Online Appendices to Congestion Information and Efficiency: An Experiment

Charles N. Noussair* Liang Qiao†

June 13, 2025

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 Partial is presented in Appendix G. Appendix H provides a more detailed analysis of the 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

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

†School of Economics & The Laboratory for Economic Behavior and Policy Simulation, Nankai University, Tianjin, China. liangqiao@nankai.edu.cn

after period t for player i and is denoted as $N_i(t)$.

$$N_i(t) = \phi_i(t) \cdot N_i(t-1) + 1, \quad (1)$$

[1]

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

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

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 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 $S_{-i} = \{0, 1, 2, 3, 4, 5\}$.[2]

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 $A^j(0) = 1, \forall j \in \{0,1\}$.[3]

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

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

[2] 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.

[3] 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.

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 uninformed 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.[4], 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

[4] 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.

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

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

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 1 to present an outline of the EWA-lite learning algorithm we have used in our analysis.

Table 1: 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, $UInf$ and 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 $t = t + 1$

A.2 The Inference Made by the Uninformed

We summarize the inference by the uninformed players in Tables 2 and 3. In particular, Table 2 shows the belief transition matrices between realized payoffs and the resulting beliefs on others' actions for uninformed players when their choice is route a ($s_i = 1$). Table 3 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).

Table 2: 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 3: 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.

A.3 Simulation results over 40 Periods

This section includes two parts. Section A.3.1 compares the No, Partial, and Full treatments with respect to efficiency and route choice under various response sensitivities.[5] 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 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 1. Panel 1a shows the average traffic cost and panels 1b and 1c illustrate the average and the standard deviation of the number of players who select route a.

Figure 1 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 2 illustrate the average cost over the first ten periods, periods 31 to 40, and all 40 periods, respectively.

Figure 2 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 1b and 1c help explain the intuition behind the above findings. First, for

[5] The contrast between the outcomes of informed and uninformed players under the Partial treatment is discussed in Appendix G.

(a) Average cost

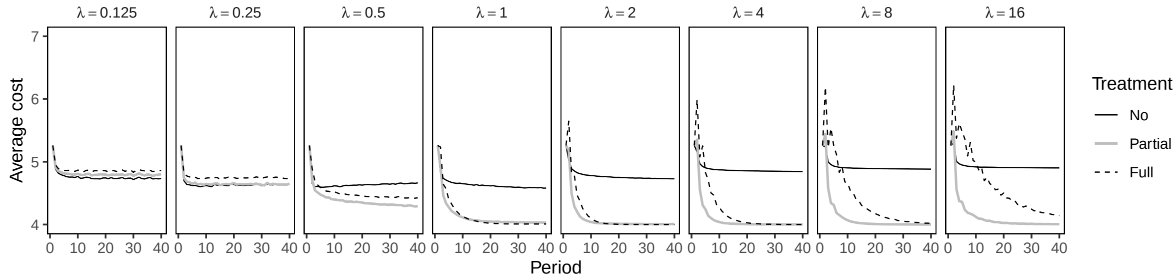

(b) Average number of players choosing route a

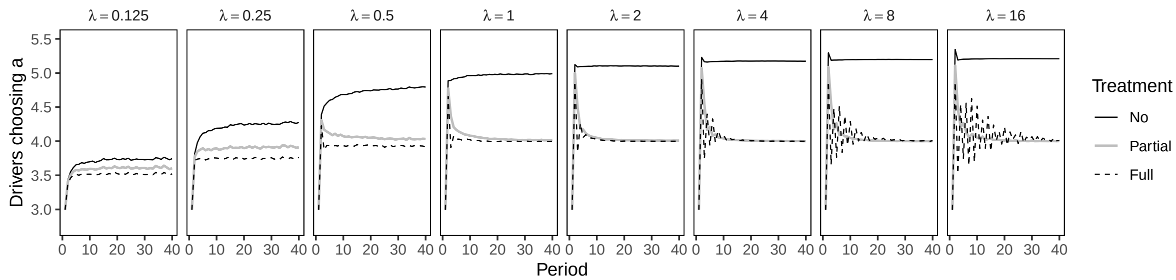

(c) Standard deviation of the number of players choosing route a

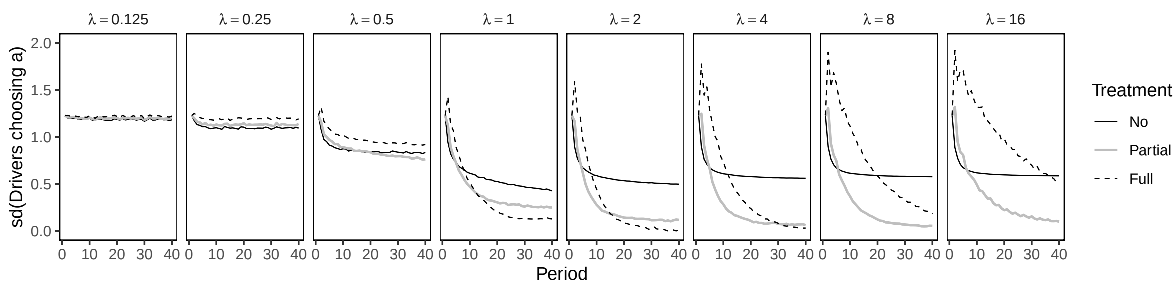

Figure 1: Dynamics of cost and traffic patterns over 40 periods. 10,000 simulations for a six-player group are conducted under each of the No, Partial, and Full treatments for $\lambda \in \{0.125, 0.25, 0.5, 1, 2, 4, 8, 16\}$. Panel (a) plots the average cost per individual. Panel (b) shows the average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

$\lambda = 0.125$, 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

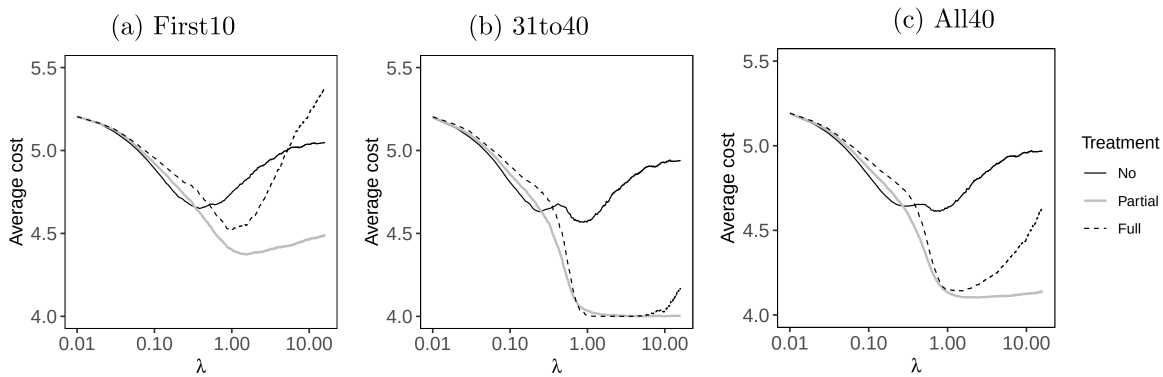

Figure 2: The average cost in 1,000 simulations for each of the three treatments when λ is between 0.01 and 16 with a step size of 0.01. Panels (a) to (c) show the average cost for the first ten periods, periods 31 to 40, and all 40 periods, respectively. The horizontal axis is on a logarithmic scale.

is to say, on average, more than four players select route $\mathtt{a}$. This bias toward route $\mathtt{a}$ is caused by the imperfect inference of uninformed players about the route choices of other players. As Tables 2 and 3 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 1b 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 1c, 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 2, 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 3 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 3 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

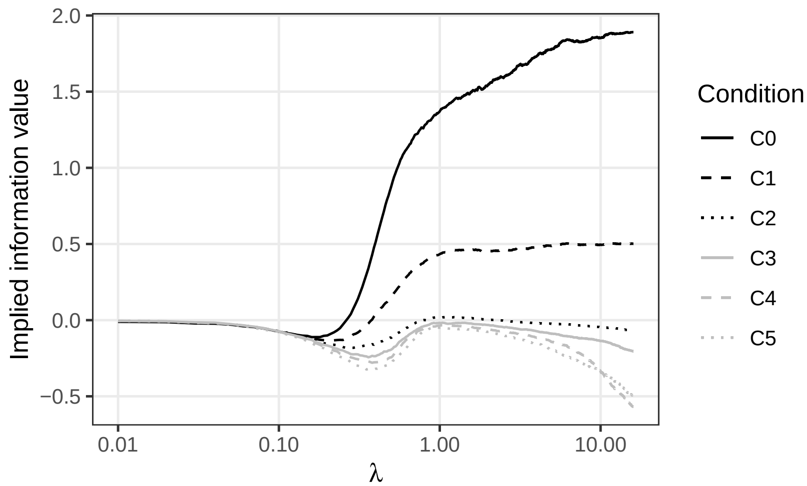

Figure 3: The implied value of information per period during 40 periods. C0 to C5 represent six conditions where the number of informed others is between 0 and 5. It shows the implied information value in 1,000 simulations under each condition for λs between 0.01 to 16 with a step size of 0.01. The horizontal axis is on a logarithmic scale.

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, $\mathsf{Partial}$, 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, $\mathsf{Partial}$, 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 6 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 4 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.

Table 4: Reliance on small samples learning algorithm summary

Given q, the single parameter of the model "sample of q"
{*Given* q*, the single parameter of the model "sample of q or less"*}
Assign types ($\mathtt{UInf}$ and $\mathtt{Inf}$) to all i
{*Assign types (*$\mathtt{UInf}$ *and* $\mathtt{Inf}$*) and sample sizes (*$q_i \in \{1, ..., q\}$ *with an equal chance) to all* i}
For $t = 1$ do
For all i do
Sample the chosen action $s_i(1) \in \{0, 1\}$ with an equal chance
Get $\sum_{k=1}^{6} s_k(1)$
For all i do
Independently randomly draw value, $\nu_i(1)$
Assess profit $\pi_i(s_i(1), s_{-i}(1))$ and $\hat{\pi}_i(s_i^j, s_{-i}(1)), j \in \{0, 1\}$
For $t > 1$ to length T do
For all i with the type $\mathtt{Inf}$ do
Sample with replacement of q trials from all past trials {1,...,t-1} with an equal chance
{*Sample with replacement of* q_i *trials from all past trials {1,...,t-1} with an equal chance*}
Select the action $s_i(t)$ with the highest average payoff in that small sample
For all i with the $\mathtt{UnInf}$ type do
If only one action has been taken beforehand
Sample the chosen action $s_i(t)$ with an equal chance
If both pure strategies have been taken beforehand
For all action j
Sample with replacement of q trials from past trials in which j has been chosen
{*Sample with replacement of* q_i *trials from past trials in which* j *has been chosen*}
Select the action $s_i(t)$ with the highest average payoff in its small sample
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 $\hat{\pi}_i(s_i^j, s_{-i}(t))$
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 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.

B.2.1 Efficiency and Route Choice

The sole free parameter, the size of a small sample denoted as 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 $\mathsf{q} \in \{1, 2, ..., 30\}$, under the No, $\mathsf{Partial}$, and Full treatments. The six panels in Figure 4 illustrate the average cost and the number of drivers selecting route a over the first ten periods, periods 31 to 40, and the entire 40-period duration, respectively.

Panels (a) to (c) in Figure 4 reveal some prominent patterns. First, in the early periods, Full information adoption is the least efficient, while the $\mathsf{Partial}$ condition is the most efficient except when $\mathsf{q} = 1$ 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 $\mathsf{q} \geq 13$. 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, $\mathsf{Partial}$, 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 $\mathsf{q} \geq 3$, 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. $\mathsf{Partial}$ 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 a for each condition is close to the socially optimal number of four.

B.2.2 The Value of Information

Figure 5 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

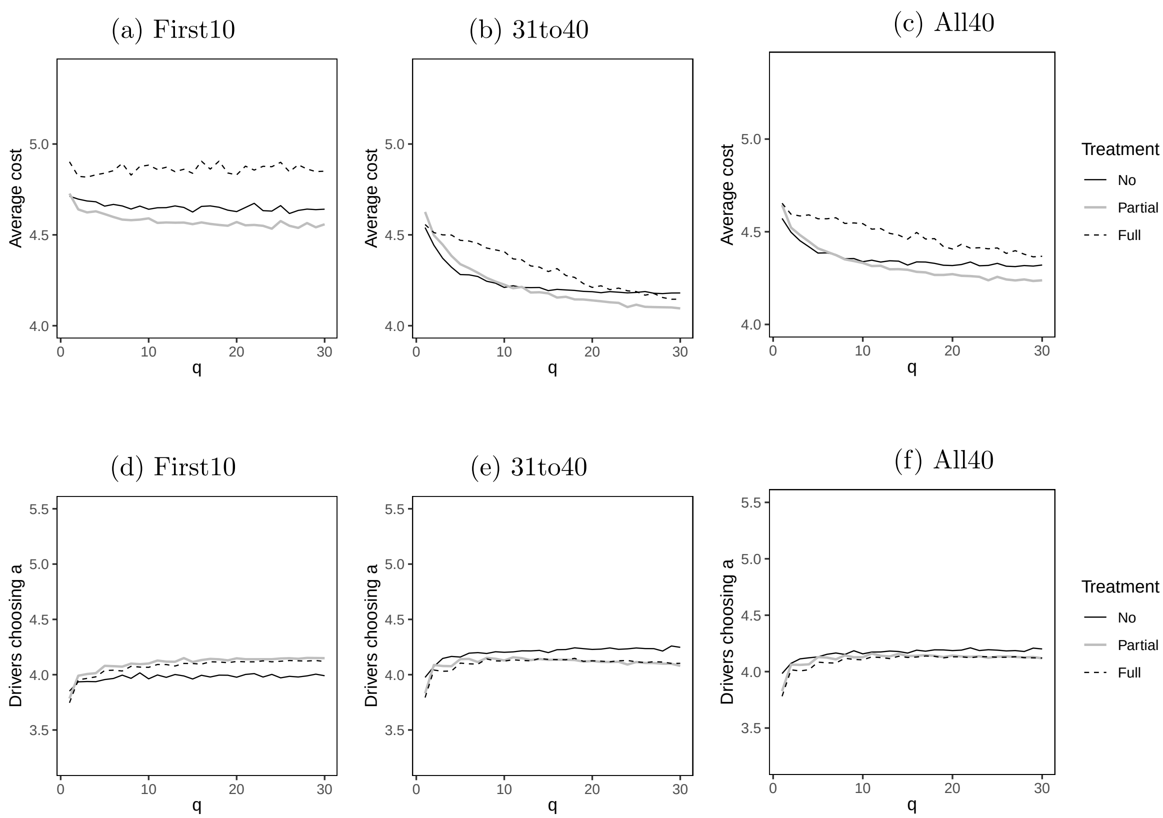

Figure 4: The average cost and the route choice in 1,000 simulations for each of the three treatments by the Sample of q model when q is between 1 and 30 with a step size of 1. Panels (a) to (c) show the average cost for the first ten periods, periods 31 to 40, and all 40 periods, respectively. Panels (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.

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 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 $q \geq 2$. The cost savings do not exceed 0.2 points per period across the entire range of q. Under condition C1, when one other player is informed and four are uninformed, becoming informed is profitable for a player when $q > 15$. However, in

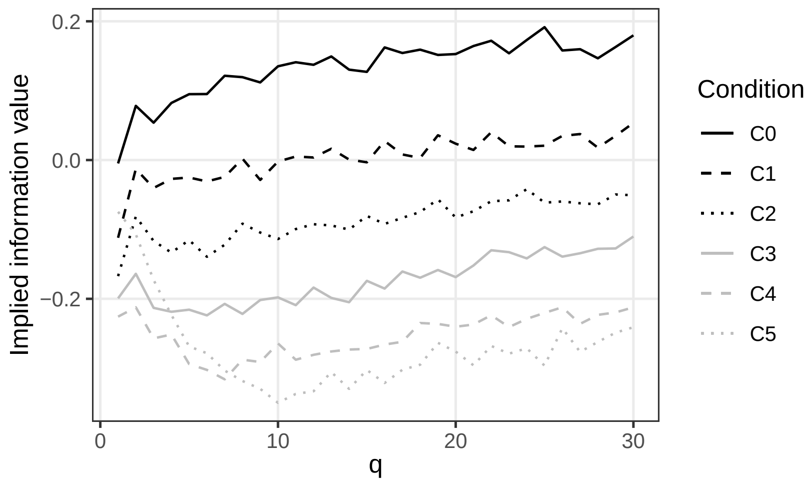

Figure 5: The implied value of information per period during 40 periods for the Sample of q model. C0 to C5 represent six conditions where the number of informed others varies from 0 to 5. It shows the implied information value in 1,000 simulations under each condition when q varies from 1 to 30 with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

C2 to C5, the implied information value is always negative.

C The Observed Transition Probability Matrix

Table 5 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 5: Transition probability matrix for relative costs of the two routes

	First 10 Periods			Last 10 Periods			All 39 Periods		
(1) Segment 1									
a. No	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$
$C^a_t < C^b_t$	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									
a. No	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$	$C^a_{t+1} < C^b_{t+1}$	$C^a_{t+1} = C^b_{t+1}$	$C^a_{t+1} > C^b_{t+1}$
$C^a_t < C^b_t$	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).[6] 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.[7]

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 6. Panels (a) and (b) use data from Segment 1 and from Segments 1 - 3, respectively.

[6] 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.

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

Table 6: Model Fitting for the EWA-lite and SSM Models and the Mean Normalized Squared Distance Scores

Model	Treatment	Fitted Parameter	Normalized Squared Distance scores: Average cost, Block 1			Average cost, Block 2			Number of drivers choosing a, Block 1			Number of drivers choosing a, Block 2			Mean
			N	P	F	N	P	F	N	P	F	N	P	F	
(a) Segment 1															
EWA-lite	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
	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
SSMq	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
	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
SSM1-q	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
	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															
EWA-lite	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
	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
SSMq	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
	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
SSM1-q	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
	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.[8] 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 6, 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 6 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

[8] 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).

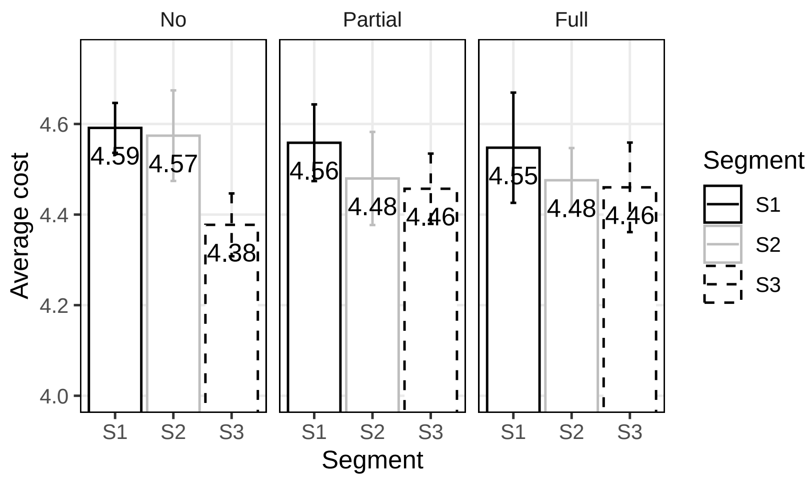

Figure 6: Carryover effects on the average cost over 40 periods. The average cost for each of the No, Partial, and Full treatments in Segments 1 - 3 (S1 to S3) and the corresponding 95% confidence intervals are presented.

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 7 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 7 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 7: 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 in the main text. 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 a across the three treatments are not significantly different. The mean standard deviation of the number of drivers on route a (MSDNA) for No, $\mathsf{Partial}$, and 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 Full. In the later periods, the value is highest under No but lowest under 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 No, $\mathsf{Partial}$, and Full, respectively. Both the bootstrapped pairwise t and the Wilcoxon signed-rank test statistics suggest that the average cost under $\mathsf{Partial}$ and 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 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 No, $\mathsf{Partial}$, and 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 $\mathsf{Partial}$ 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 7 illustrates the differences between the informed and uninformed with regard to congestion cost and the selection rate of route a, as predicted by simulations from the EWA-lite and Sample of 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 q for Sample of 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 $\mathsf{Partial}$ 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

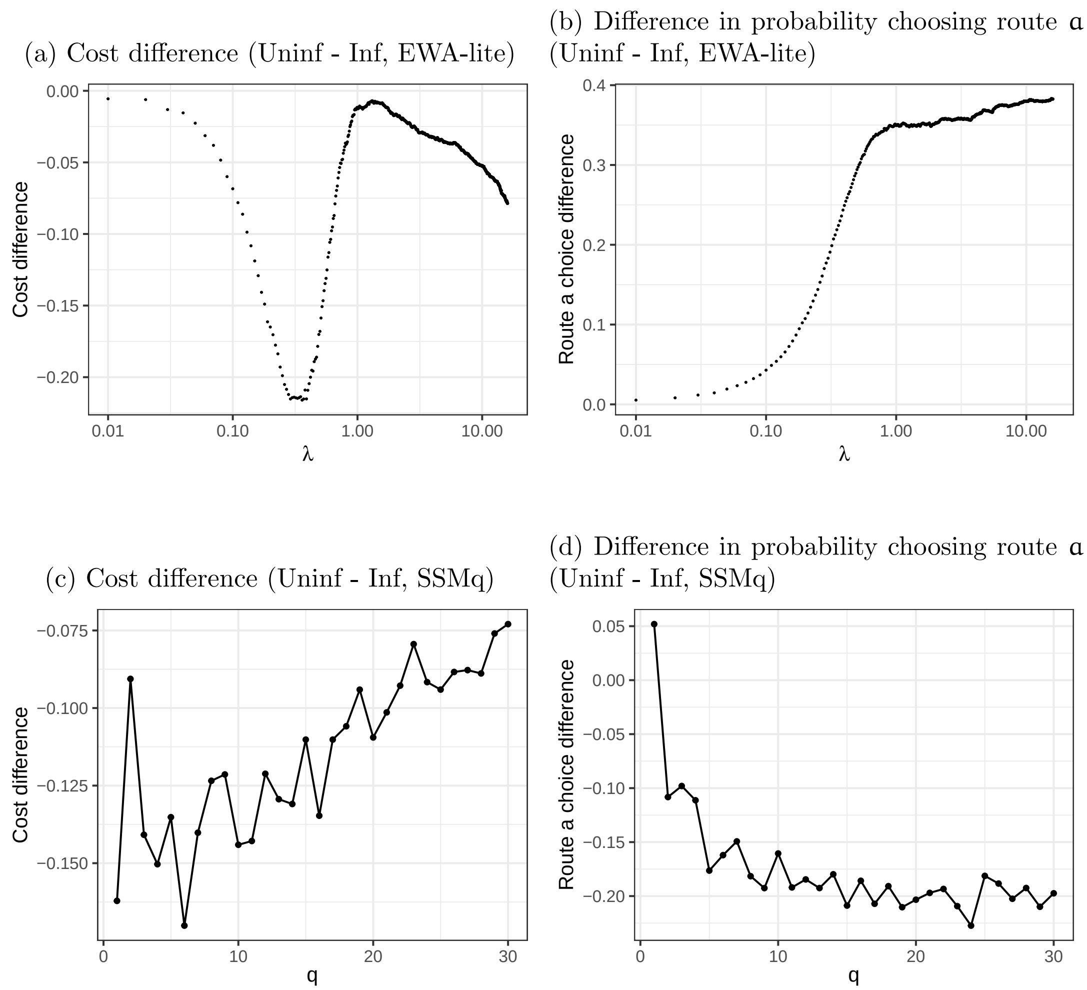

Figure 7: The differences between the uninformed and the informed under the Partial treatment as predicted by the EWA-lite and Sample of q models are illustrated. For each λ, the parameter for EWA-lite, ranging from 0.01 to 16 with a step size of 0.01, and each q, the parameter for Sample of q (SSMq), 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 $\mathtt{a}$ choice difference consistently maintains a strictly positive value, implying that the uninformed players are more likely to select route $\mathtt{a}$ than the informed players, irrespective of the value of λ. However, the Sample of $\mathtt{q}$ model predicts that the route $\mathtt{a}$ choice difference is strictly negative when $\mathtt{q} \geq 2$.

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 $\mathtt{t}$-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 $\mathtt{t}$-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 ($\mathtt{p}$-value=0.14), failing to reject the null hypothesis that the two types bear an equal average cost.

The average probability of choosing route $\mathtt{a}$ 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 $\mathtt{a}$ 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 $\mathtt{a}$ 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 $\mathtt{t}$-tests fail to reject the null hypothesis of equality between the two types regarding average cost and route $\mathtt{a}$ 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 $\mathtt{a}$ 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 in the main text. 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 8. 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.

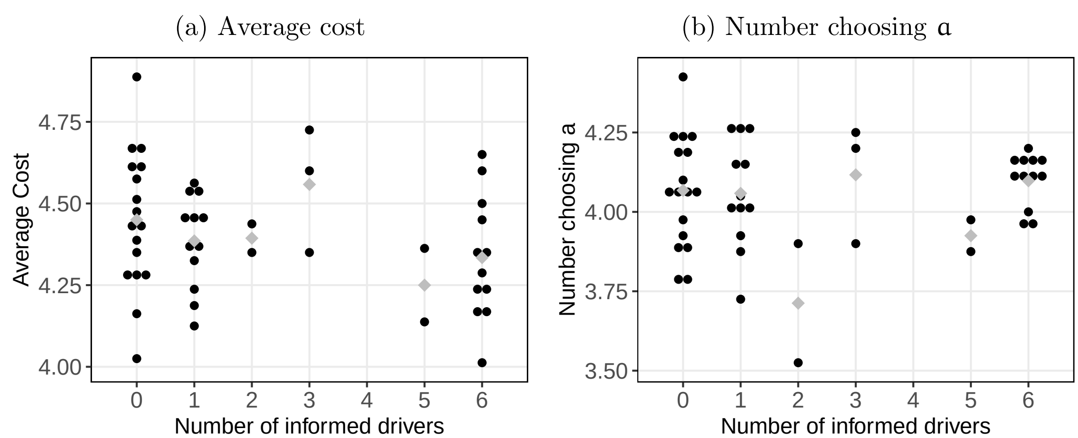

Figure 8: Plots of the average travel cost and the route choice of each of the 48 groups under the Endogenous treatment over all 40 periods. (a) plots average cost; (b) plots the number of players choosing route a. The gray diamonds represent the averages of groups with the same number of informed participants.

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 9 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 8 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.

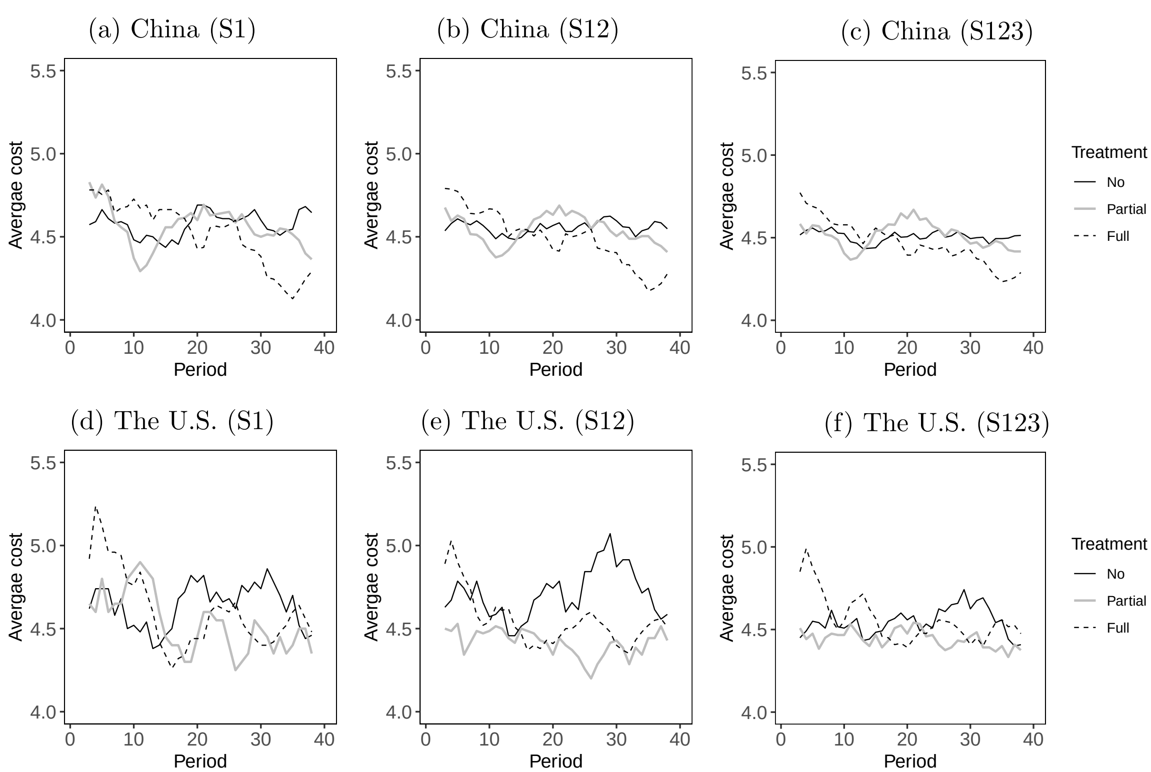

Figure 9: The realized average cost of each treatment over time. Panels (a) to (c) show the Chinese data from Segments 1, Segments 1 and 2, and Segments 1 - 3 of the experiment, respectively. In Panels (d) to (f), the U.S. data from Segments 1, Segments 1 and 2, and Segments 1 - 3 of the experiment are presented. The connected lines indicate five-period moving averages.

The combined insights derived from Figure 9 and Table 8 involve two patterns shared by the two sites. First, the Full condition consistently exhibits the lowest efficiency in

Table 8: 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.

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 9 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 9: 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 a under No than under Partial (S1: U=108.5 with p-value=0.09; S123: V=329 with p-value=0.05, PT=2.27 with p-value=0.04). ii) Among the U.S. participants in the last ten periods, significantly more drivers selected route a under No than under Partial (S123: PT=1.67 with p-value=0.09) and under Full (S123: V=61.5 with p-value=0.08, PT=1.91 with p-value=0.08). Over the entire 40 periods, more drivers select route a under No than under Partial in the U.S. sample (S123: V=63 with p-value=0.06 and PT=1.94 with 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 No, Partial, and Full.[9] 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.

[9] 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 10 below illustrates the traffic task.

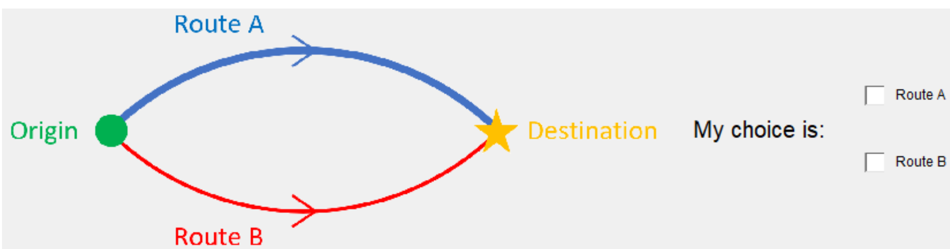

Figure 10: Traffic task

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:

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

your earnings = your travel value (minus) 2 × number of drivers on route $\mathfrak{b}$ + 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:

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

your earnings = your travel value - 2 × number of drivers on route b - access fee, if you choose route 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 a, if you choose route a

your earnings = your travel value - 2 × number of drivers on route b, if you choose route 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 11:

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 11: 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?

<u>Stage 2: Type revealing</u>

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.

References

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

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.

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. E. Roth (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.

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.

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