^{Inf}(s_{-i}^k, s_{-i}(t))$ is an indicator function: it equals 1 if $s_{-i}^k = s_{-i}(t)$ and 0 otherwise. The immediate history of other players' actions s_{-i}^k for the $UInf$-type player is denoted as $r_i^{UInf,k}(t) = I^{UInf}(s_{-i}^k, s_{-i}(t)) \in \{0, 1/5, 1/4, 1/3, 1/2, 1\}$. Here, $I^{UInf}(s_{-i}^k, s_{-i}(t))$ is the belief function, and its value represents the chance that $s_{-i}^k = s_{-i}(t)$ under a uniform belief assumption. This assumption on beliefs is made to reflect the fact that an un-

[23]It is common to make simplifying assumptions to reduce the set of other players' strategies that are included in the model's computation. Anderson and Camerer (2000) assume that players use agent-normal-form reasoning in a sender-receiver game to simplify sender's strategy; Ioannou and Romero (2012) focus on finite automata to constitute strategies. These papers simplify the strategy set to reduce the computational burden and to reflect bounded rationality, specifically the notion that a player may not consider all feasible strategies but rather limit himself to less-complex strategies.

[24]Ho et al. (2007) mention four ways to pin down the initial attractions. Among those alternatives, we assume uniformly distributed first-period choices to make the EWA-lite model more comparable to the Small Sample Models.

informed player can not directly observe $s_{-i}(t)$, but can imperfectly infer $s_{-i}(t)$ from their received payoff $\pi_i^j(t)$ and their selected strategy $s_i(t)$. In particular, we assume that uninformed players have a uniform belief over the number of other players taking each route that could lead to the payoffs they have observed.[25], and infer their expected foregone payoff based on this belief.

The cumulative history of others' actions s_{-i}^k for player i with type $y \in \{UInf, Inf\}$ after period t is

$$h_i^{y,k}(t) = \frac{\sum_{\tau=1}^{t} I^y(s_{-i}^k, s_{-i}(t))}{t}.$$

The surprise index, $S_i(t)$, measures the difference between the most recent observation and the historical average. To be specific, $S_i(t)$ sums up the squared differences between the cumulative history vector and the immediate history vector; that is,

$$S_i(t) = \sum_{k=1}^{m_{-i}} (h_i^{y,k}(t) - r_i^{y,k}(t))^2,$$

where m_{-i} equals 6, representing the number of possible total numbers of other drivers taking route a. $S_i(t) \in [0, 2)$. It is zero when the last action profile of others is the same one they have always chosen previously. It approaches two when the other players suddenly switch to a brand new action profile in the last period after consistently playing a specific action profile.

The change-detector function, $\phi_i(t)$, transforms the surprise index into a weight between zero and one, according to the relation:

$$\phi_i(t) = 1 - 0.5 S_i(t).$$

$\phi_i(t) \in (0, 1]$ is the weight accorded to lagged attraction, which is interpreted with the intuition that less weight should be put on distant experience when a player senses that other players are changing their strategies.

The attention function, $\delta_i^j(t)$, specifies the reinforcement weight on the actual, foregone, and expected foregone payoffs. Intuitively, players with limited attention are likely to focus on strategies that would have yielded (weakly) higher payoffs than the actual payoff they received. Accordingly, $\delta_i^j(t)$ is defined as

[25] Camerer et al. (2002) make the same assumption in a p-beauty contest game in which subjects who lost did not know the winning number. The average payoff rule used to calculate beliefs about the unobserved foregone payoff in Ho et al. (2008) is similar in spirit. Replacing the uniform belief with any centrosymmetric belief has no impact on the simulation results for the game studied here.

$$\delta_i^j(t) = \begin{cases} 1 \text{ if } \hat{\pi}_i(s_i^j, s_{-i}(t)) \geq \pi_i(t), \\ 0, \text{ otherwise}. \end{cases}$$

That is, the selected strategies, as well as all unselected strategies that could have yielded weakly higher payoffs compared to the received payoff, are reinforced with a weight of one. Conversely, players do not reinforce unchosen strategies which would have yielded payoffs strictly lower than the one that they chose. For informed players, the other route is reinforced if choosing it would have led to a greater payoff than they received for the period. For the uninformed, it is reinforced if it would have led to a greater expected payoff under uniform beliefs about the number of drivers taking each route than they received for the period.

The self-tuning system reflects two properties of learning. First, the attention function is an update of behavior in the direction of the ex-post best response. Second, the change detector function is a refocusing of attention after a change in the decision environment. Players explore a wide range of strategies and reinforce superior strategies before equilibration takes place and strategy choices are locked in. When the environment changes, players reallocate their attention and reinitiate an exploration process.

At the beginning of each period, players choose actions with probabilities determined by attractions. More attractive actions are more likely to be chosen. We follow the literature by using the logit functional form to model the choice probabilities as a function of the attractions.

$$P_i^j(t+1) = \frac{e^{\lambda \cdot A_i^j(t)}}{\sum_{k=1}^{m_i} e^{\lambda \cdot A_i^k(t)}}, \tag{3}$$

with $m_i = 2$ representing the number of actions for a player in our game.

The parameter λ measures the response of players to the difference between attractions: $\lambda = 0$ represents a situation in which each action is equally likely to be chosen. As λ becomes large, the function converges to choosing the best response given one's attractions with probability 1. λ is the only parameter that we estimate in our data.

Following Romero and Rosokha (2019), we include Table 9 to present an outline of the EWA-lite learning algorithm we have used in our analysis.

A.2 The Inference Made by the Uninformed

We summarize the inference by the uninformed players in Tables 10 and 11. In particular, Table 10 shows the belief transition matrices between realized payoffs and the resulting

Table 9: EWA-lite learning algorithm summary

Given λ, set initial values of $N(0) = 1$, and $A^j(0) = 1$, $\forall j \in \{0, 1\}$
Assign types, $\mathtt{UInf}$ and $\mathtt{Inf}$, to all i
For $t = 1$ to length T do
    For all i do
        Determine probability $P_i^j(t)$ from $A_i(t-1)$ and λ
        Sample the chosen strategy $s_i(t)$ according to $P_i(t)$
    Get $\sum_{k=1}^{6} s_k(t)$
    For all i do
        Independently randomly draw value, $\nu_i(t)$
        Assess profit $\pi_i(s_i(t), s_{-i}(t))$ and $\hat{\pi}_i(s_i^j, s_{-i}(t))$
        Evaluate $\phi_i(t)$ and $\delta_i^j(t)$
        Obtain $N_i(t)$ and $A_i^j(t)$
    Set $t = t + 1$

beliefs on others' actions for uninformed players when their choice is route a ($s_i = 1$). Table 11 provides the case when the route choice for the uninformed is route b ($s_i = 0$). Column (2) of each table lists a series of payoff ranges. Values within each range share an identical belief about others' actions, $s_{-i} \in \{0, 1, 2, 3, 4, 5\}$, shown in columns (3) to (8). Column (9) reports the expected number of others selecting action 1 under the uniform belief assumption. The inferred foregone payoff and the corresponding value of the attention function of the unchosen action are given in columns (10) and (11). For ease of comparison, the value of the attention function for an informed player is also displayed in column (12).

A.3 Simulation results over 40 Periods

This section includes two parts. Section A.3.1 compares the $\mathtt{No}$, $\mathtt{Partial}$, and $\mathtt{Full}$ treatments with respect to efficiency and route choice under various response sensitivities.[26] Section A.3.2 explores how the implied value of information depends on the configuration of information held by others.

A.3.1 Efficiency and Route Choice

The unique free parameter of EWA-lite, the response sensitivity λ, determines the learning process in the repeated game. We first focus on a few representative points in a relatively large range between zero and 16, which covers the plausible λ values for

[26]The contrast between the outcomes of informed and uninformed players under the $\mathtt{Partial}$ treatment is discussed in Appendix G.

Table 10: The belief transition matrix for an uninformed player i choosing $s_i = 1$ (route a).

$s_i = 1$:		(3)–(8) Others' action s_{-i}						(9)	(10)	(11)	(12)
(1)	(2)	= 0	= 1	= 2	= 3	= 4	= 5	$E(s_{-i})$	$E(\pi_i^0)$	δ_i^0	δ_i^0 for an informed i
	(14,15]	1	0	0	0	0	0	0	$\pi_i^1 - 11$	0	0
	(13,14]	1/2	1/2	0	0	0	0	0.5	$\pi_i^1 - 9.5$	0	0
	(12,13]	1/3	1/3	1/3	0	0	0	1	$\pi_i^1 - 8$	0	0
	(11,12]	1/4	1/4	1/4	1/4	0	0	1.5	$\pi_i^1 - 6.5$	0	0
	(10,11]	1/5	1/5	1/5	1/5	1/5	0	2	$\pi_i^1 - 5$	0	=1 if $s_{-i} = 4$
π_i^1	[5,10]	1/6	1/6	1/6	1/6	1/6	1/6	2.5	$\pi_i^1 - 3.5$	0	=1 if $s_{-i} = 4, 5$
	[4,5)	0	1/5	1/5	1/5	1/5	1/5	3	$\pi_i^1 - 2$	0	=1 if $s_{-i} = 4, 5$
	[3,4)	0	0	1/4	1/4	1/4	1/4	3.5	$\pi_i^1 - 0.5$	0	=1 if $s_{-i} = 4, 5$
	[2,3)	0	0	0	1/3	1/3	1/3	4	$\pi_i^1 + 1$	1	=1 if $s_{-i} = 4, 5$
	[1,2)	0	0	0	0	1/2	1/2	4.5	$\pi_i^1 + 2.5$	1	1
	[0,1)	0	0	0	0	0	1	5	$\pi_i^1 + 4$	1	1

Notes: Column (2) divides the realized payoff π_i^1 into 11 ranges. Payoff values within each range share an identical belief regarding others' actions, s_{-i}. Columns (3) to (8) specify the probabilities that player i assigns to each $s_{-i} \in \{0, 1, 2, 3, 4, 5\}$. Column (9) reports the expected number of other players choosing route a. The expected foregone payoff for choosing route b ($s_i = 0$) is given in column (10). Column (11) presents the value of the attention function of route b for the uninformed player i, δ_i^0. Finally, column (12) shows the value of the attention function for a counterfactual player i who is informed.

Table 11: The belief transition matrix for an uninformed player i choosing $s_i = 0$ (route b).

$s_i = 0$:		(3)–(8) Others' action s_{-i}						(9)	(10)	(11)	(12)
(1)	(2)	= 0	= 1	= 2	= 3	= 4	= 5	$E(s_{-i})$	$E(\pi_i^1)$	δ_i^1	δ_i^1 for an informed i
	(12,14]	0	0	0	0	0	1	5	$\pi_i^0 - 4$	0	0
	(10,12]	0	0	0	0	1/2	1/2	4.5	$\pi_i^0 - 2.5$	0	0
	(8,10]	0	0	0	1/3	1/3	1/3	4	$\pi_i^0 - 1$	0	=1 if $s_{-i} = 3$
	(6,8]	0	0	1/4	1/4	1/4	1/4	3.5	$\pi_i^0 + 0.5$	1	=1 if $s_{-i} = 2, 3$
	(4,6]	0	1/5	1/5	1/5	1/5	1/5	3	$\pi_i^0 + 2$	1	=1 if $s_{-i} = 1, 2, 3$
π_i^0	4	1/6	1/6	1/6	1/6	1/6	1/6	2.5	$\pi_i^0 + 3.5$	1	=1 if $s_{-i} = 0, 1, 2, 3$
	[2,4)	1/5	1/5	1/5	1/5	1/5	0	2	$\pi_i^0 + 5$	1	=1 if $s_{-i} = 0, 1, 2, 3$
	[0,2)	1/4	1/4	1/4	1/4	0	0	1.5	$\pi_i^0 + 6.5$	1	1
	[-2,0)	1/3	1/3	1/3	0	0	0	1	$\pi_i^0 + 8$	1	1
	[-4,-2)	1/2	1/2	0	0	0	0	0.5	$\pi_i^0 + 9.5$	1	1
	[-6,-4)	1	0	0	0	0	0	0	$\pi_i^0 + 11$	1	1

Notes: Column (2) divides the realized payoff π_i^0 into 11 ranges. Payoff values within each range share an identical belief regarding others' actions, s_{-i}. Columns (3) to (8) specify the probabilities that player i assigns to each $s_{-i} \in \{0, 1, 2, 3, 4, 5\}$. Column (9) reports the expected number of other players choosing route a. The expected foregone payoff for choosing route a ($s_i = 1$) is given in column (10). Column (11) presents the value of the attention function of route a for the uninformed player i, δ_i^1. Finally, column (12) shows the value of the attention function for a counterfactual player i who is informed.

our experiment. We simulate 40 periods of EWA-lite learning 10,000 times for each $\lambda \in \{0.125, 0.25, 0.5, 1, 2, 4, 8, 16\}$ under the No, Partial, and Full treatments. The average results are presented in Figure 8. Panel 8a shows the average traffic cost and panels 8b and 8c illustrate the average and the standard deviation of the number of players who select route a.

(a) Average cost

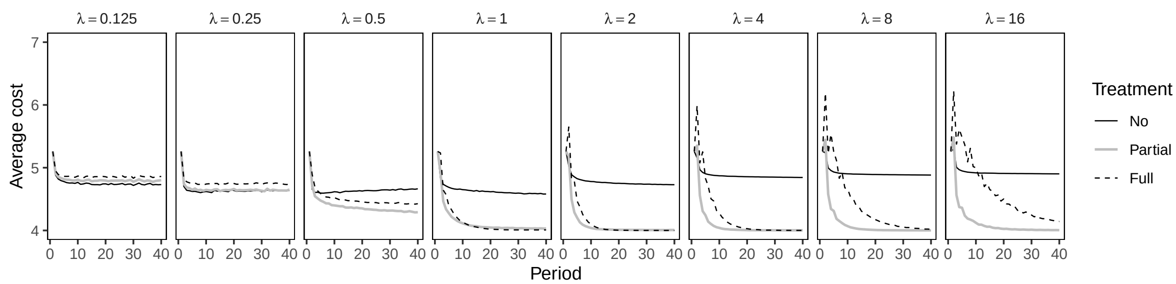

(b) Average number of players choosing route a

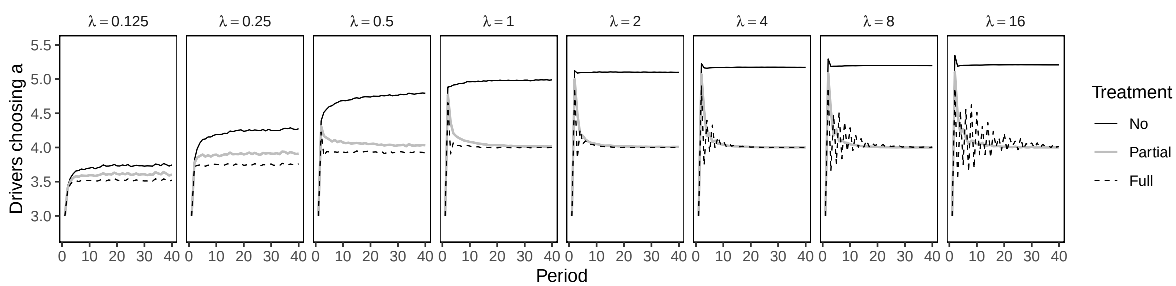

(c) Standard deviation of the number of players choosing route a

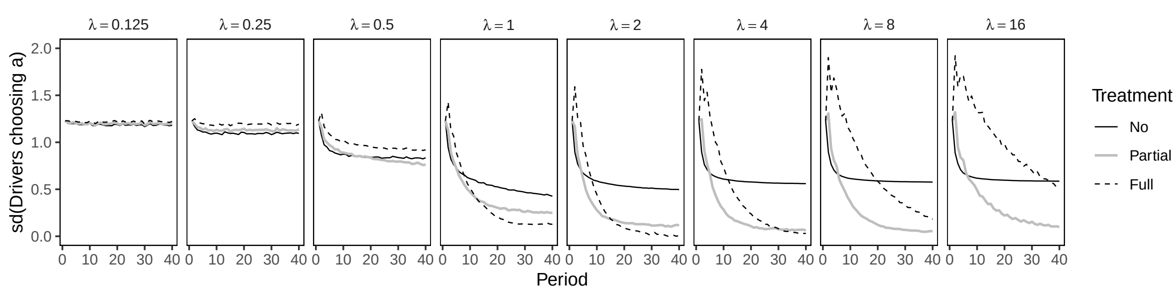

Figure 8: Dynamics of cost and traffic patterns over 40 periods. 10,000 simulations for a six-player group are conducted under each of the No, Partial, and Full treatments for $\lambda \in \{0.125, 0.25, 0.5, 1, 2, 4, 8, 16\}$. Panel (a) plots the average cost per individual. Panel (b) shows the average number of players choosing route a out of a group of six. Panel (c) presents the standard deviation of the number of players in a group choosing route a.

Figure 8 suggests that comparisons of the travel cost under the three treatments depend on the sensitivity parameter, as well as the number of periods that have elapsed.

We define the first ten periods as the short run and periods 31 to 40 as the long run. To have a more comprehensive view of the model predictions, we conduct 1,000 simulations for each λ value between 0.01 and 16 with a step size of 0.01. The three panels in Figure 9 illustrate the average cost over the first ten periods, periods 31 to 40, and all 40 periods, respectively.

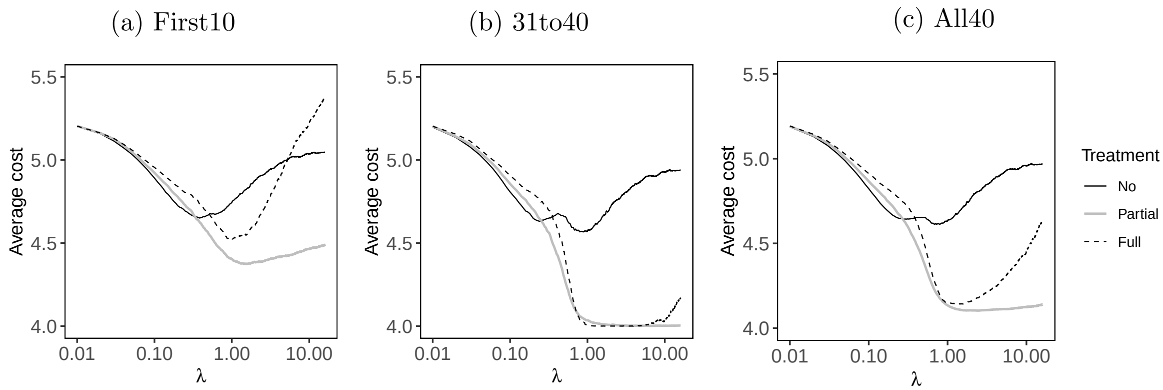

Figure 9: The average cost in 1,000 simulations for each of the three treatments when λ is between 0.01 and 16 with a step size of 0.01. Panels (a) to (c) show the average cost for the first ten periods, periods 31 to 40, and all 40 periods, respectively. The horizontal axis is on a logarithmic scale.

Figure 9 reveals some interesting patterns. First, No adoption of past information can be consistently better than both Partial and Full information adoption when λ takes values lower than 0.25, regardless of the length of experience individuals have. Therefore, providing congestion information leads to efficiency loss for small λs. Second, when λ is greater than 0.34, providing congestion information to one-half of the people is beneficial. Partial adoption leads to a lower average cost than No adoption, as seen from the three panels. Moreover, Full adoption is not optimal except for some λ values in the long run. Indeed, Full results in the lowest efficiency when λ is small and is always suboptimal to Partial in the short run or during all 40 periods.

Figures 8b and 8c help explain the intuition behind the above findings. First, for $\lambda = 0.125$, the traffic flow on route a under the No treatment is closer to the socially optimal level of four than under either Partial or Full. Nevertheless, the degree of fluctuation in route choices is similar among the three treatments. Therefore, No is optimal for $\lambda = 0.125$ and values nearby. Nevertheless, when $\lambda \geq 0.25$, players under the No treatment systematically choose route a more than is socially optimal. That is to say, on average, more than four players select route a. This bias toward route a

is caused by the imperfect inference of uninformed players about the route choices of other players. As Tables 10 and 11 in Appendix A.2 suggest, uninformed players are less likely to reinforce route $\mathtt{b}$ when choosing route $\mathtt{a}$ yet are more likely to reinforce route $\mathtt{a}$ when choosing route $\mathtt{b}$, compared to what informed players do. The cost of route $\mathtt{a}$ is less sensitive to others' decisions than the cost of choosing route $\mathtt{b}$. If too many players choose route $\mathtt{b}$, the cost increases rapidly, discouraging its use.

Second, Figure 8b suggests that Partial and Full do not exhibit significant systematic route choice bias. When $\lambda \geq 1$, Partial and Full gradually converge to an average of four players selecting route $\mathtt{a}$. As seen from Figure 8c, the main difference between the Full and Partial conditions is the degree of fluctuation in choices. Generally, the standard deviation of the number of individuals choosing route $\mathtt{a}$ is greater under Full than that under Partial, especially in earlier periods. The intuition behind this is that informed players move in the direction of the ex-post best response, which only improves the traffic pattern if the number of informed players is small. If too many drivers are informed, too many people change their behavior in response to congestion, creating congestion on new routes.

A greater sensitivity of behavior to attractions does not necessarily lead to more efficient outcomes. Specifically, a greater λ means choices are more responsive to past payoffs, particularly among informed players. When λ is close to zero, the randomization element in the logit response rule gains importance, and players can not exploit the more attractive option. On the other hand, when λ is too high, highly sensitive behavioral responses will exacerbate fluctuations, delaying convergence to the final distribution of traffic flows and eroding the value of past traffic information. Indeed, the optimal λ varies by the extent of the availability of past traffic information. More informed players mean a higher likelihood of overreaction but a lower likelihood of underreaction. According to Figure 9, the lowest average cost over the entire 40 periods under No is 4.59, which occurs for $\lambda = 0.77$. Partial and Full attain their lowest costs of 4.09 and 4.12 at $\lambda = 2.33$ and $\lambda = 1.42$, respectively. Hence, the information adoption level and the scale of the sensitivity parameter together determine the performance of a traffic system.

A.3.2 The Value of Information

In an interactive environment, the impact of having information depends on the proportion of informed players. Figure 10 presents the implied information value per period during 40 periods under six conditions where the number of others being informed is between zero and five (C0 to C5). The implied value of information equals the difference in cost when a player is informed vs. uninformed for a given number of others with information. It is calculated by subtracting the average cost of the informed from the average cost of the uninformed in a group with one fewer informed player. The implied information value in 1,000 simulations under each condition for λs between 0.01 and 16 with a step size of 0.01 is presented.

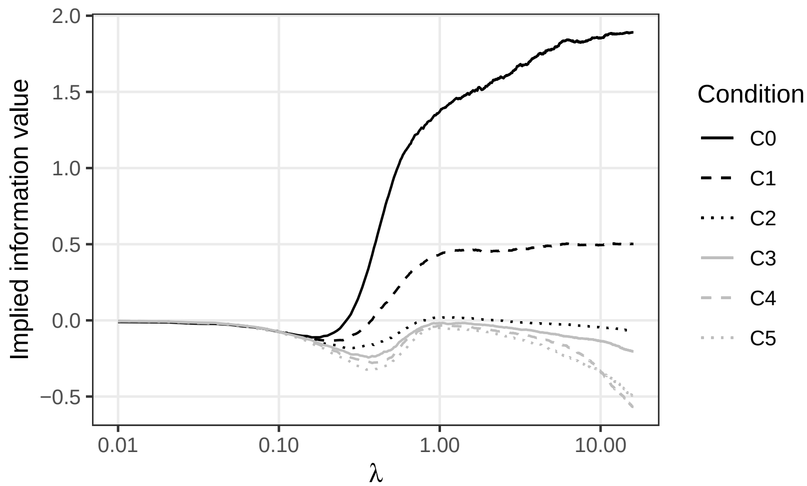

Figure 10: The implied value of information per period during 40 periods. C0 to C5 represent six conditions where the number of informed others is between 0 and 5. It shows the implied information value in 1,000 simulations under each condition for λs between 0.01 to 16 with a step size of 0.01. The horizontal axis is on a logarithmic scale.

Figure 10 shows that the implied value of information depends on both the sensitivity parameter and the number of informed others. When $\lambda < 0.26$, the implied values of all six conditions are *negative*, indicating that no uninformed players can lower their cost by accessing information. Under condition C0, when the five others are uninformed, an uninformed player can lower their cost by becoming informed if $\lambda \geq 0.26$. The player's cost savings can be greater than 1.5 per period for λ sufficiently high. Under condition C1, when one other player is informed and four are uninformed, accessing information can reduce the cost of an uninformed player if $\lambda \geq 0.37$. The player's cost saving is no greater than 0.5 per period. Under condition C2, choosing to be informed can help an uninformed player decrease their cost by at most 0.02 per period when $0.82 \leq \lambda \leq 2.46$. For conditions C3 to C5, however, uninformed players do not benefit from acquiring information for any value of λ.

There are two other interesting observations that merit some emphasis. First, the model predicts that the information can have a negative value. Thus the model predicts that for a certain range of λs, individuals could refuse the information, even if it were free. Second, for $\lambda \in [0.13, 10.96]$, an expansive range, the value of information is decreasing in the number of other drivers who also have it.

A.4 MLE Estimation

We estimate the EWA-lite learning parameter, λ, on the experimental data from Part 1 using structural maximum likelihood estimation (MLE). The MLE criterion serves as a robustness check for the MNSD criterion adopted in Appendix D in estimating the EWA-lite learning parameter. As introduced beforehand, given λ, the EWA-lite learning model needs to specify three initial values: the initial experience level $N(0) = 1$ and attraction values for the two pure strategies $A^j(0) = 1, \forall j \in \{0, 1\}$. The three values are updated after each period according to equations (1) and (2). The probability of choosing a strategy $j \in \{0, 1\}$ in period t for player i, $P_i^j(t)$, is determined by equation (3). The likelihood function of having route choice $y_{it} \in \{0, 1\}$ for player $i = 1, ..., n$ and period $t = 1, ..., T$ is

$$\mathcal{L}(y_{it}; \lambda) = \prod_{i=1}^{n} \prod_{t=1}^{T} [1 - P_i^{j=1}(t)]^{1-y_{it}} [P_i^{j=1}(t)]^{y_{it}}$$

and the log-likelihood is

$$\ell(y_{it}; \lambda) = \sum_{i=1}^{n} \sum_{t=1}^{T} [(1 - y_{it}) \ln(1 - P_i^{j=1}(t)) + y_{it} \ln(P_i^{j=1}(t))] \tag{4}$$

The maximum likelihood estimator of λ is the value of λ that maximizes the log-likelihood function in equation (4). We compute the pooled EWA-lite estimate over all treatments by treating the entire sample from Part 1 as one "representative" individual. We obtain a point estimate of $\hat{\lambda} = 0.27$ with a standard error of 0.01 using data from Segment 1. The standard error is computed by applying the bootstrap-based clustered error (Cameron et al., 2008), with clustering at the level of the group. We also estimate the parameter for each treatment separately. The point estimates of $\hat{\lambda}$ are 0.24 (se=0.02), 0.27 (se=0.02), and 0.32 (se=0.03) for No, Partial, and Full, respectively. Using data from Segments 1 - 3, we get a point estimate of $\hat{\lambda} = 0.31$ with a standard error of 0.01. The treatment-specific point estimates of $\hat{\lambda}$ are 0.28 (se=0.01), 0.30 (se=0.02), and 0.35 (se=0.02) for No, Partial, and Full. The estimated λs using the MLE and MNSD criteria are comparable. Let $\lambda(T)$ denote the EWA-lite estimate under treatment

T. The relationship of the maximum likelihood estimates for the three treatments is consistent with that based on the MNSD criterion shown in Table 14 in Appendix D: $\lambda(\mathsf{No}) < \lambda(\mathsf{Partial}) < \lambda(\mathsf{Full})$.

B The Small Sample Models

B.1 The Models

We consider two learning models from the family of Small Sample Models (Erev and Roth, 2014). These models assume that players sample from past trials to decide on their actions. In the *Sample of q* model, players draw q trials from the past in which the player receives information concerning the payoff from the relevant option. In the *Sample of q or Less* model, each player i draws $q_i \in \{1, 2, ..., q\}$ past trials in which player i receives information concerning the payoff from the relevant option. q is the sole free parameter of the Small Sample models and q_i is drawn from a discrete uniform distribution $\{1, 2, ..., q\}$ for player i in a Sample of q or Less model. Following (Erev and Roth, 2014), we focus on two models: "Sample of 5" (SSM5) and "Sample of 9 or Less" (SSM1-9). Table 12 summarizes the learning algorithm.

In the SSM5 model, since informed players have the congestion information, the exact foregone payoff of the unselected action is available to them. Therefore, an informed player randomly draws 5 trials from all past trials and selects the option with the highest average payoff in the small sample. Unlike the EWA-lite model, the SSM5 model does not assume the uninformed players infer their foregone payoffs. Instead, if both actions have been selected in the past, for each action, an uninformed player draws a small sample of size 5 from past trials in which the action was chosen. Then, the uninformed player compares the average payoff of the two small samples and selects the action with the highest average payoff in its small sample. If only one action has been selected in the past, the uninformed player will select an action randomly in the current trial.

B.2 Simulations over 40 Periods

This section focuses on the predictions of the Sample of q models. Section B.2.1 compares the No, $\mathsf{Partial}$, and Full treatments with respect to efficiency and route choice across various small sample sizes. Section B.2.2 investigates how the implied value of information is contingent on the arrangement of information held by others.

Table 12: Reliance on small samples learning algorithm summary

1:	Given $\mathtt{q}$, the single parameter of the model "sample of q"
{1':	*Given* $\mathtt{q}$*, the single parameter of the model "sample of q or less"*}
2:	Assign types ($\mathtt{UInf}$ and $\mathtt{Inf}$) to all $\mathtt{i}$
{2':	*Assign types (*$\mathtt{UInf}$ *and* $\mathtt{Inf}$*) and sample sizes (*$q_i \in \{1, ..., q\}$ *with an equal chance) to all* $\mathtt{i}$}
3:	For $t = 1$ do
4:	For all $\mathtt{i}$ do
5:	Sample the chosen action $s_i(1) \in \{0, 1\}$ with an equal chance
6:	Get $\sum_{k=1}^{6} s_k(1)$
7:	For all $\mathtt{i}$ do
8:	Independently randomly draw value, $\nu_i(1)$
9:	Assess profit $\pi_i(s_i(1), s_{-i}(1))$ and $\hat{\pi}_i(s_i^j, s_{-i}(1)), j \in \{0, 1\}$
10:	For $t > 1$ to length T do
11:	For all $\mathtt{i}$ with the type $\mathtt{Inf}$ do
12:	Sample with replacement of $\mathtt{q}$ trials from all past trials {1,...,t-1} with an equal chance
{12':	*Sample with replacement of* q_i *trials from all past trials {1,...,t-1} with an equal chance*}
13:	Select the action $s_i(t)$ with the highest average payoff in that small sample
14:	For all $\mathtt{i}$ with the $\mathtt{UnInf}$ type do
15:	If only one action has been taken beforehand
16:	Sample the chosen action $s_i(t)$ with an equal chance
17:	If both pure strategies have been taken beforehand
18:	For all action $\mathtt{j}$
19:	Sample with replacement of $\mathtt{q}$ trials from past trials in which $\mathtt{j}$ has been chosen
{19':	*Sample with replacement of* q_i *trials from past trials in which* $\mathtt{j}$ *has been chosen*}
20:	Select the action $s_i(t)$ with the highest average payoff in its small sample
21:	Get $\sum_{k=1}^{6} s_k(t)$
22:	For all $\mathtt{i}$ do
23:	Independently randomly draw value, $\nu_i(t)$
24:	Assess profit $\pi_i(s_i(t), s_{-i}(t))$ and $\hat{\pi}_i(s_i^j, s_{-i}(t))$
25:	Set $t = t + 1$

Note: This table summarizes the learning algorithm of the model "Sample of q" and "Sample of q or Less" (italics inside curly braces).

B.2.1 Efficiency and Route Choice

The sole free parameter, the size of a small sample denoted as $\mathtt{q}$, plays a crucial role in shaping the learning dynamics in the repeated game. To gain insight into the models' general predictions, we conduct 1,000 simulations, each spanning 40 periods, for every $q \in \{1, 2, ..., 30\}$, under the $\mathtt{No}$, $\mathtt{Partial}$, and $\mathtt{Full}$ treatments. The six panels in Figure 11 illustrate the average cost and the number of drivers selecting route $\mathtt{a}$ over the first ten periods, periods 31 to 40, and the entire 40-period duration, respectively.

Panels (a) to (c) in Figure 11 reveal some prominent patterns. First, in the early

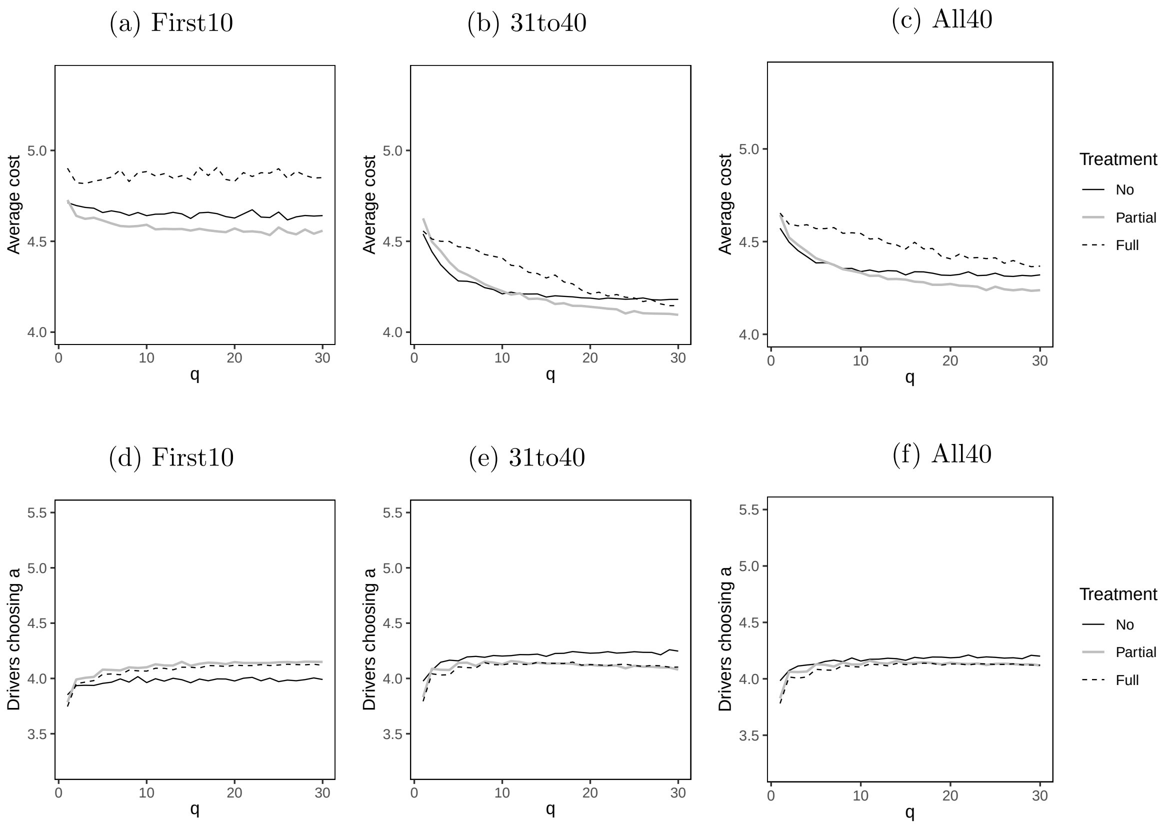

Figure 11: The average cost and the route choice in 1,000 simulations for each of the three treatments by the Sample of q model when q is between 1 and 30 with a step size of 1. Panels (a) to (c) show the average cost for the first ten periods, periods 31 to 40, and all 40 periods, respectively. Panels (d) to (f) present the number of drivers choosing route a for the first ten periods, periods 31 to 40, and all 40 periods, respectively.

periods, Full information adoption is the least efficient, while the Partial condition is the most efficient except when $q = 1$ in the early periods. In other words, providing congestion information to one-half of users or no users is more beneficial than to all. Second, in the late periods, providing information to some or all users is more efficient than providing information to nobody if and only if $q \geq 13$. Third, the comparison between the first ten and last ten periods for a given q indicates that traffic efficiency improves over time under any of the No, Partial, and Full conditions. Fourth, Full is never the most efficient condition. Finally, as the sample size q increases within the interval explored, efficiency improves under all three conditions.

Regarding the route choice behavior depicted in Panels (d) to (f), we find that when $q \geq 3$, the No condition has the least drivers on route a in the first ten periods but has the most drivers on route a in the last ten periods. Partial and Full do not exhibit significant systematic route choice bias irrespective of the duration of individual experience. Additionally, the average number of drivers choosing route $\mathtt{a}$ for each condition is close to the socially optimal number of four.

B.2.2 The Value of Information

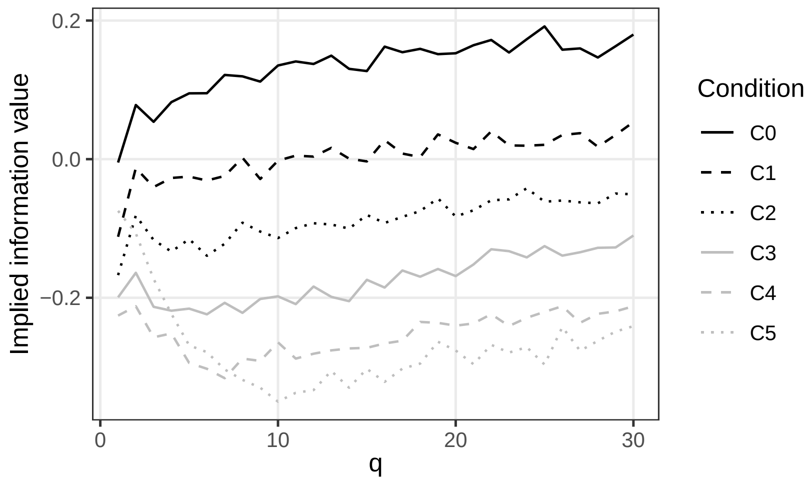

Figure 12: The implied value of information per period during 40 periods for the Sample of $\mathtt{q}$ model. C0 to C5 represent six conditions where the number of informed others varies from 0 to 5. It shows the implied information value in 1,000 simulations under each condition when $\mathtt{q}$ varies from 1 to 30 with a step size of 1.

Figure 12 presents the implied information value per period over 40 periods across six conditions (C0 to C5), varying the number of others informed from zero to five. The implied value of information represents the cost difference between a player being informed vs. uninformed, considering a specific number of others with information. This value is determined by subtracting the average cost of an informed player from the average cost of an uninformed player in a group with one fewer informed player. The implied information values are based on 1,000 simulations for each condition when $\mathtt{q}$ varies from 1 to 30 with a step size of 1.

The figure demonstrates that the implied information value is influenced by the small sample size and the number of informed other drivers. In condition C0, where all five others are uninformed, an uninformed player can reduce their cost by accessing information as long as $\mathtt{q} \geq 2$. The cost savings do not exceed 0.2 points per period across the entire range of $\mathtt{q}$. Under condition C1, when one other player is informed and four are uninformed, becoming informed is profitable for a player when $\mathtt{q} > 15$. However, in C2 to C5, the implied information value is always negative.

C The Observed Transition Probability Matrix

Table 13 illustrates the observed relative-cost transition probabilities under each treatment for data in Segment 1 and Segments 1 - 3. A relative-cost transition probability is the probability of observing the relationship between the congestion cost of routes a and b in Period t (denoted as C_t^a and C_t^b) and their relationship in Period $t+1$, for $t \in \{1, ..., 39\}$. The probabilities are reported for the first ten, the last ten, and all 39 periods under No, $\mathsf{Partial}$, and Full, respectively.

Table 13: Transition probability matrix for relative costs of the two routes

(1) Segment 1

	First 10 Periods			Last 10 Periods			All 39 Periods		
	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$
a. No									
$C^a_t < C^b_t$	0.08	0.09	0.09	0.11	0.08	0.11	0.11	0.10	0.09
$C^a_t = C^b_t$	0.07	0.10	0.14	0.11	0.11	0.11	0.11	0.11	0.11
$C^a_t > C^b_t$	0.11	0.13	0.21	0.08	0.14	0.17	0.08	0.12	0.16
b. Partial									
$C^a_t < C^b_t$	0.09	0.11	0.08	0.08	0.13	0.09	0.11	0.12	0.09
$C^a_t = C^b_t$	0.09	0.17	0.10	0.12	0.16	0.12	0.12	0.15	0.10
$C^a_t > C^b_t$	0.09	0.14	0.14	0.11	0.11	0.08	0.09	0.11	0.10
c. Full									
$C^a_t < C^b_t$	0.13	0.09	0.08	0.09	0.06	0.07	0.10	0.10	0.07
$C^a_t = C^b_t$	0.04	0.06	0.16	0.09	0.38	0.08	0.08	0.24	0.11
$C^a_t > C^b_t$	0.14	0.09	0.20	0.05	0.12	0.07	0.09	0.10	0.10

(2) Segments 1 - 3

	First 10 Periods			Last 10 Periods			All 39 Periods		
	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$
a. No									
$C^a_t < C^b_t$	0.10	0.10	0.09	0.10	0.10	0.09	0.11	0.11	0.08
$C^a_t = C^b_t$	0.09	0.15	0.12	0.10	0.18	0.11	0.10	0.17	0.12
$C^a_t > C^b_t$	0.10	0.10	0.16	0.08	0.12	0.12	0.08	0.11	0.13
b. Partial									
$C^a_t < C^b_t$	0.09	0.13	0.06	0.09	0.13	0.08	0.10	0.13	0.08
$C^a_t = C^b_t$	0.10	0.18	0.13	0.10	0.19	0.13	0.12	0.18	0.12
$C^a_t > C^b_t$	0.09	0.12	0.10	0.09	0.11	0.08	0.09	0.11	0.09
c. Full									
$C^a_t < C^b_t$	0.14	0.10	0.07	0.07	0.08	0.08	0.10	0.10	0.07
$C^a_t = C^b_t$	0.06	0.12	0.14	0.09	0.35	0.10	0.08	0.23	0.11
$C^a_t > C^b_t$	0.11	0.09	0.16	0.06	0.11	0.07	0.09	0.10	0.10

Notes: A relative-cost transition probability is the probability of observing the relationship between C^a_t and C^b_t specified in a given row and the relationship between C^a_{t+1} and C^b_{t+1} specified in a given column, for $t \in \{1, ..., 39\}$. Observations are at the group level. Panels a, b, and c show the probability during the first ten, the last ten, and all 39 periods, for No, Partial, and Full, respectively. Parts (1) and (2) use data from Segment 1 and Segments 1 - 3, respectively.

D Parameter Estimates

We estimate the parameters of the EWA and SSM models for our data. In the estimation, we adopt the grid search method utilizing the mean normalized squared distance (MNSD) criterion, close to the one used in the prediction competition of Erev et al. (2010).[27] Specifically, the MNSD score is the mean of the normalized squared distance (SD) between the model predictions and experimental observations of the route a selection rate and average cost across all three treatments. We divide the data into block 1 (Periods 1 to 20) and block 2 (Periods 21 to 40). The search method aims to find the best-fitting parameter, $\lambda \in \{0.01, 0.02, ..., 10\}$ for EWA-lite and $q \in \{0, 1, ..., 30\}$ for Small Sample models, that minimizes the average of the 12 normalized SD.[28]

The process involves five steps to obtain each normalized SD for a given model parameter: (i) Simulate the model 1,000 times. (ii) Calculate values of the 12 objects in each simulation ($\hat{x}_{ij}$, $1 \in \{1, ..., 1000\}$ and $j \in \{1, ..., 12\}$), treating the average value of each object over the 1,000 simulations as the simulated model prediction. (iii) Compute the SD between the simulated model prediction and the observed statistic (x_j). (iv) Normalize the SD by dividing by the variance of the observed statistic ($D(x_j)$), derived from 1,000 bootstrapping of the experimental data. (v) Calculate the average of the normalized SD over the 12 statistics. Conducting simulations helps mitigate the effect of randomness inherent in the learning model due to sampling, as noted by see Chen et al. (2011). Additionally, normalization ensures comparability among the 12 SD values. In summary, the outlined process solves the following problem:

$$\min_{\lambda \text{ or } q} \frac{1}{12} \sum_{j=1}^{12} \frac{\left(\frac{1}{1000}\sum_{i=1}^{1000} \hat{x}_{ij} - x_j\right)^2}{D(x_j)}$$

The fitted parameters and the resulting MNSD scores for the pooling data, as well as the treatment-specific data, are presented in Table 14. Panels (a) and (b) use data from Segment 1 and from Segments 1 - 3, respectively.

[27]In Appendix A.4, we also estimate the EWA-lite parameter using the Maximum Likelihood Estimation (MLE) method as a robustness check. The estimated λs using the MLE and MNSD criteria are comparable.

[28]In the Small Sample models, $q = 0$ means participants randomly decide between two routes in each period, regardless of previous experience.

Table 14: Model Fitting for the EWA-lite and SSM Models and the Mean Normalized Squared Distance Scores

Model	Treatment	Fitted Parameter	Average cost Block 1 N	Average cost Block 1 P	Average cost Block 1 F	Average cost Block 2 N	Average cost Block 2 P	Average cost Block 2 F	Number of drivers choosing a Block 1 N	Number of drivers choosing a Block 1 P	Number of drivers choosing a Block 1 F	Number of drivers choosing a Block 2 N	Number of drivers choosing a Block 2 P	Number of drivers choosing a Block 2 F	Mean
(a) Segment 1															
	Pooled	0.32	4.57	0.15	0.41	0.41	0.74	26.80	21.35	0.63	26.15	22.61	0.14	13.39	9.8
EWA-lite	N	0.22	8.14			0.35			0.00			0.74			2.3
	P	0.33		0.05			0.31			0.96			0.28		0.4
	F	0.52			1.94			0.07			4.84			0.34	1.8
	Pooled	2	0.03	0.77	0.00	20.82	0.26	3.70	0.38	9.28	0.02	0.05	3.96	3.95	3.6
SSMq	N	1	1.16			3.91			0.66			0.82			1.6
	P	2		0.77			0.26			9.28			3.96		3.6
	F	3			0.01			2.71			0.02			2.19	1.2
	Pooled	2	0.08	0.11	0.03	10.35	0.25	1.77	0.24	0.56	7.93	0.02	0.00	0.26	1.8
SSM1-q	N	1	1.16			4.52			0.70			1.04			1.9
	P	2		0.11			0.25			0.56			0.00		0.2
	F	4			0.36			0.63			0.99			0.57	0.6
(b) Segments 1 - 3															
	Pooled	0.32	30.43	12.42	11.82	15.27	8.25	97.20	49.53	0.03	48.30	145.03	0.56	50.31	39.1
EWA-lite	N	0.21	46.90			16.28			0.90			6.57			17.7
	P	0.37		4.49			0.02			1.53			3.38		2.4
	F	0.53			1.81			0.08			4.20			4.39	2.6
	Pooled	2	3.86	1.63	3.68	6.42	0.10	17.86	0.57	12.52	1.07	4.96	17.13	3.40	6.1
SSMq	N	2	3.86			6.42			0.57			4.96			4.0
	P	3		0.12			2.65			12.79			15.11		7.7
	F	3			3.24			14.06			1.07			1.08	4.9
	Pooled	3	4.17	2.74	1.50	5.91	0.49	9.29	0.04	0.57	2.45	4.48	1.78	0.80	2.9
SSM1-q	N	2	7.24			0.76			0.96			1.99			2.7
	P	3		2.74			0.49			0.57			1.78		1.4
	F	4			0.52			5.15			0.12			0.01	1.5

Notes: This table presents the fitted parameters for EWA-lite, the Sample of q model (SSMq), and the Sample of q or less model (SSM1-q), according to the mean normalized squared distance (MNSQD) criterion. The normalized squared distance scores of the average cost and number choosing route a across block 1 (periods 1 -20) and block 2 (periods 21-40) under treatments No (N), Partial (P) and Full (F) are reported. Panels (a) and (b) show the results using data from Segment 1 only and Segments 1 - 3, respectively.

For the pooled from all treatments data, regardless of the data set, the estimated EWA-lite parameter is $\hat{\lambda} = 0.32$. The Sample of q model (SSMq) and the Sample of q or less model (SSM1-q) share the same estimated parameter, $\hat{q} = 2$, with an exception that $\hat{q} = 3$ for SSM1-q using data from Segments 1 to 3. Using data from Segment 1, the MNSD scores for EWA-lite, SSMq, and SSM1-q are 9.8, 3.6, and 1.8, respectively. The corresponding scores are 39.1, 6.1, and 2.9 for data from Segments 1 - 3. Therefore, when pooling the data across treatments, the Small Sample models produce a lower MNSD score compared to EWA-lite. SSM1-q performs better than SSMq.

For the treatment-specific data, SSM1-q always performs best, yielding a lower score than both SSMq and EWA-lite.[29] Comparing SSMq and EWA-lite, we find that EWA-lite outperforms SSMq under Partial, but SSMq outperforms EWA-lite under No. Their comparative performance under Full depends on the data set analyzed. According to Table 14, the underperformance of EWA-lite is driven in large part by the low accuracy for the average cost under the No treatment. The Full treatment has the highest estimated parameter under both models, perhaps reflecting the existence of more historical information, which makes it easier to sample previous data and estimate counterfactual payoffs.

E Carryover Effect Tests

In this appendix, we test for carryover effects between the No, Partial, and Full treatments. We have six orderings of the three treatments in the first three segments of the experiment. Seven or nine different groups of participants were exposed to each treatment order. 48 groups of data are analyzed.

Figure 13 displays the average cost for all three treatments that appeared in Segments 1 - 3 (S1 to S3). For the Partial and Full treatments, the average cost seems consistent among segments. However, the average cost under the No treatment experiences a sharp decrease in Segment 3.

We use the Mann-Whitney U test and the bootstrapped two-sample t-test to consider whether there are order effects for the three treatments in Part 1. The Mann-Whitney U test statistic for the hypothesis that the average cost in the No treatment comes from the same distribution in Segments 1 and 2 is 150.5 (p-value=0.41). Thus, we can not reject the null hypothesis that the two distributions of the average cost of No in Segments

[29] When $q = 1$, the two Small Sample models are the same. Erev and Roth (2014) note that the Small Sample model with $q = 1$ results in probability matching (Estes, 1950).

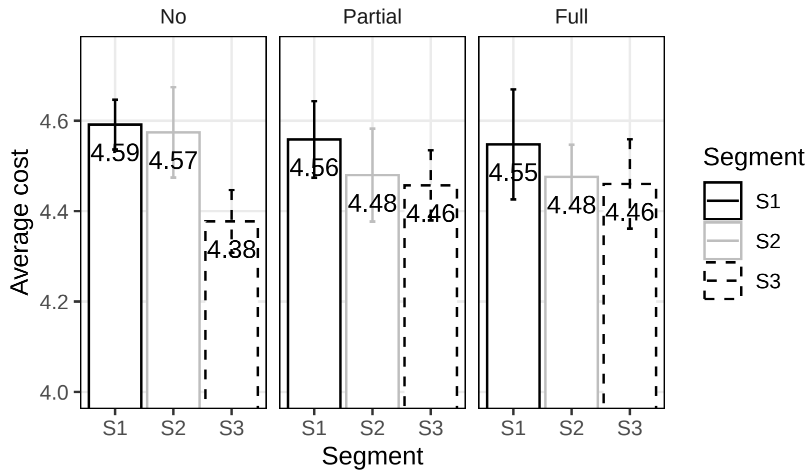

Figure 13: Carryover effects on the average cost over 40 periods. The average cost for each of the No, Partial, and Full treatments in Segments 1 - 3 (S1 to S3) and the corresponding 95% confidence intervals are presented.

1 and 2 are equal. The bootstrapped two-sample t-test, with the null hypothesis that the mean of the average cost of No is equal in Segments 1 and 2, yields z=0.30 (p-value=0.76). Hence, we do not observe a significant order effect between Segments 1 and 2 for the No treatment. However, the Mann-Whitney U test statistics between Segments 1 and 3, and between Segments 2 and 3 are 231 (p-value<0.01) and 203 (p-value<0.01), respectively. The corresponding bootstrapped two-sample t-test statistics are 4.94 (p-value<0.01) and 3.93 (p-value<0.01). Both tests imply a significant order effect for the No treatment in Segment 3 at the significance level of 0.01.

The two tests are also conducted for the Partial treatment. The Mann-Whitney U test statistics for Segments 1 and 2, Segments 1 and 3, and Segments 2 and 3 for Partial are 154 (p-value=0.33), 174 (p-value=0.09) and 148 (p-value=0.46), respectively. The corresponding bootstrapped two-sample t-test statistics are 1.29 (p-value=0.20), 1.68 (p-value=0.09), and 0.34 (p-value=0.74). None of the tests are significant at $p < 0.05$.

Regarding the Full treatment, there is no significant carryover effect. The Mann-Whitney U test statistics for Segments 1 and 2, Segments 1 and 3, and Segments 2 and 3 of Full are 150.5 (p-value=0.41), 158.5 (p-value=0.26), and 136 (p-value=0.78), respectively. The corresponding bootstrapped two-sample t-test statistics are 1.11 (p-value=0.27), 1.19 (p-value=0.23), and 0.26 (p-value=0.80). None of the null hypotheses of equality of any pair of segments can be rejected.

The preceding analysis suggests that there is a carryover effect for the No treatment in Segment 3. Nevertheless, it is possible that there was a carryover effect in the initial periods only, which dissipated later on within the segment. To evaluate this claim, we first focus on the No treatment and compare the average costs in Periods 1 to 10 of Segment 2 to the average cost over all 40 periods in Segment 1. We also compare the average costs in Periods 1 to 10 in Segment 3 to the average cost over all 40 periods in Segment 2. The Mann-Whitney U test statistics between Segments 1 and 2, and between Segments 2 and 3 are 126 (p-value=0.95) and 184.5 (p-value=0.03), respectively. The corresponding bootstrapped two-sample t-test statistics are 0.01 (p-value=0.99) and 1.86 (p-value=0.03). Both tests imply a significant carryover effect for the No treatment in Segment 3 at a significance level of 5%.

We now consider the Partial treatment, comparing the average cost in Periods 1 to 10 in Segment 2 with the average cost over all 40 periods of Segment 1. Another comparison between the average cost in Periods 1 to 10 in Segment 3 with the average cost over all 40 periods in Segment 2 is also conducted. The respective Mann-Whitney U test statistics between Segments 1 and 2, and between Segments 2 and 3 are 169.5 (p-value=0.12) and 143.5 (p-value=0.57). The corresponding bootstrapped two-sample t-test statistics are 1.55 (p-value=0.12) and 0.30 (p-value=0.77). There is no significant carryover effect for the Partial treatment during Periods 1 to 10.

In the Full treatment. The Mann-Whitney U test statistics between Segments 1 and 2, and between Segments 2 and 3 are 98.5 (p-value=0.27) and 102.5 (p-value=0.35), respectively. The respective bootstrapped two-sample t-test statistics are -1.08 (p-value=0.28) and -1.69 (p-value=0.09). None of the test statistics are significant at the significance level of 5%, indicating no evidence of carryover effects.

In the main text, we report results using data from both Segment 1 only and Segments 1 - 3. Using data from Segment 1 generates results free of the carryover effect. Employing data from Segments 1 - 3 allows us to present the totality of the experimental data. Additionally, we conduct pairwise comparisons among the three treatments using data from Segments 1 and 2 in the next Appendix, F, as a robustness check.

F Robustness to Data Set Choice

Table 15 presents summary statistics and pairwise comparisons for the No, Partial, and Full treatments, utilizing data from Segments 1 and 2. Both the Wilcoxon signed-rank and the bootstrap pairwise t-tests are performed for the comparison.

Table 15 reveals that No is significantly more efficient than Full in the first ten periods, while in the last ten periods, Full is significantly more efficient than No. These

Table 15: Average cost and the number of drivers choosing route a using data from Segments 1 and 2

(a) Mean	Average cost			Number choosing route a			Obs.
Periods:	First10	Last10	All40	First10	Last10	All40	
No	4.60	4.56	4.58	4.09	4.12	4.02	32
	(0.21)	(0.23)	(0.15)	(0.37)	(0.31)	(0.19)	
Partial	4.57	4.44	4.52	4.00	3.99	3.96	32
	(0.30)	(0.19)	(0.18)	(0.29)	(0.26)	(0.16)	
Full	4.72	4.34	4.51	3.99	3.98	3.95	32
	(0.34)	(0.26)	(0.19)	(0.28)	(0.22)	(0.13)	
(b) Tests							
No vs. Partial:							
Wilcoxon signed-rank	89	59.5	84	67.5	62	83	16
	[0.29]	[1.00]	[0.42]	[0.13]	[0.26]	[0.20]	
Bootstrapped pairwise t	0.72	0.07	0.72	1.42	1.26	1.40	16
	[0.60]	[0.95]	[0.51]	[0.22]	[0.23]	[0.15]	
No vs. Full:							
Wilcoxon signed-rank	10	119	89	41	74	74.5	16
	[0.01]	[0.01]	[0.29]	[0.78]	[0.44]	[0.18]	
Bootstrapped pairwise t	-3.48	3.23	0.50	-0.16	1.04	1.58	16
	[0.00]	[0.04]	[0.71]	[0.88]	[0.30]	[0.20]	
Partial vs. Full:							
Wilcoxon signed-rank	49.5	99	76	56	57	59.5	16
	[0.35]	[0.11]	[0.70]	[0.84]	[0.89]	[0.68]	
Bootstrapped pairwise t	-0.84	1.58	0.52	-0.29	-0.29	-0.31	16
	[0.46]	[0.22]	[0.60]	[0.80]	[0.79]	[0.78]	

Notes: This table uses data from Segments 1 and 2 for treatments No, Partial, and Full. In panel (a), the units of observation for the first three columns are the average cost for a participant for the first ten periods, for the last ten periods, and for all 40 periods in a given treatment. The next three columns show the average number of players choosing route a in a group for the first ten periods, for the last ten periods, and for all 40 periods in a given treatment. Standard deviations are in parentheses. A standard deviation is calculated by averaging the cost of the members of each group in each period, averaging over the 10 or 40 periods for each group, and then computing the between-group standard deviation of these averages. In panel (b), within-subject pairwise comparisons are conducted with respect to the average cost and the number of players choosing route a among No, Partial, and Full with two statistical tests: the Wilcoxon signed-rank test and the bootstrapped pairwise t-test. All tests are two-sided. The p-values are in square brackets.

patterns align with those only using data solely from Segment 1 and from Segments 1 - 3, as presented in Table 5. Furthermore, although the average cost comparisons between Partial and Full in the first ten periods, between No and Partial in both early and late periods, and between Partial and Full in the last ten periods become statistically insignificant, the sign of their cost differences never reverses. The average numbers of participants choosing route $\mathtt{a}$ across the three treatments are not significantly different. The mean standard deviation of the number of drivers on route $\mathtt{a}$ (MSDNA) for $\mathtt{No}$, $\mathtt{Partial}$, and $\mathtt{Full}$ are 1.09, 1.05, and 1.15 during the first ten periods and are 1.05, 0.94, and 0.82 during the last ten periods. In the early periods, MSDNA is highest under $\mathtt{Full}$. In the later periods, the value is highest under $\mathtt{No}$ but lowest under $\mathtt{Full}$. These findings are consistent with those found in the main text.

When comparing the behavior between the first ten and the last ten periods within data from Segments 1 and 2, the average cost decreases by 0.04 (sd=0.29), 0.13 (sd=0.29), and 0.38 (sd=0.41) for $\mathtt{No}$, $\mathtt{Partial}$, and $\mathtt{Full}$, respectively. Both the bootstrapped pairwise $\mathtt{t}$ and the Wilcoxon signed-rank test statistics suggest that the average cost under $\mathtt{Partial}$ and $\mathtt{Full}$ are lower in the long run compared to those in the short run at significance levels of 5% and 1%, respectively. Concerning the change in the number of drivers on route $\mathtt{a}$, none of the test statistics are significantly different from zero. Regarding MSDNA, its value decreases by 0.04, 0.11, and 0.33 under $\mathtt{No}$, $\mathtt{Partial}$, and $\mathtt{Full}$, respectively. The change increases with the information adoption level. These findings remain consistent with those reported in the main text, where two other data subsets are analyzed.

G Comparison of Uninformed and Informed Players under the Partial Treatment

Under the $\mathtt{Partial}$ treatment, half of the players in a group are informed, while the other half remain uninformed. This setup enables a direct comparison between the two types within the same group. First, we examine how the models predict behavior gaps. Figure 14 illustrates the differences between the informed and uninformed with regard to congestion cost and the selection rate of route $\mathtt{a}$, as predicted by simulations from the EWA-lite and Sample of $\mathtt{q}$ models. Specifically, we vary the parameter λ for EWA-lite from 0.01 to 16 with a step size of 0.01, and the parameter $\mathtt{q}$ for Sample of $\mathtt{q}$ from 1 to 30 with a step size of 1.

In Panel (a) and (c), the cost difference is presented, obtained by subtracting the average cost of the informed players from that of the uninformed players under $\mathtt{Partial}$ as predicted by both models. This difference represents the cost saving for an informed player compared to an uninformed player. The cost difference is always strictly negative

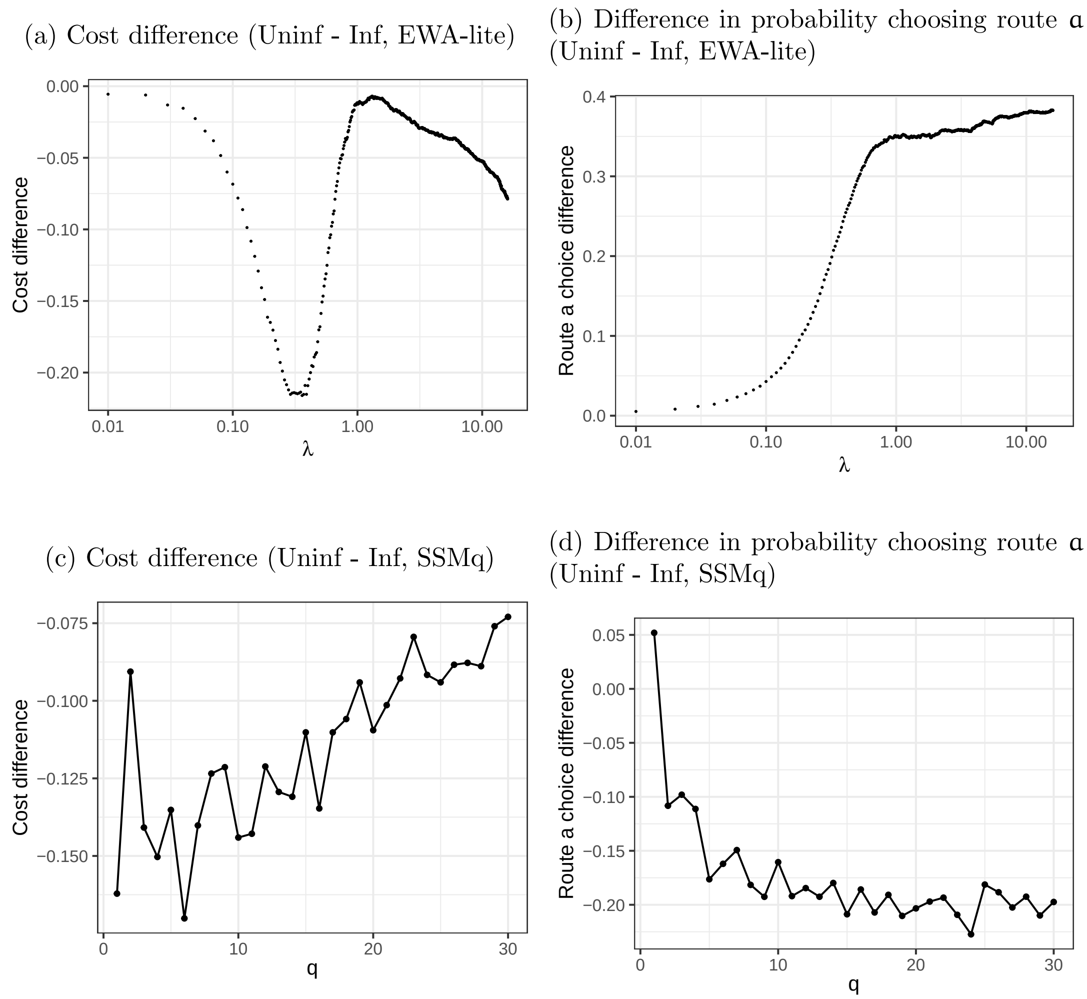

(a) Cost difference (Uninf - Inf, EWA-lite)
(b) Difference in probability choosing route a (Uninf - Inf, EWA-lite)
(c) Cost difference (Uninf - Inf, SSMq)
(d) Difference in probability choosing route a (Uninf - Inf, SSMq)

Figure 14: The differences between the uninformed and the informed under the Partial treatment as predicted by the EWA-lite and Sample of q models are illustrated. For each λ, the parameter for EWA-lite, ranging from 0.01 to 16 with a step size of 0.01, and each q, the parameter for Sample of q (SSMq), ranging from 1 to 30 with a step size of 1, 1,000 simulations are conducted. Panels (a) and (c) display the cost difference, calculated by subtracting the average cost of the informed players from that of the uninformed players in the Partial treatment. Panels (b) and (d) present the route a choice difference, derived by subtracting the probability of choosing route a of the informed players from that of the uninformed players under Partial. The horizontal axis in (a) and (b) is presented on a log scale.

regardless of the value λ and q, indicating that the uninformed pay a lower cost on average than the informed. In other words, past information brings a relative disadvantage to players who adopt.

Panel (b) and (d) showcase the route a choice difference predicted by both models. This difference is derived by subtracting the probability of choosing route a of the informed players from that of the uninformed players under $\mathtt{Partial}$. The EWA-lite model predicts that the route a choice difference consistently maintains a strictly positive value, implying that the uninformed players are more likely to select route a than the informed players, irrespective of the value of λ. However, the Sample of q model predicts that the route a choice difference is strictly negative when $q \geq 2$.

Next, we turn to the experimental data. Over 40 periods, using data from Segment 1, the average costs for uninformed and informed drivers are 4.63 (sd=0.18) and 4.49 (sd=0.22), respectively. Both the Wilcoxon signed-rank and pairwise t-tests find that the informed players bear a lower cost than the uninformed players under $\mathtt{Partial}$, at a significance level of 5%. When using data from Segments 1 - 3, the respective average costs are 4.53 (sd=0.10) and 4.47 (sd=0.09). The two-tailed pairwise t-test suggests that the informed players pay a lower cost than the uninformed players at a significance level of 10%. The corresponding Wilcoxon signed-rank test statistic is 676.5 (p-value=0.14), failing to reject the null hypothesis that the two types bear an equal average cost.

The average probability of choosing route a for the uninformed and informed are 0.67 (sd=0.12) and 0.64 (sd=0.12) in Segment 1 and 0.65 (sd=0.12) and 0.66 (sd=0.12) in Segments 1 - 3. Neither test provides evidence of a significant difference between the uninformed and informed, regardless of the data sets. Thus, under the $\mathtt{Partial}$ treatment, the selection rate of route a for informed and uninformed drivers remains similar over the entire span of 40 periods.

Additionally, we compare the two types in the early and late periods separately. During the first ten periods, the average costs for uninformed and informed drivers are 4.74 (sd=0.34) and 4.64 (sd=0.43) in the data from Segment 1 and are 4.54 (sd=0.33) and 4.53 (sd=0.33) in Segments 1 - 3. The average probability of choosing route a for uninformed and informed drivers are 0.69 (sd=0.11) and 0.65 (sd=0.10) with Segment 1 and are 0.66 (sd=0.13) and 0.66 (sd=0.13) with Segments 1 - 3. Both the Wilcoxon signed-rank and pairwise t-tests fail to reject the null hypothesis of equality between the two types regarding average cost and route a selection rate in each data set.

During the last ten periods, the average costs for uninformed and informed drivers are 4.53 (sd=0.25) and 4.38 (sd=0.22) in Segment 1 and 4.48 (sd=0.30) and 4.37 (sd=0.23) in Segments 1 - 3. The two tests show that informed drivers pay a lower average cost than uninformed drivers at significance levels of 10% and 5% in Segment 1 and Segments 1 - 3, respectively. The average probability of choosing route a for uninformed and informed drivers are 0.68 (sd=0.13) and 0.65 (sd=0.14) in Segment 1 and are 0.66 (sd=0.16) and 0.67 (sd=0.15) in Segments 1 - 3. Both tests fail to reject the null hypothesis that the two types are equally likely to select a route using each data set.

To summarize, the uninformed and informed have a similar probability of selecting route a, regardless of the length of their experience. However, the informed pay significantly less in congestion costs than the uninformed in the late periods.

H The Endogenous Treatment

Under the Endogenous treatment, the realized price for the information depends on the outcome of the random price draw process introduced in Section 4.2. The numbers of groups that contain zero to six informed participants are 17, 12, 2, 3, 0, 2, and 12, respectively. Thus, the majority of groups endogenously produce No Information and Full Information conditions. The average cost and number of participants choosing route a for each group are plotted in Figure 15. The means of groups are displayed with gray diamonds. We can not draw broad conclusions due to the small samples. Therefore, we focus on conditions with at least 12 observations, i.e., conditions with the number of informed drivers equal to 0, 1, and 6. Based on the Wilcoxon signed-rank and the bootstrapped t-tests, the three conditions do not differ from each other significantly in terms of the average cost or the number of drivers choosing route a.

(a) Average cost (b) Number choosing a

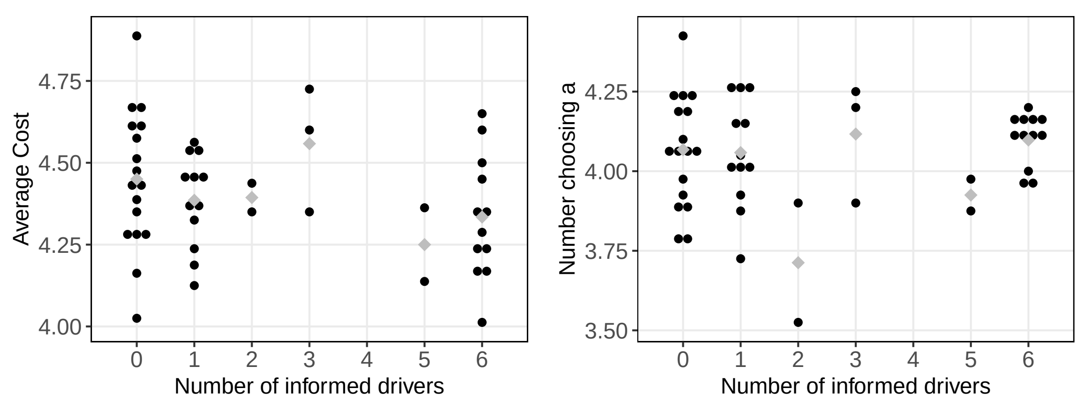

Figure 15: Plots of the average travel cost and the route choice of each of the 48 groups under the Endogenous treatment over all 40 periods. (a) plots average cost; (b) plots the number of players choosing route a. The gray diamonds represent the averages of groups with the same number of informed participants.

I Comparison Between The Two Locations

This section compares the data from the two locations. We will refer to the two locations as the Chinese and U.S. samples for simplicity.

I.1 Efficiency

To begin with, we analyze the average cost. Figure 16 presents information on the average cost across the three treatments, utilizing data from the Chinese and U.S. samples within Segment 1, Segments 1 and 2, and Segments 1 - 3, respectively. For a comprehensive overview, Table 16 summarizes the average cost for the first ten, last ten, and all 40 periods in each of the three treatments, for the various subsets of the data.

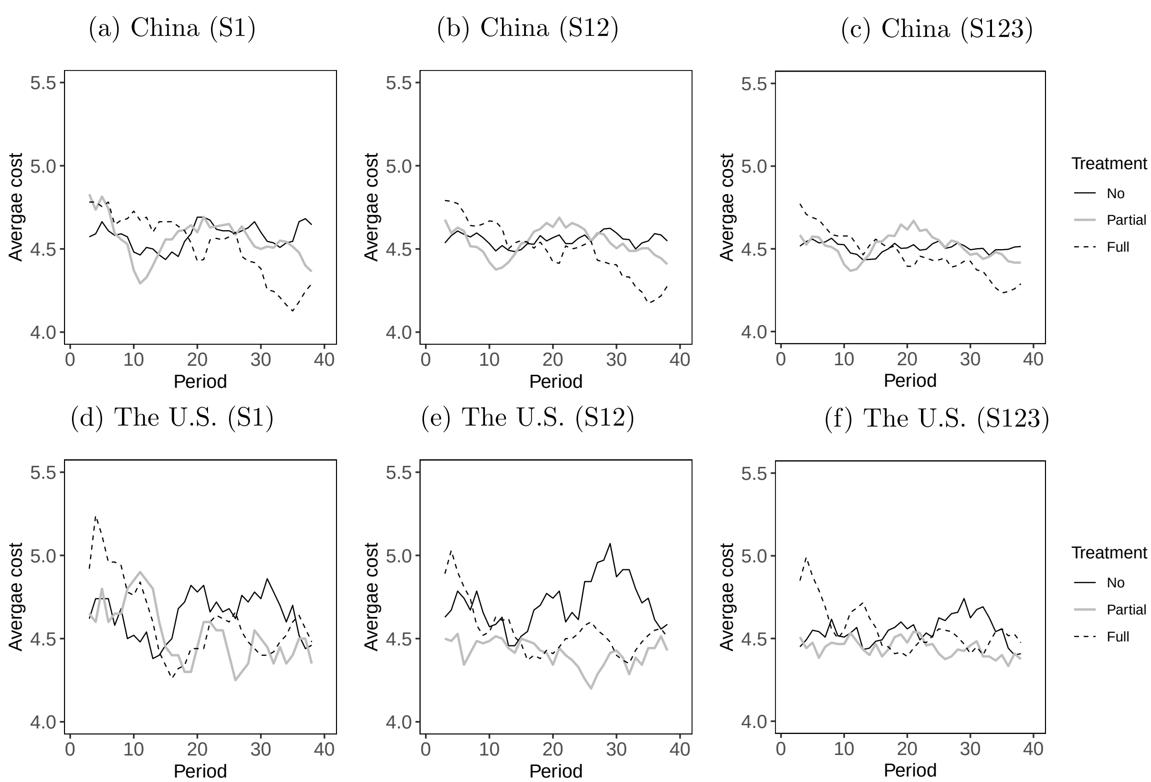

Figure 16: The realized average cost of each treatment over time. Panels (a) to (c) show the Chinese data from Segments 1, Segments 1 and 2, and Segments 1 - 3 of the experiment, respectively. In Panels (d) to (f), the U.S. data from Segments 1, Segments 1 and 2, and Segments 1 - 3 of the experiment are presented. The connected lines indicate five-period moving averages.

The combined insights derived from Figure 16 and Table 16 involve two patterns shared by the two sites. First, the `Full` condition consistently exhibits the lowest effi-

Table 16: Average cost for groups in China and the U.S.

Treatment	China: average cost				The U.S.: average cost			
	First10	Last10	All40	Obs.	First10	Last10	All40	Obs.
No (S1)	4.58	4.58	4.58	11	4.65	4.58	4.63	5
	(0.19)	(0.23)	(0.11)		(0.20)	(0.20)	(0.10)	
Partial (S1)	4.69	4.46	4.57	14	4.65	4.40	4.51	2
	(0.33)	(0.16)	(0.14)		(0.49)	(0.35)	(0.32)	
Full (S1)	4.73	4.25	4.52	11	4.74	4.47	4.55	5
	(0.34)	(0.21)	(0.26)		(0.42)	(0.20)	(0.14)	
No (S12)	4.57	4.52	4.55	25	4.71	4.69	4.70	7
	(0.20)	(0.22)	(0.13)		(0.19)	(0.25)	(0.17)	
Partial (S12)	4.59	4.45	4.55	25	4.49	4.41	4.42	7
	(0.30)	(0.18)	(0.16)		(0.28)	(0.22)	(0.22)	
Full (S12)	4.71	4.28	4.49	22	4.74	4.47	4.55	10
	(0.31)	(0.27)	(0.21)		(0.42)	(0.20)	(0.14)	
No (S123)	4.54	4.49	4.51	36	4.53	4.52	4.54	12
	(0.21)	(0.23)	(0.14)		(0.28)	(0.30)	(0.24)	
Partial (S123)	4.55	4.43	4.52	36	4.49	4.38	4.45	12
	(0.28)	(0.21)	(0.15)		(0.28)	(0.25)	(0.21)	
Full (S123)	4.68	4.30	4.47	36	4.70	4.48	4.56	12
	(0.33)	(0.26)	(0.19)		(0.39)	(0.25)	(0.16)	

Notes: The average cost for a Chinese or a U.S. participant for the first ten periods, for the last ten periods, and for all 40 periods in a given treatment. Results are presented across three subsets of the data: S1 (Segment 1), S12 (Segments 1 and 2), and S123 (Segments 1 - 3). Standard deviations are in parentheses. A standard deviation is calculated by averaging the cost of the members of each group in each period, averaging over the 10 or 40 periods for the group, and then computing the between-group standard deviation of these averages.

ciency in the earlier periods, whereas the No condition tends to be the least efficient in the later periods. Second, there is a consistent trend for the average cost to decrease under both the Partial and Full treatments, while no significant change is observed under the No treatment across periods.

The primary distinction between the two locations lies in the relative performance of the Partial treatment. In the U.S. samples, Partial consistently emerges as the most efficient condition, regardless of the timeframe considered. In the Chinese sample, Partial demonstrates an efficiency level that falls between No and Full during the early and late periods, never ranking as the best or worst.

To assess the aforementioned observations, we employ the Wilcoxon signed-rank and bootstrapped pairwise t-tests to test for within-subject data in Segments 1 and 2, as well as in Segments 1 - 3. For the between-subject data from Segment 1 only, we utilize the Mann-Whitney U and bootstrapped t-tests. Due to the relatively small sample size in the U.S., the tests on the U.S. sample are only conducted on the full data from Segments 1 - 3. For simplicity, we denote the statistics for Wilcoxon signed-rank, bootstrapped pairwise t, Mann-Whitney U, and bootstrapped t-tests as V, PT, U, and T, respectively. All tests are two-sided.

During the first ten periods, the tests suggest: i) No is significantly more efficient than Full in the Chinese sample (S12: V=4 with p-value=0.03, PT=-2.80 with p-value=0.01; S123: V=173.5 with p-value=0.06, PT=-2.15 with p-value=0.03). The comparison is of the same sign but insignificant in the U.S. sample. ii) Partial is significantly more efficient than Full using the S123 data from both samples (China: V=186.5 with p-value=0.06; the U.S.: V=16.6 with p-value=0.08, PT=-2.17 with p-value=0.04).

In the last ten periods, the tests establish that: i) Partial is not significantly more efficient than No using either the Chinese sample or the U.S. sample. ii) Full is significantly more efficient than No using the Chinese sample (S1: U=105.5 with p-value<0.01, T=3.52 with p-value<0.01; S12: V=56 with p-value=0.05; S123: V=470.5 with p-value<0.01, PT=3.98 with p-value<0.01), while the comparison is of the same sign but is insignificant in the U.S. sample. iii) Full is significantly more efficient than Partial using the Chinese sample (S1: U=124 with p-value=0.01, T=2.73 with p-value=0.04; S12: V=52.5 with p-value=0.09; S123: V=511.5 with p-value<0.01, PT=2.58 with p-value=0.07). Within the U.S. sample, however, Full is insignificantly less efficient than Partial.

The only significant difference observed over the entire 40 periods is between Partial and Full using the U.S. samples from Segments 1 - 3 (V=16 with p-value=0.08). The comparison is insignificant for the Chinese sample.

Regarding the efficiency dynamic within a specific treatment, which is the difference in average cost between the first ten and last 10 periods, the tests indicate: i) The average cost decreases significantly under Partial within samples from both China (S1: V=87 with p-value=0.03, PT=-2.56 with p-value=0.01; S12: V=249.5 with p-value=0.02, PT=-2.34 with p-value=0.02; S123: V=418.5 with p-value=0.04, PT=-2.22 with p-value=0.03) and the U.S. (S123: V=57 with p-value=0.04). ii) The average cost decreases significantly under Full utilizing both the Chinese sample (S1: V=63 with p-value<0.01, PT=3.70 with p-value<0.01; S12: V=217.5 with p-value<0.01, PT=4.99 with p-value<0.01; S123: V=528 with p-value<0.01, PT=5.83 with p-value<0.01) and the U.S. sample (S123: V=61 with p-value=0.09, PT=2.04 with p-value=0.04).

Overall, the data from the two locations display similar patterns. This is perhaps not surprising since we are unaware of any cultural differences in behavior in environments requiring coordination.

I.2 Route Selection

In this subsection, we focus on the route selection in the two locations. Table 17 displays route selection details for the three treatments for both samples using different data subsets. Notably, in the Chinese sample, there is a tendency for more drivers under No to choose route a than under the other two treatments at the beginning, regardless of the data subset analyzed. In the U.S. sample, however, there tend to be more drivers under No choosing route a than under the other two treatments in the late periods.

Table 17: Number of users choosing route a in the groups of China and the U.S.

Treatment	China: number choosing a				The U.S.: number choosing a			
	First10	Last10	All40	Obs.	First10	Last10	All40	Obs.
No (S1)	4.23	4.02	4.03	11	4.18	4.24	4.09	5
	(0.31)	(0.32)	(0.22)		(0.46)	(0.30)	(0.19)	
Partial (S1)	4.01	4.01	3.94	14	4.00	3.70	3.88	2
	(0.30)	(0.19)	(0.18)		(0.57)	(0.56)	(0.35)	
Full (S1)	4.09	4.05	4.02	11	3.98	3.80	3.90	5
	(0.23)	(0.14)	(0.09)		(0.31)	(0.07)	(0.14)	
No (S12)	4.12	4.08	4.02	25	3.99	4.24	4.03	7
	(0.33)	(0.27)	(0.20)		(0.50)	(0.40)	(0.20)	
Partial (S12)	3.98	4.02	4.97	25	4.07	3.87	3.95	7
	(0.27)	(0.24)	(0.17)		(0.34)	(0.32)	(0.17)	
Full (S12)	4.01	4.03	3.98	22	3.95	3.86	3.89	10
	(0.28)	(0.18)	(0.12)		(0.30)	(0.25)	(0.15)	
No (S123)	4.10	3.99	4.03	36	3.93	4.12	3.99	12
	(0.37)	(0.33)	(0.22)		(0.41)	(0.40)	(0.18)	
Partial (S123)	3.95	4.00	3.98	36	4.02	3.88	3.89	12
	(0.28)	(0.22)	(0.16)		(0.27)	(0.30)	(0.15)	
Full (S123)	3.99	4.06	4.00	36	3.98	3.87	3.91	12
	(0.27)	(0.24)	(0.12)		(0.27)	(0.23)	(0.15)	

Notes: This table shows the average number of participants choosing route a for a Chinese or a U.S. group for the first ten periods, the last ten periods, and all 40 periods in a given treatment. Results are presented across three subsets of the data: S1 (Segment 1), S12 (Segments 1 and 2), and S123 (Segments 1 - 3). Standard deviations are in parentheses. A standard deviation is calculated by counting the number choosing route a of each group in each period, averaging over the 10 or 40 periods for the group, and then computing the between-group standard deviation of these averages.

We employ the same tests as in Appendix I.1 to examine these patterns and find the following: i) Among the Chinese participants and during the first ten periods, significantly more drivers selected route $\mathtt{a}$ under $\mathtt{No}$ than under $\mathtt{Partial}$ (S1: $\mathtt{U}$=108.5 with $\mathtt{p}$-value=0.09; S123: $\mathtt{V}$=329 with $\mathtt{p}$-value=0.05, $\mathtt{PT}$=2.27 with $\mathtt{p}$-value=0.04). ii) Among the U.S. participants in the last ten periods, significantly more drivers selected route $\mathtt{a}$ under $\mathtt{No}$ than under $\mathtt{Partial}$ (S123: $\mathtt{PT}$=1.67 with $\mathtt{p}$-value=0.09) and under $\mathtt{Full}$ (S123: $\mathtt{V}$=61.5 with $\mathtt{p}$-value=0.08, $\mathtt{PT}$=1.91 with $\mathtt{p}$-value=0.08). Over the entire 40 periods, more drivers select route $\mathtt{a}$ under $\mathtt{No}$ than under $\mathtt{Partial}$ in the U.S. sample (S123: $\mathtt{V}$=63 with $\mathtt{p}$-value=0.06 and $\mathtt{PT}$=1.94 with $\mathtt{p}$-value=0.04). The findings in the two locations are not in disagreement. Furthermore, there is no significant change in the route selection rate over time, regardless of the time frame analyzed.

J Instructions and Quizzes for the Experiment

This Appendix contains the instructions that were displayed and read to participants during the experiment and the quiz questions in the questionnaire used at the end of the experiment. There were six versions of the instructions. Three or five groups of participants out of 24 groups received the version reprinted here. In particular, for this version, the sequence of treatments introduced in Parts 1 to 3 is $\mathtt{No}$, $\mathtt{Partial}$, and $\mathtt{Full}$.[30] The other three or five groups of participants received an identical version, except for the order in which some material in the instructions for Parts 1 to 3 appeared. Specifically, in the other five versions, the order of the treatments introduced in Parts 1 to 3 shown in the instructions was five shuffles of Parts 1 to 3 in this version.

Instructions for experiment

General Instructions

Welcome to an experiment on traffic route choice. During this experiment, all of you will make many decisions. If you follow the instructions carefully and make good decisions, you will earn a considerable amount of money that will be paid to you at the end of this experiment. Please raise your hand if you have any questions, then one of the supervisors will come to assist you. Note that communication with other participants and phone using are prohibited during the experiment.

The experiment consists of four parts, and each part includes 40 periods.

In the experiment, your payoff will be counted in points. At the end of the experiment, one part will be selected randomly by the computer to count toward your final payment.

[30] Parts 1 to 4 used in instructions correspond to Segments 1 to 4 in the main text.

Each part has an equal chance to be chosen. You will be paid in cash for your total earnings for all the 40 periods of the selected part of the experiment at an exchange rate of \$1=14 points. Furthermore, you will receive a show-up fee of \$10 for attendance.

Your earnings for a period may be negative only when you are incredibly unlucky. If your total earnings for the 40 periods of the selected part are negative, which is extremely unlikely, your total earnings will be rounded to zero, and you will get the \$10 show-up fee.

The instructions for part 1 will be given now, and those for parts 2 to 4 will be given later.

Instructions for part 1

In the role of drivers, the same six participants will be grouped for the duration of the experiment. In each period of this part, you will participate in the following traffic task.

Traffic task

In each period, you will be asked to choose either route a or b to go from the origin to the destination. The other five drivers in your group will do the same. No one knows which route other drivers have chosen when they make their own decision. Figure 17 below illustrates the traffic task.

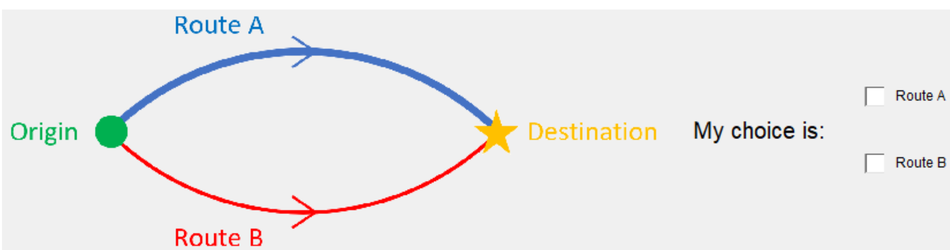

Figure 17: Traffic task

There are two routes, route a (the upper curve in blue) and route b (the lower curve in red), between the origin (the left circle in green) and the destination (the right star in yellow). Route a is wider than route b, so it is faster to take route a than route b for the same number of drivers on each route.

To decide for one of the routes, you can select it by clicking the small white rectangle in front of the route option at the right part of the figure.

For reaching the destination, you must pay a cost from your earnings, representing the delay and inconvenience caused by the traffic on your route. Specifically, if you decide on route a, the traffic cost equals the number of drivers on route a. If you choose route b, the cost equals **two times** the number of drivers on route b. For example, if two drivers choose route $\mathfrak{a}$ and four drivers choose route $\mathfrak{b}$, each driver on route $\mathfrak{a}$ pays a cost of 2 points, and each driver on route $\mathfrak{b}$ pays $2 \times 4 = 8$ points. Please note that the traffic cost is susceptible to congestion. It increases proportionally with the amount of traffic. And route $\mathfrak{b}$ is more sensitive than route $\mathfrak{a}$: for the same number of drivers on both routes, route $\mathfrak{b}$'s traffic cost is double that of route $\mathfrak{a}$.

Reaching the destination earns you an unknown amount of money, called the travel value. In particular, the travel value is an unknown number between 6 and 16, with one decimal place. Each number in this range has an equal chance to be your travel value. Regardless of your route choice, the travel value for a given period is fixed. In other words, the travel value for choosing route $\mathfrak{a}$ is the same as that for choosing route $\mathfrak{b}$. However, the travel value varies among drivers and periods. Specifically, your travel values in any two periods have no relationship with each other; any two drivers' travel values in the same period or any two periods are not related to each other.

Your earnings in points equal the difference between your travel value and the cost of the route you choose in each period:

$$\text{your earnings} = \text{your travel value - cost of your route}$$

where

$$\text{cost of your route} = \text{number of drivers on route } \mathfrak{a}, \text{if you choose route } \mathfrak{a}$$
$$\text{cost of your route} = 2 \times \text{number of drivers on route } \mathfrak{b}, \text{if you choose route } \mathfrak{b}$$

Thus, your earnings in each period are:

$$\text{your earnings} = \text{your travel value - number of drivers on route } \mathfrak{a}, \text{ if you choose route } \mathfrak{a}$$
$$\text{your earnings} = \text{your travel value - } 2 \times \text{number of drivers on route } \mathfrak{b}, \text{ if you choose route } \mathfrak{b}$$

Procedure

There is **one driver type** in this part of the experiment: **uninformed**. That is, all of you are uninformed of the traffic cost at the end of each period.

At the beginning of each period of this part, you will be asked to choose one of the two available routes $\mathfrak{a}$ and $\mathfrak{b}$. Once all participants in your group decide, you will receive a feedback screen presenting your current period's earnings, **but not** the travel value **or** the cost for any route.

You will perform this task for 40 periods. At any time, you can see a table on your screen indicating your route choice and earnings in all previous periods.

Payment

The chance of part 1 counting toward your final payment is one in four or 25%. If part 1 is randomly selected by the computer at the end of the experiment, your total earnings for the 40 periods in part 1 will be paid at an exchange rate of $1 = 14 points.

Practice part

There will be 20 practice periods now to help you become more comfortable with the traffic task. In the practice part (period -19 to period 0), each of you will be grouped with **five computerized bots** instead of five real participants that you will be grouped with in parts 1 to 4. Specifically, in every practice period, the five bots randomly choose their route, and their likelihood of choosing route $\mathfrak{a}$ is the same in a given period, but this likelihood is unknown to you. Furthermore, a bot's route choice (choose route $\mathfrak{a}$ or B) has no other relationship with the route choice of other bots. Your earnings in the practice part will not count toward your final payment. Part 1 will start after the practice part.

Instructions for part 2

For the duration of this part of the experiment, the same six participants as in part 1 will be grouped to act as drivers. In each period of this part, you will participate in the same traffic task introduced in part 1.

Procedure

There are **two driver types** in this part of the experiment: **informed** and **uninformed**. In particular, three drivers are informed of the traffic cost at the end of each period, and the other three drivers are uninformed. Whether you are informed or not is randomly assigned and remains unchanged throughout the 40 periods in this part of the session. You will learn your type at the beginning of this part.

At the beginning of each period of this part, you will be asked to choose one of the two available routes $\mathfrak{a}$ and $\mathfrak{b}$. Once all participants decide, each uninformed driver will receive a feedback screen presenting own current period's earnings, **but not** the travel value **or** the traffic cost for any route. However, each informed driver will receive a feedback screen presenting own current period's earnings, **and traffic cost information**, which includes cost of own route and the number of drivers on each route.

You will perform this task for 40 periods. At any time, you can see a table on your screen indicating your route choice and earnings in all previous periods. And if you are an informed driver, you can also observe traffic cost information including the cost of your route and the numbers of drivers on each route in all previous periods.

Payment

The chance of part 2 counting toward your final payment is one in four or 25%. If part 2 is randomly selected by the computer at the end of the experiment, your total earnings for the 40 periods in part 2 will be paid at an exchange rate of $1 = 14 points.

Instructions for part 3

For the duration of this part of the experiment, the same six participants as in parts 1 and 2 will be grouped to act as drivers. In each period of this part, you will participate in the same traffic task introduced in part 1.

Procedure

There is **one driver type** in this part of the experiment: **informed**. That is, you are all informed of the traffic cost at the end of each period.

At the beginning of each period of this part, you will be asked to choose one of the two available routes $\mathtt{a}$ and $\mathtt{b}$. Once all participants decide, you will receive a feedback screen presenting your current period's earnings, **and traffic cost information**, which includes cost of your route and the number of drivers on each route.

You will perform this task for 40 periods. At any time, you can see a table on your screen indicating your route choice and earnings in all previous periods. And if you are an informed driver, you can also observe traffic cost information including the cost of your route and the numbers of drivers on each route in all previous periods.

You will perform this task for 40 periods. At any time, you can see a table on your screen indicating your route choice, earnings, and traffic cost information including the cost of your route and the numbers of drivers on each route in all previous periods.

Payment

The chance of part 3 counting toward your final payment is one in four or 25%. If part 3 is randomly selected by the computer at the end of the experiment, your total earnings for the 40 periods in part 3 will be paid at an exchange rate of $1 = 14 points.

Instructions for part 4

For the duration of this part of the experiment, the same six participants as all previous parts will be grouped to act as drivers. In each period of this part, you will participate in the same traffic task introduced in part 1.

Procedure

There **can be two driver types** in this part of the experiment: **informed** and **uninformed**. In particular, informed drivers can observe the traffic cost at the end of each period, but uninformed drivers cannot. At the beginning of this part, you will go through a driver-type decision process in which you will decide whether to be informed or uninformed of the traffic cost at the end of each period. Your driver-type decision stays the same throughout the 40 periods of this part. The driver-type decision process will be introduced shortly. You will have a randomly assigned driver ID number in your group, a unique number in 1 to 6. The ID number keeps unchanged throughout this part of the session. You will learn your ID number and your type at the beginning of this part.

At the beginning of each period of this part, you will be asked to choose one of the two available routes $\mathfrak{a}$ and $\mathfrak{b}$. Once all participants decide, each uninformed driver will receive a feedback screen presenting own current period's earnings, **but not** the travel value **or** the traffic cost for any route. However, each informed driver will receive a feedback screen presenting own current period's earnings, **and traffic cost information**, which includes the cost of own route and the number of drivers on each route.

You will perform this task for 40 periods. At any time, you can see a table on your screen indicating your route choice and earnings in all previous periods. And if you are an informed driver, you can also observe traffic cost information including the cost of your route and the number of drivers on each route in all previous periods.

You will perform this task for 40 periods. At any time, you can see a table on your screen indicating your route choice, earnings, and traffic cost information including the cost of your route and the numbers of drivers on each route in all previous periods.

Payment

If you decide to be an informed driver so that you will be able to observe the traffic cost at the end of each period, you will receive (or pay) a certain **access bonus (or access fee)** per period. Specifically, the access bonus (or access fee) for traffic cost information is added to (or deducted from) the account of informed drivers. Your earnings in points are calculated as follows.

If an access bonus were provided per period for being informed, and you decide to be an **informed** driver, your earnings in each period will be:

$$\text{your earnings} = \text{your travel value - number of drivers on route } \mathfrak{a} + \text{access bonus, if you choose route } \mathfrak{a}$$

$$\text{your earnings} = \text{your travel value (minus) } 2 \times \text{number of drivers on route } \mathfrak{b} + \text{access bonus,if you choose route } \mathfrak{b}$$

If an access fee were charged per period for being informed, and you decide to be an **informed** driver, your earnings in each period will be:

$$\text{your earnings} = \text{your travel value - number of drivers on route } \mathfrak{a} \text{ - access fee, if you choose route } \mathfrak{a}$$

your earnings = your travel value - 2 × number of drivers on route $\mathtt{b}$ - access fee, if you choose route $\mathtt{b}$

If an access fee were charged per period for being informed, and you decide to be an **informed** driver, your earnings in each period will be:

your earnings = your travel value - number of drivers on route $\mathtt{a}$, if you choose route $\mathtt{a}$

your earnings = your travel value - 2 × number of drivers on route $\mathtt{b}$, if you choose route $\mathtt{b}$

If you decide to be an uninformed driver, you will not receive (or pay) any access bonus (or fee). Your earnings in each period will be:

The chance of part 4 counting toward your final payment is one in four or 25%. If part 4 is randomly selected by the computer at the end of the experiment, your total earnings for the 40 periods in part 4 will be paid at an exchange rate of \$1 = 14 points.

Driver type decision

Before you do the traffic task, we ask each of you to make two classes of decisions in a decision table to determine your type. The first class of decision is an unconditional decision column which includes 13 decisions. The second class of decision consists of six conditional decision columns, each containing 13 decisions under a specified condition. You will have to make both classes of decisions without knowing the others' decisions. Once determined, your type, informed or uninformed, cannot be changed during the 40 periods of this part.

Stage 1: Decision table

The following decision table will be displayed on your computer screen. There are 13 offers from row 1 to row 13 which provide you an access bonus or charge you an access fee for being informed. Please indicate in columns (3) to (9) whether you are willing to accept each offer to be informed. The decision table will look like Figure 18:

Column (1) shows the decision row number from row 1 to row 13, and each corresponds to an offer.

Column (2) presents the access bonus (or the access fee) for the traffic cost information for the offer in each decision row: the per period bonus (or fee) for being informed. A bonus means you will receive some points if you choose to be informed; a fee means you will need to pay some points if you choose to be informed. For example, the offer in decision row 3 will provide a driver who accepts the offer an access bonus of 0.50 points per period. The offer in decision row 6 will charge a driver who accept the offer an access fee of 0.25 points per period.

(1)	(2)	(3) *unconditional*	(4) [conditional]	(5) [conditional]	(6) [conditional]	(7) [conditional]	(8) [conditional]	(9) [conditional]
Decision row	Access bonus/fee	Will you accept the offer?	If **no** others accept the offer, will you accept it?	If **one** other accepts the offert, will you accept it?	If **two** others accept the offer, will you accept it?	If **three** others accept the offer, will you accept it?	If **four** others accept the offer, will you accept it?	If **five** others accept the offer, will you accept it?
1	Bonus: you receive 1.00 point per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
2	Bonus: you receive 0.75 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
3	Bonus: you receive 0.50 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
4	Bonus: you receive 0.25 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
5	You receive/pay 0 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
6	Fee: you pay 0.25 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
7	Fee: you pay 0.50 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
8	Fee: you pay 0.75 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
9	Fee: you pay 1.00 point per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
10	Fee: you pay 1.25 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
11	Fee: you pay 1.50 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
12	Fee: you pay 1.75 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No
13	Fee: you pay 2.00 points per period	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No	Yes ◯◯ No

Figure 18: Type decision

Column (3) is the **unconditional** decision column which asks whether you will accept the offer, regardless of other group members' decisions.

Columns (4) to (9) are six **conditional** decision columns that ask whether you will accept the offer, dependent on other group members' decisions.

You will be asked to select between 'Yes' and 'No' in columns (3) to (9) for each decision row. For example, row 6 column (3) asks: if the offer is that you need to pay 0.25 points per period for being informed, are you willing to accept the offer? And row 6 column (6) asks: if the offer is that you need to pay 0.25 points for being informed, and two other drivers in your group accept the offer in order to be informed, are you willing to accept it?

Stage 2: Type revealing

Once all participants in your group make all decisions, the computer will randomly generate an 'R' number, a number in 1 to 13 with an equal chance, to determine which row of the table will count. Then, the computer will randomly generate a 'D' number, a number in 1 to 6 with an equal chance, to indicate the selected driver ID in your group. The five unselected drivers' type will depend on how they filled in column (3) in the 'R' row. If they chose 'Yes', they are informed drivers; If they chose 'No', they are uninformed drivers. The selected 'D' driver's type is determined by how he/she answered 'Yes' or 'No' in the 'R' row and the column indicating how many of the five unselected drivers chose 'Yes' in column (3) of the 'R' row.

For example, suppose that the 'R' number that determines the row that counts is 8, and the 'D' number indicating the selected driver ID is 4. Then, drivers 1, 2, 3, 5, and 6's decision in row 8 column (3) will count. Driver 4's type will be determined as follows. We look at how many of the other drivers chose 'Yes' in row 8 column (3). That number determines the column that counts for driver 4. For example, if drivers 1, 2, and 3 answered 'Yes' and drivers 5 and 6 answered 'No' in row 8 column (3), driver 4's decision in row 8 column (7) will count. In addition, informed drivers will pay an access fee of 0.75 points per period, and they will be informed of the traffic cost at the end of each period in this part.

Finally, you will learn your type and the numbers of informed and uninformed drivers.

Quiz before the practice type decision process

This is a quiz to test your understanding of the instructions. Please answer them carefully.

Question 1A: Assuming line 2 (price = -0.75) is implemented and you are not the selected driver, you select "**Yes**" in row 2 column (3) (unconditional column). Then your driver type is:

- Informed
- Uninformed

Question 1B: Assuming line 2 (price = -0.75) is implemented and you are not the selected driver, you select "**Yes**" in row 2 column (3) (unconditional column). How many points do you pay or do you get in each period for the information?

- You will pay 0.75 points in each period for the information.
- You will get 0.75 points in each period for the information.
- You will pay zero and get zero in each period for the information.

Question 1C: Assuming line 2 (price = -0.75) is implemented and you are not the selected driver, you select "**Yes**" in row 2 column (3) (unconditional column). In addition, if 3 drivers, including you, choose Route A in a specific period and your trip value is 10. Then your payoff for that period is: _____

Question 2A: Assuming line 2 (price = -0.75) is implemented and you are not the selected driver, you select "**No**" in row 2 column (3) (unconditional column). Then your driver type is:

- Informed
- Uninformed

Question 2B: Assuming line 2 (price = -0.75) is implemented and you are not the selected driver, you select "**No**" in row 2 column (3) (unconditional column). How many points do you pay or do you get in each period for the information?

- You will pay 0.75 points in each period for the information.
- You will get 0.75 points in each period for the information.
- You will pay zero and get zero in each period for the information.

Question 2C: Assuming line 2 (price = -0.75) is implemented and you are not the selected driver, you select "**No**" in row 2 column (3) (unconditional column). In addition, if 2 drivers, including you, choose Route A in a specific period and your trip value is 10. Then your payoff for that period is: _____

Question 3A: Assuming line 8 (price = 0.75) is implemented and you are not the selected driver, you select "**Yes**" in row 8 column (3) (unconditional column). Then your driver type is:

- Informed
- Uninformed

Question 3B: Assuming line 8 (price = 0.75) is implemented and you are not the selected driver, you select "**Yes**" in row 8 column (3) (unconditional column). How many points do you pay or do you get in each period for the information?

- You will pay 0.75 points in each period for the information.
- You will get 0.75 points in each period for the information.
- You will pay zero and get zero in each period for the information.

Question 3C: Assuming line 8 (price = 0.75) is implemented and you are not the selected driver, you select "**Yes**" in row 8 column (3) (unconditional column). In addition, if 4 drivers, including you, choose Route A in a specific period and your trip value is 10. Then your payoff for that period is: _____

Question 4A: Assuming line 7 (price = 0.5) is implemented and you **are** the selected driver. One other driver selected "**Yes**" and four others selected "**No**" in their row 7 column (3). You select "**No**" in row 7 column (3) (unconditional column) but "**Yes**" in row 7 column (5) (the conditional column corresponding to the number of other drivers choosing "**Yes**": one driver). Then your driver type is:

- Informed
- Uninformed

Question 4B: Assuming row 7 (price=0.5) is implemented and you **are** the selected driver. One other driver selected “**Yes**” and four others selected “**No**” in their row 7 column (3). You select “**No**” in row 7 column (3) (unconditional column) but “**Yes**” in row 7 column (5) (the conditional column corresponding to the number of other drivers choosing “**Yes**”: one driver). How many points do you pay or do you get if you receive the information?

- You will pay 0.5 points in each period for the information.
- You will get 0.5 points in each period for the information.
- You will pay zero and get zero in each period for the information.

Question 4C: Assuming row 7 (price=0.5) is implemented and you **are** the selected driver. One other driver selected “**Yes**” and four others selected “**No**” in their row 7 column (3). You select “**No**” in row 7 column (3) (unconditional column) but “**Yes**” in row 7 column (5) (the conditional column corresponding to the number of other drivers choosing “**Yes**”: one driver). In addition, if 5 drivers, including you, choose Route A in a specific period and your trip value is 10. Then your payoff for that period is: _____

Practice type decision process

There will be a practice type decision process now to help you become more comfortable with the type decisions. In the practice process, each of you will be grouped with **five computerized bots** instead of five real participants. Specifically, each bot randomly decides between ‘Yes’ or ‘No’ in the decision table. Furthermore, a bot’s decision has no relationship with the decisions of other bots. Your decisions in the practice decision process will not count.

Quiz in the questionnaire

Bonus question

What is the lowest cost (in points) that the six drivers of a group can have on average in one given period?

Your answer: _____

The bonus in points you get in this question:

$$\text{Your bonus}=15 - 2\times(\text{your answer - the right answer})^2 \text{ if } 15 - 2\times(\text{your answer - the right answer})^2 > 0;$$

$$\text{Your bonus}=0 \text{ if } 15 - 2\times(\text{your answer - the right answer})^2 \leq 0.$$

In other words, the smaller the gap between your answer and the correct answer, the more bonus you can get. Moreover, the bonus is non-negative, so you will not lose money by answering this question.

Non-bonus question

Which of the following two options do you think is better?

A. Try to change the route you chose frequently and keep people guessing/uncertain about your route.

B. Try to find a good assignment of drivers and routes and stick with the same route in different periods.