Experimental Economics (2025), **28**, 600–624
doi:10.1017/eec.2025.10025

economic science association

CAMBRIDGE UNIVERSITY PRESS

SPECIAL ISSUE ARTICLE

The Downs-Thomson Paradox with Endogenously Determined Departure Times

Hironori Otsubo[1], Eyran J. Gisches[2] and Amnon Rapoport[3]

[1]Faculty of Global Management, Chuo University, Hachioji, Tokyo, Japan
[2]Department of Management Information System, Eller College of Management, University of Arizona, Tucson, AZ, USA
[3]Department of Management and Organization, Eller College of Management, University of Arizona, Tucson, AZ, USA
Corresponding author: Hironori Otsubo; Email: otsubo.76t@g.chuo-u.ac.jp

(Received 30 November 2023; revised 25 June 2025; accepted 13 July 2025)

Abstract

We introduce a novel scenario that embeds the Downs-Thomson paradox in the context of departure-time choice during the morning commute. Commuters, departing from a common origin and traveling to a common destination, must choose between a congestible mode (car, road) and a non-congestible mode (train, railway). Those choosing the road must also select their departure times independently and anonymously. This decision involves a trade-off between the cost of queuing at the bottleneck and the cost of schedule delay (i.e., deviation from the desired arrival time). We numerically derive a symmetric mixed-strategy equilibrium that characterizes both mode and departure-time choices. We then examine how improvements to either the road or the railway affect mean travel costs. Our laboratory experiment shows that, consistent with the paradox, improving the railway lowers mean travel cost; however, contrary to the paradox, improving the road also reduces mean travel cost. These findings suggest that the Downs-Thomson paradox may fail to emerge fully when commuters must coordinate multiple strategic dimensions under intertemporal congestion externalities.

Keywords: Downs-Thomson paradox; queuing; endogenous departure-time choice; bottleneck; experiment

JEL Codes: C72; C92

1. Introduction

In a seminal paper, Arnott and Small (1994) noted that traffic congestion in metropolitan areas has become a major plague of modern society. Thirty-one years later, the populations of large cities around the world have grown significantly, the number of cars on the road has proliferated, delay times have increased markedly, and despite billions of dollars, yen, and euros spent expanding the capacity of road networks and constructing new roads, the negative effects of congestion have worsened. There seems to be no quick fix for traffic congestion. The magnitude of the congestion problem in large metropolitan areas is so enormous that no single policy can address it. Rather, the problem requires multiple strategies, where each strategy represents a component of a broader solution.

The most obvious strategies include expanding the capacity of existing roads; constructing new roads, railways, and bridges; assigning priority by designating special lanes for public transport and carpooling; and imposing driving restrictions on selected days. Although these strategies may seem intuitive, under certain conditions, some can yield counterintuitive results. Several paradoxical

© The Author(s), 2025. Published by Cambridge University Press on behalf of the Economic Science Association. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

outcomes have been discussed in the literature on transportation science, operations research, and computer science. These studies show that, under specific conditions related to network architecture, the number of users, and the costs of alternative modes of transportation, attempts to reduce congestion by expanding traffic networks may prove ineffective. Specifically, the *Pigou-Knight-Downs* (P-K-D) paradox (Pigou, 1920, Knight, 1924, Downs, 1962) states that, under certain conditions, expanding road capacity may not achieve the goal of reducing travel cost. The *Braess Paradox* (Braess, 1968) indicates that expanding a traffic network by adding one or more routes may paradoxically increase the equilibrium cost of commuting. The *Downs-Thomson* (D-T) paradox (Downs, 1962, Thomson, 1978), which is the focus of our study, posits that under certain conditions, increasing road capacity when commuters face a choice between private and public transportation modes may prove ineffective as it may *not* decrease the overall cost of travel.

The three paradoxes above have been presented, discussed, and illustrated numerically in the theoretical literature (see, for example, Arnott and Small, 1994, Rapoport and Mak, 2018). Empirical evidence dates back to Mogridge (1990), who reported that increased road capacity in central London coincided with a decline in mass transit systems, supporting the D-T paradox. Indirect evidence supporting the Braess paradox has been reported, among others, by Murchland (1970) and Youn et al. (2008).

Horowitz (1984) noted that field data on route choices or choices between private and public transportation cannot resolve whether equilibrium solutions accurately describe behavior in real traffic networks. In their study of the D-T paradox, Dechenaux et al. (2014) argue that "The variety and complexity of ways in which current travel decisions may depend on past network performance, and the complete lack of empirical information about the form of this dependence, renders the usefulness of 'natural experiments' in real networks doubtful" (p.463). An alternative way to test the validity of the equilibrium solutions of these paradoxes is through controlled laboratory experiments, in which participants are incentivized by monetary payoffs contingent on their performance. These experiments present a powerful technique for answering theoretical questions regarding the descriptive power of equilibrium solutions. One of their advantages is their ability to manipulate key variables in traffic networks (e.g., transportation cost, group size, scheduling of traffic, and information structure) that influence commuters' choices and may give rise to the paradoxes described above.

In the theoretical and experimental studies of the D-T paradox, the population typically consists of homogeneous agents: all commuters face the same cost structure and make independent choices between public and private transportation. Departing from this standard setting, we propose a novel scenario that combines elements of the D-T paradox with elements of queuing theory. As in the D-T paradox, our model features two distinct transportation modes, private car and railway, that connect a common origin (e.g., residential area) and a common destination (e.g., central business area). Unlike the standard D-T setting, the road in our scenario is known to have a single bottleneck with a fixed capacity that compels drivers to trade off the cost of queuing delay at the bottleneck with the cost of schedule delay (i.e., the cost of arriving early or late). In our model, each commuter is asked independently to make two decisions: which mode of transportation to use, and if choosing to travel by car, at what time to depart from the origin. An important implication of this scenario is that improving the road (e.g., by reducing service time at the bottleneck) may affect not only the number of commuters choosing to travel by car but also the distribution of their departure times. In equilibrium, although arrival times may differ, travel costs must be equal across commuters.

Another departure from the standard D-T scenario is our assumption that the railway operator's response to changes in passenger volume is perfectly inelastic. Improving the road may attract train commuters from the railway, thereby reducing revenue for the railway operator. If revenue falls while costs remain fixed or rise, the operator may need to increase fares or reduce service frequency, further pushing commuters toward the road. Therefore, road improvement may completely dissipate its benefits and even worsen overall travel cost. However, we assume that the railway operator does

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

not modify its service quality from day to day to avoid disruptions to commuters. Consequently, the travel cost of train commuters is independent of passenger volume and depends solely on train travel time. If the railway operator shortens train travel time, some car commuters may switch from the road to the railway. Therefore, the D-T hypothesis we set to test is that improving the railway may decrease mean travel cost, but improving the road may keep mean travel cost unchanged.

The goal of this study is to experimentally investigate a novel scenario that embeds the D-T paradox in the context of departure-time choice in the morning commute (for example, Vickrey, 1969, Arnott et al., 1990, Arnott et al., 1993). Commuters opting to travel by car must also choose when to depart from the origin to trade off the costs of queuing delay due to congestion at the bottleneck and the schedule delay costs associated with their time of arrival at their common destination. We numerically derive a symmetric mixed-strategy equilibrium for the probability of choosing between private (congestible) and public (non-congestible) transportation, for the probability distribution of (discrete) departure times from the common origin for commuters who choose traveling by car, and for the effects of improving either the road or railway on the equilibrium cost of travel.

The remainder of this paper is organized as follows. Section 2 presents a brief literature review of theoretical and experimental studies related to the present study. Section 3 introduces the notation, formally describes the model, illustrates it with a numerical example, and presents the symmetric mixed-strategy equilibrium solution. Section 4 outlines the experimental design and methods. Section 5 presents and discusses the experimental results. Section 6 proposes a reinforcement learning model. Section 7 concludes the study with a summary of the results.

2. Literature review

The D-T paradox was introduced by Downs (1962) and subsequently studied, among others, by Thomson (1978), Mogridge (1990), Mogridge et al. (1987), Calvert (1997), and Afimeimounga et al. (2005). Ding and Song (2012) have presented a technical introduction that allows a comparison between the P-K-D, Braess, and D-T paradoxes. The D-T paradox concerns a directed network with two parallel routes (or modes of transportation) that connect a common origin with a common destination. Travel on the *private route* generates negative externalities, whereas travel on the *public route* generates either positive externalities or no externalities. Under certain conditions on the cost structure and size of the user population (see, for example, Ding & Song, 2012), the expansion of the road produces no change in or even a decline in social welfare due to a shift in the equilibrium. Network users shift from the public to the private mode of transportation, thereby generating additional negative externalities, which, in turn, may reduce social welfare.

Denant-Boèmont and Hammiche (2010) conducted a study of the D-T paradox that included a two-stage coordination game. In Stage 1, a single participant (called "operator") chose the capacity of the public route, whereas the cost of the private route was chosen exogenously by the experimenter. In stage 2, multiple group members were asked to choose between the private and public modes of transportation. The road capacity was systematically manipulated across iterations. The authors reported three major findings. First, the D-T paradox was realized. Second, the mean route choices approached the equilibrium solution across iterations of the stage game. Third, there was no support for equilibrium play by the operator in Stage 1 of the D-T game. Dechenaux et al. (2014) conducted a second, single-stage, game in which the cost of the public route was determined exogenously and remained fixed over iterations of the stage game whereas the cost of the private route was changed systematically. Travel costs and group sizes were manipulated in the experiment. The authors report evidence in support of the D-T paradox as well as the systematic effects of group size and travel costs on public route choices.

Our study differs markedly from these two previous studies in its setup. The players in our study make two simultaneous decisions: which mode of transportation to choose (congestible versus. non-congestible), and, for those choosing the congestible mode, when to depart from their common

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

origin. Unlike the previous two studies, we do not merely focus on the number of group members who choose to travel by car or train. Rather, we test both the choice of transport and changes in the *distributions* of departure time. The cost structure of the model (presented in Section 3) reflects these two decisions. In addition, whereas the travel cost on the public route ("metro") in the experiment reported by Dechenaux et al. (2014) was set to decrease in the number of users, it was fixed exogenously in our study.

Another major difference from the two previous studies is in what happens on the congestible road and the method used to derive the symmetric mixed-strategy equilibrium. Whereas the cost of traversing the road is commonly determined by a simple function of the number of users, we employ the first-come, first-served (FCFS) queue discipline at the bottleneck, which states that players are handled in the order of their arrival and as soon as space becomes available. This means that the cost to a car commuter departing at time t depends not only on the number of other car commuters choosing the same departure time, but also on the departure time distribution of other car commuters that left at or before time $t-1$. This implies that future departures are irrelevant to the travel costs of past and current departures. As in the two previous studies, we search for the equilibrium using the key property of equalizing travel costs across departure modes and times. However, owing to the FCFS discipline, we must sequentially search for equilibrium probabilities from the earliest departure time through the last departure time.

3. Model and equilibrium

3.1. The model

Two alternative transportation modes connect a common origin and a common destination. Each of n symmetric commuters wishes to travel from the origin to the destination by either car (congestible mode) or train (non-congestible mode). All commuters wish to arrive at their destination at time t^*. If a commuter chooses to travel by car, they select a departure time from a commonly known set of discrete departure times $\{1, 2, \dots, t^*, \dots, t_L\}$. If they choose to travel by train, they board a train that departs at a fixed scheduled time. Commuters independently choose both the mode and the time of departure from the origin; each of them chooses a time $t \in T = \{1, 2, \dots, t^*, \dots, t_L\} \cup \{\text{train}\}$, where "train" stands for the choice of using the public service. Commuters cannot observe the decisions of other group members when they make their own choices.

3.1.1. Travel cost: Private car commuters

First, consider the case of a commuter who chooses to travel by private car. Following Arnott et al. (1990), Arnott et al. (1993), we model congestion as a queue forming at a single section, referred to as a *bottleneck* (e.g., a road segment under construction, a tunnel, a single-lane bridge, or an inspection station), where the FCFS queuing discipline is applied. The bottleneck serves only one car at a time, with a service time of s (> 1) time units per car. Thus, if a driver enters the bottleneck at time t, only that driver occupies the bottleneck for s time units. As in Arnott et al. (1990), Arnott et al. (1993), this model assumes zero travel time between the origin and bottleneck entrance and between the bottleneck exit and destination. This assumption implies that a commuter traveling by car arrives at the bottleneck as soon as they depart from the origin and arrives at the destination as soon as they depart the bottleneck. If multiple commuters depart simultaneously (and consequently arriving at the bottleneck at the same time), entry order is determined randomly with equal probability.

In addition to the fixed service time s, car commuters may have to wait for an additional time at the bottleneck.[1] Suppose that car commuter i departs from the origin at time t. Then, their waiting time depends on the number of car commuters who have already departed before time t and have not exited the bottleneck, and the ones departing at time t. For each t, we denote by $\omega_{t-1} \in \Omega$ the

[1] Following Arnott et al. (1990), Arnott et al. (1993), service time is kept fixed. Fixed service time also considerably simplifies the experiment and helps avoid unnecessary behavioral noise introduced by additional uncertainty.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

state of the road at the end of $t-1$, which is characterized by a vector with three components $[r\ q\ e]$, where

- r: number of other commuters remaining at the origin,
- q: number of other car commuters forming a queue at the bottleneck,
- e: time elapsed since the bottleneck became occupied.

Suppose that $\omega_{t-1} = [r\ q\ e]$, that the d other commuters (out of r) also depart by car at time t, and that k of them are ahead of commuter i after the tie is broken with equal probability. Then, commuter i's waiting time falls into one of the eight exhaustive cases below:

$$W(t,\omega_{t-1},k) = \begin{cases} 0 & \text{if } r=0, q=0, e=0 \\ s(q-1)+(s-1) & \text{if } r=0, q>0, e=0 \\ sq+(s-1-e) & \text{if } r=0, \text{any } q, 0<e<s \\ sq & \text{if } r=0, \text{any } q, e=s-1 \\ sk & \text{if } r>0, q=0, e=0 \\ s(q-1+k)+(s-1) & \text{if } r>0, q>0, e=0 \\ s(q+k)+(s-1-e) & \text{if } r>0, \text{any } q, 0<e<s \\ s(q+k) & \text{if } r>0, \text{any } q, e=s-1 \end{cases}$$

Commuter i's travel time by car is equal to the sum of their service and waiting times. Then, given ω_{t-1} and a value of k, i's travel time from departing at time t is

$$T(t,\omega_{t-1},k) = s + W(t,\omega_{t-1},k).$$

Car commuters experience two sources of disutility (i.e., cost). The first source results from travel time: the longer the travel time, the higher the cost to the commuters. Additionally, each commuter may incur the cost of arriving early or late, which is known as the *schedule delay cost*. This second source is because a car commuter must wait until work starts if they arrive at work too early, and they bear a penalty if they arrive too late. The degree of schedule delay is measured relative to the common desired arrival time t^*. Therefore, for given ω_{t-1} and k, car commuter i's total travel cost is given by

$$\begin{aligned} C(t,\omega_{t-1},k) = {} & \alpha T(t,\omega_{t-1},k) \\ & + \beta \max\left\{0, t^* - \left(t + T(t,\omega_{t-1},k)\right)\right\} \\ & + \gamma \max\left\{0, \left(t + T(t,\omega_{t-1},k) - t^*\right)\right\} \end{aligned}$$

where the parameters α, β, and γ denote the unit costs of traveling, arriving early, and arriving late, respectively. In accordance with Small (1982), we assume $\gamma > \alpha > \beta$.

3.1.2. Travel cost: Public train commuter

Unlike the standard D-T scenario, the travel cost for train commuters is assumed to be fixed. There is a single commuter train that departs from the origin at time t_{train} ($1 \le t_{\text{train}} < t^*$) and arrives at the common destination precisely at the desired arrival time t^*, regardless of how crowded it is. The cost of traveling by train is given by

$$C_{\text{train}} = \alpha(t^* - t_{\text{train}}),$$

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

Table 1. Numerical example for $(n, t^*, t_L, t_{train}, s, \alpha, \beta, \gamma) = (20, 50, 60, 26, 4, 10, 5, 25)$

			Travel time			
Commuter	Transportation mode	Departure time	Waiting time	Service time	Arrival time	Travel cost
11	Car	1	0	4	5	265
7	Car	5	0	4	9	245
5	Car	11	0	4	15	215
20	Car	16	0	4	20	190
9	Car	17	3	4	24	200
14	Car	19	5	4	28	200
8	Car	20	8	4	32	210
12	Car	20	12	4	36	230
16	Car	22	14	4	40	230
6	Car	24	16	4	44	230
19	Car	26	18	4	48	230
4	Car	34	14	4	52	230
1	Car	40	12	4	56	310
3	Car	40	16	4	60	450
13	Car	45	15	4	64	540
2	Train	26	24		50	240
18	Train	26	24		50	240
17	Train	26	24		50	240
15	Train	26	24		50	240
10	Train	26	24		50	240

which reflects only the time spent in transit; it does not include train fares and assumes no economies of scale.[2] We also assume that $t_{train} > s$ to ensure that there is no symmetric equilibrium in which all commuters travel by train.[3]

3.2. A numerical example

Table 1 presents a numerical example that illustrates the computation of travel cost in our study. Columns 2 through 7 of Table 1 list the transportation mode, departure time from the origin, waiting time at the bottleneck, service time, arrival time at the destination, and individual travel cost, respectively. In this example, the parameter tuple is $(n, t^*, t_L, t_{train}, s, \alpha, \beta, \gamma) = (20, 50, 60, 26, 4, 10, 5, 25)$. The table shows that departing by car does not always result in a lower travel cost, that departing too late can lead to a high cost, and that traveling by train is not always optimal.

The earliest departure occurs at $t = 1$, by commuter 11. Since this commuter encounters no congestion, the travel time is 4 units (i.e., zero waiting time plus 4 units of service time). Commuter 11 arrives at the destination at time $t = 5$, 45 time units before $t^* = 50$ (the desired arrival time). Consequently, the travel cost is 265 ($= 10 \times 4 + 5 \times 45$). Car commuters 8 and 12 depart at time

[2]Tabuchi (1993) considers several travel cost structures for train commuters. The travel cost per train commuter, consisting of both marginal and fixed costs, exhibits economies of scale with respect to the number of train commuters. He analyzes a variety of railroad fare pricing schemes, including marginal cost pricing, and average cost pricing.

[3]Let p_{train} denote the probability of traveling by train. If there were a symmetric mixed-strategy equilibrium with $p_{train} = 1$, then at least one train commuter would benefit from deviating unilaterally to another mixed strategy in which, for example, $p_{t^*-s} = 1$. This leads to a contradiction.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

$t = 20$. At that time, commuter 9 is being served at the bottleneck, and commuter 14 is waiting in the queue. The tie between commuters 8 and 12 is resolved randomly with equal probability, in favor of commuter 8. Commuter 8 waits 8 units at the bottleneck until the services for commuters 9 and 14 are completed. The travel time for commuter 8 is 12 units (8 units of waiting time plus 4 units of service time), leading to an arrival time of $t = 32$, which is 18 units before time t^*. The travel cost is 210 ($= 10 \times 12 + 5 \times 18$). Commuter 12, who enters the bottleneck after commuter 8, waits 12 units of time until the services for commuters 9, 14, and 8 are completed. The travel time is 16 units (12 units of waiting time plus 4 units of service time), resulting in an arrival at $t = 36$, 14 units before t^*. The travel cost in this case is 230 ($= 10 \times 16 + 5 \times 14$).

Five commuters (2, 10, 15, 17, and 18) travel by train, departing at $t = 26$ and arriving at $t^* = 50$ (the desired arrival time). Since the travel time for each train commuter is 24 units, the cost for each is 240 ($= 10 \times 24$). The right-hand column of Table 1 shows that commuter 20, who departs at $t = 16$, incurs the lowest travel cost ($C = 190$), whereas commuter 13, who departs at $t = 45$ by car, incurs the highest travel cost ($C = 540$). The mean travel cost among the 15 car commuters is $249.7 > 240$.

3.3. Symmetric mixed-strategy equilibrium

This study considers the symmetric equilibrium in mixed strategies for two reasons. First, the commuters are ex-ante symmetric; the same set of strategies and cost structure hold for every commuter. Second, in asymmetric equilibria, commuters must behave differently, resulting in asymmetry in their travel costs. This implies that, without any opportunity for pre-play communication, commuters must tacitly reach an agreement on who incurs a higher travel cost and who enjoys a lower travel cost.

Figure 1 displays the equilibrium cumulative distributions of strategies, 60 discrete departure times and the train option, for the three experimental conditions: Baseline, Railway Improved, and Road Improved.[4] In the figure, the train option appears at the far right of the horizontal axis. The three experimental conditions differ in two key parameters: the service time at the bottleneck (s) and the schedule train departure time (t_{train}): $(s, t_{\text{train}}) = (4, 26)$ in the Baseline condition, $(4, 30)$ in the Railway Improved condition (representing a shorter train travel time), and $(3, 26)$ in the Road Improved condition (representing a reduced bottleneck service time). The corresponding equilibrium costs of travel are 240 in the Baseline and Road Improved conditions and 200 in the Railway Improved condition. The equilibrium probabilities of choosing to travel by train are 0.467 for the Baseline condition, 0.595 for the Railway Improved condition, and 0.206 for the Road Improved condition. The equilibrium support for departure times ranges from $t = 7$ to $t = 47$ in the Baseline condition, $t = 15$ to $t = 46$ in the Railway Improved condition, and $t = 6$ to $t = 50$ in the Road Improved condition.

Figure 1 illustrates two key implications of the equilibrium solution. First, shortening the travel time by train from 24 to 20 results in a shift of car commuters to the improved railway and a consequent redistribution of the departure times of car commuters until the travel costs on both modes equalize at 200. Thus, improving the congestion-free railway reduces the travel costs on both transportation modes. Second, improving the road by reducing the duration of service time at the bottleneck from 4 units to 3 units induces a shift of train commuters to the upgraded road and a redistribution of the departure times of car commuters until the travel costs for both the road and railway are equalized at 240. In other words, the benefits resulting from expanding the bottleneck are completely dissipated, in accordance with the D-T paradox. Based on the equilibrium solution, we derive the following testable hypothesis concerning the D-T paradox:

Hypothesis. (D-T Hypothesis): *Railway improvement decreases the mean travel cost, whereas road improvement does not change it.*

[4]See Sections 1 and 2 of the supplementary material for a detailed discussion about how the symmetric mixed-strategy equilibrium was derived, and Section 3 for the equilibrium probabilities.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

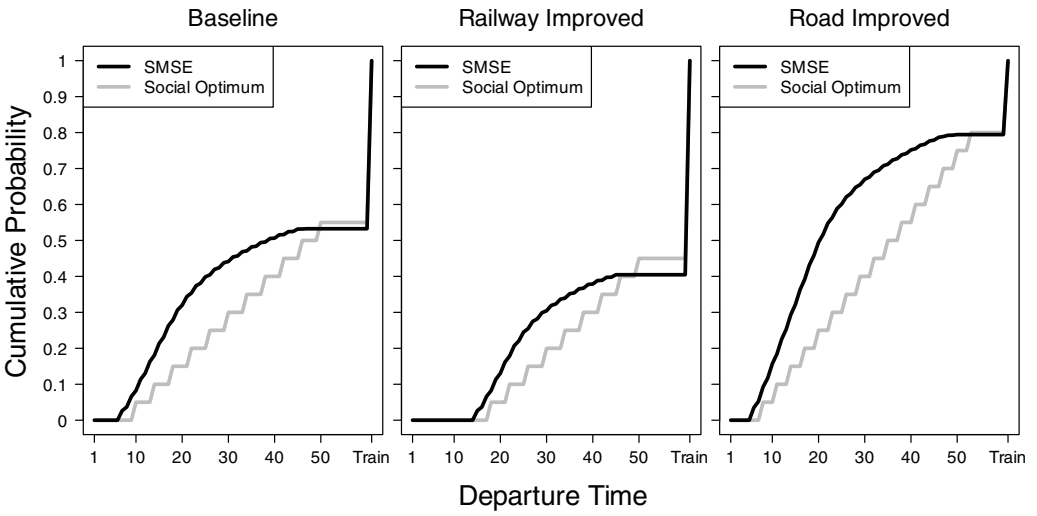


Figure 1. Equilibrium cumulative probability distributions of strategies

Table 2. Number of car commuters, their departure times, and the mean, minimum, and maximum travel costs in the social optimum for $(n, t^*, t_L, \alpha, \beta, \gamma) = (20, 50, 60, 10, 5, 25)$

			Travel cost		
Condition	No. of car commuters	Departure times	Mean	Minimum	Maximum
Baseline $(s, t_{\text{train}}) = (4, 26)$	11	10, 14, ..., 46, 50	180	40	240
	12	6, 10, ..., 46, 50		40	
	12	7, 11, ..., 47, 51		55	
	12	8, 12, ..., 48, 52		50	
	12	9, 13, ..., 49, 53		45	
	12	10, 14, ..., 50, 54		40	
	13	6, 10, ..., 50, 54		40	
Railway Improved$(s, t_{\text{train}}) = (4, 30)$	9	18, 22, ..., 46, 50	161	40	200
	10	14, 18, ..., 46, 50			
Road Improved$(s, t_{\text{train}}) = (3, 26)$	16	8, 11, ..., 50, 53	151.5	30	240
	17	5, 8, ..., 50, 53			

The symmetric mixed-strategy equilibrium is inefficient in minimizing total travel cost. As a complementary benchmark, we heuristically derived the social optimum for each condition, following Arnott et al. (1990). First, all car commuters depart at different departure times. Second, the departure times of any two car commuters are separated by an interval equal to the bottleneck service time. These two conditions ensure that, during rush hour, no queue forms at the bottleneck and the bottleneck is never idle. Third, neither the first nor the last car commuter wants to unilaterally move to the other end of the rush hour.

Table 2 presents the number of car commuters, departure times, and the mean, minimum, and maximum travel costs in the social optimum separately for each condition. We identified seven such combinations in the Baseline condition, and two combinations each in the Rail Improved and Road Improved conditions. Two observations follow. First, although total travel cost is minimized in the social optimum, individual travel costs vary considerably across commuters. This implies that commuters incurring higher travel costs, such as car commuters assigned early or late departure times or train commuters, have a strong incentive for unilateral deviation. Indeed, these social optima do not

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

constitute pure-strategy equilibria.[5] Second, the three conditions differ in the extent of inefficiency. The ratio of the equilibrium travel cost to the travel cost in the social optimum, known as the *price of anarchy* (Koutsoupias and Papadimitriou, 1999, Papadimitriou, 2001, Mak and Rapoport, 2013), is 1.33 $\left(= \frac{240}{180} \right)$ in the Baseline condition, 1.24 $\left(= \frac{200}{161} \right)$ in the Railway Improved condition, and 1.58 $\left(= \frac{240}{151.5} \right)$ in the Road Improved condition. Thus, relative to the Baseline condition, railway improvement enhances efficiency, whereas road improvement worsens efficiency.[6]

Figure 1 also displays the socially optimal cumulative relative frequency distribution alongside the symmetric mixed-strategy equilibrium for each condition.[7] The social optimum in this figure corresponds to the first combination of each condition listed in Table 2. This facilitates a visual comparison of the two distributions. Whereas the equilibrium distribution is more concentrated, the social optimum spreads departure times more evenly. This reflects a trade-off between individual gains and collective efficiency. The equilibrium leads to queuing at the bottleneck, whereas the social optimum mitigates congestion through tacit coordination among car commuters.

4. Method

4.1. Participants

The experiment was conducted in the computer laboratory of the Max Planck Institute of Economics, Jena, Germany.[8] A total of 240 undergraduate students from Friedrich Schiller University Jena voluntarily participated in a computer-controlled decision-making experiment, with monetary payoffs contingent on individual performance. We conducted 12 sessions with 20 participants each, running four sessions per experimental condition. Each session lasted approximately 90 minutes.

4.2. Procedure

In every round, each participant earns a per-round payoff equal to the difference between a fixed reward R for reaching the destination and their travel cost. They accumulated these per-round payoffs for all 40 rounds of play. The three experimental conditions shared the following parameter values: $t_1 = 1$ (8:01), $t_L = 60$ (9:00), $t^* = 50$ (8:50), $R = 500$, $\alpha = 10$, $\beta = 5$, and $\gamma = 25$. They differed in the values of service time (s) or train departure time (t_{train}). Table 3 presents our experimental design. The effect of change in train departure time (i.e., travel time on the public route) was tested by comparing the Baseline and Railway Improved conditions. The effect of a change in service time was tested by comparing the Baseline and Road Improved conditions. The three conditions also differed from one another in the conversion rate of points into euros: 1200 = 1 euro, 1300 = 1 euro, and 1200 = 1 euro for the Baseline, Railway Improved, and Road Improved conditions, respectively.

Written instructions were handed to participants at the beginning of each session.[9] The instructions describe the bottleneck game as a simulation of the decisions faced by commuters traveling from a common origin to common destination each morning. Participants were informed that they would play 40 identical rounds, in each of which they would choose between traveling by car or by train. In the former case, they would travel on a road with a bottleneck characterized by a fixed service time and would select their departure time from the origin. In the latter case, they would travel by train,

[5]For example, in the first combination of the Baseline condition in Table 2, a train commuter incurs a travel cost of 240. This commuter would benefit by switching to car travel and departing at $t = 9$: the resulting travel cost would be $4 \times 10 + 5 \times (50 - 13) = 225$, which is lower than 240.

[6]A smaller price of anarchy implies greater efficiency.

[7]The social optimum is a pure strategy, whereas the equilibrium is a mixed strategy. To display them side-by-side in the same figure, we computed the cumulative relative frequency distribution for the social optimum.

[8]The institute is now known as the Max Planck Institute of Geoanthropology.

[9]Both the original German instructions and their English translation are available at https://doi.org/10.17605/OSF.IO/R678H.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

Table 3. Experimental Design

Condition	No. of sessions	No. of participants per session	Service time (s)	Train departure time (t_{train})	Expected travel cost
Baseline	4	20	4	26	240
Railway Improved	4	20	4	30	200
Road Improved	4	20	3	26	240

which departs at a predetermined time and arrives at the destination precisely at the desired arrival time t^*. A numerical example illustrated the financial implications of choosing different departure times (see Table 1 in the instructions).

The payoff equations (see Section 2 above) were explained in detail and illustrated with a numerical example. To prevent negative cumulative payoffs (referred to as "bankruptcy"), each participant received an initial endowment of 10,000 points; tokens were added to or subtracted from this balance depending on the the payoff (positive or negative) earned in each round. Table 1 in the instructions is identical to Table 1 in the main text, except that it reports payoffs ($R - C$) rather than costs (C). The decision screen (Figure 1 in the instructions) displayed 60 alternative departure times. The outcome screen (Figure 3 in the instructions) showed the departure times of all group members, the participant's payoff for the round, and their cumulative payoff. At the end of the session, participants were paid individually and dismissed. Excluding a common show-up fee of 2.50 euros, the mean payoffs in the Baseline, Rail Improved, and Road Improved conditions were 18.92, 18.76, and 19.23 euros, respectively.

5. Results

5.1. Aggregate behavior

Figure 2 presents side-by-side comparisons of the predicted, observed, and socially optimal cumulative relative frequency distributions of strategies, aggregated at the session level. The observed distributions were computed separately for rounds 1-10, 11-20, 21-30, and 31-40. In each panel, the thick line depicts the symmetric mixed-strategy equilibrium, whereas the thin lines represent the session-level distributions. In the Baseline and Railway Improved conditions, excess usage of the road persisted throughout the sessions. In the Road Improved condition, the observed distribution consistently remained within the area bounded by the equilibrium and socially optimal distributions, suggesting that car commuters spread out their departure times more than predicted, thereby alleviating congestion at the bottleneck. Thus, the observed behavior in the Road Improved condition lay between the equilibrium prediction and the social optimum, partially improving efficiency relative to equilibrium play, yet falling short of fully reaching the social optimum.

Visual inspection suggests that the observed distribution moved closer to the equilibrium distribution over time in each condition. We hypothesized that, as participants gained experience throughout the experiment, aggregate behavior would approach equilibrium play. To measure the distance between the observed and predicted distributions, we invoke the deviation index proposed by Stein et al. (2007) for each round:

$$DI_{jr} = \sqrt{\sum_{t \in T} \left(p_{jrt} - p_t^* \right)^2}$$

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

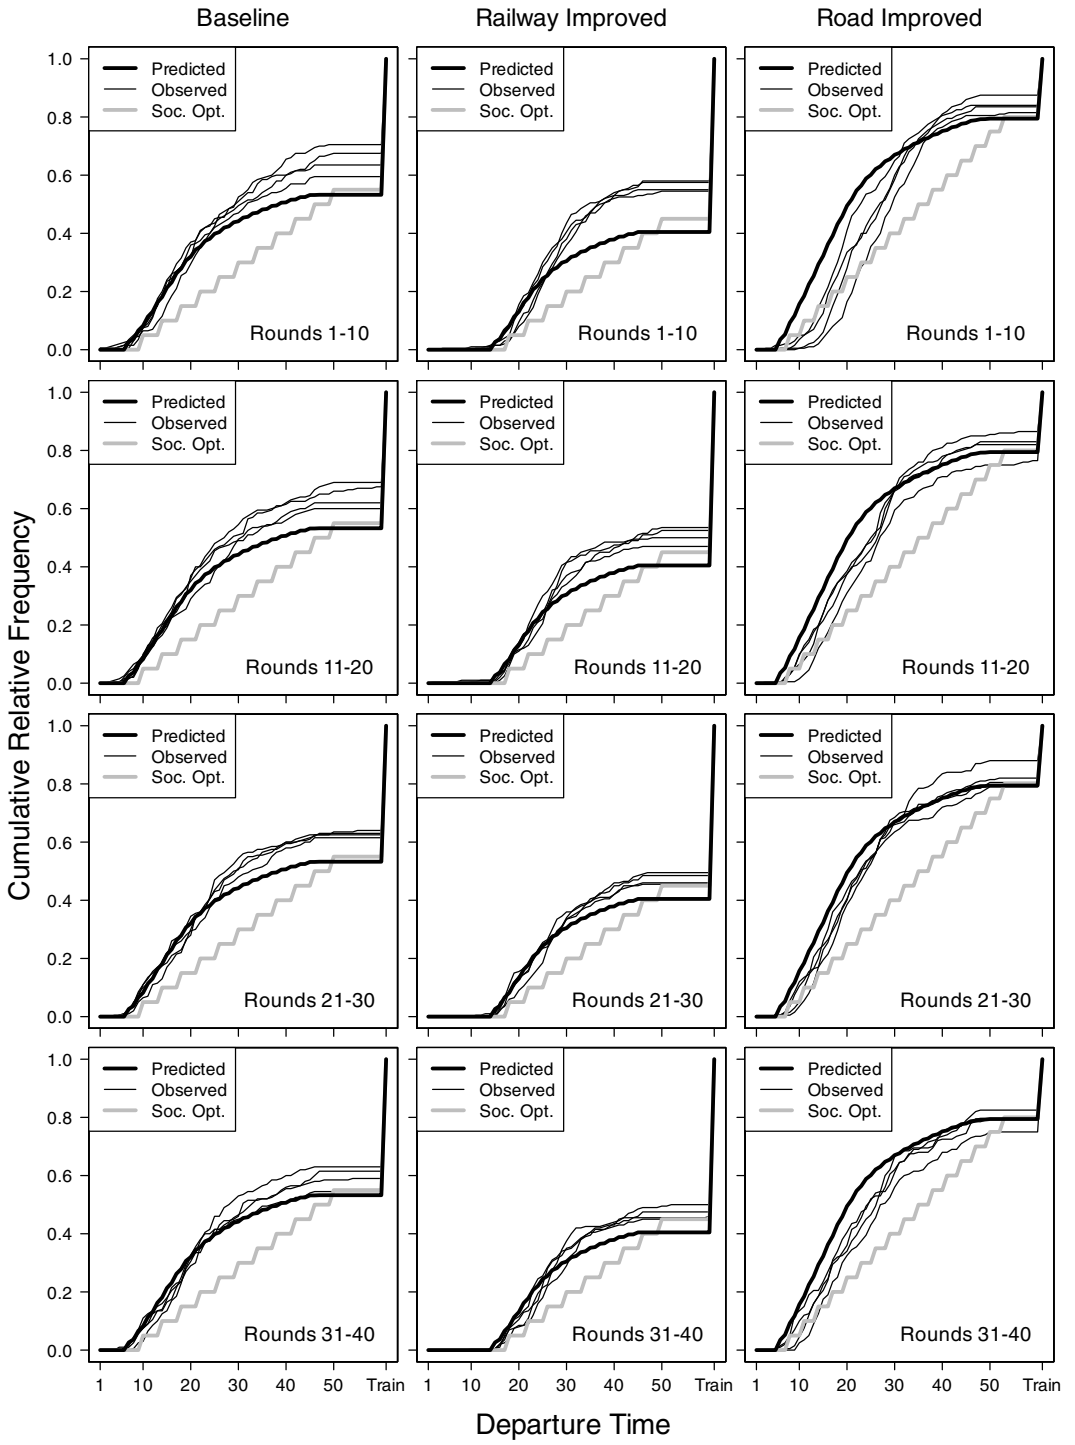


Figure 2. Predicted, observed, and socially optimal cumulative relative frequency distributions of strategies

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

where DI_{jr} is the deviation index for session j in round r, p_{jrt} is the relative frequency of departure choice $t \in T = \{1, 2, \ldots, 60\} \cup \{\text{Train}\}$ within session j in round r, and p_{t^*} is the probability of choosing t under the symmetric mixed-strategy equilibrium.[10]

We applied the one-sided Mann–Kendall trend test to assess whether the deviation index exhibited a decreasing trend over 40 rounds in each session. Each condition consists of four sessions, resulting in 12 independent trend tests. The results indicated a significant decreasing trend in 11 of the 12 sessions ($p < 0.05$). However, Session 2 (the second session in the Baseline condition) did not reject the null hypothesis ($S = -26$, $Z = -0.291$, $p = 0.385$), suggesting no clear trend over the 40 rounds in that session.[11]

Result 1. *Aggregate behavior approached equilibrium play with more experience.*

Given that most sessions exhibit a decreasing trend in the deviation index, unless stated otherwise, we focus on the last ten rounds (rounds 31-40) for the remainder of the analyses. These rounds best reflect the participants' more experienced decisions after the initial 30 rounds of play.

5.2. Testing the D-T hypothesis

The D-T hypothesis posits directional predictions regarding mean travel costs; railway improvement reduces the mean travel cost, whereas road improvement leaves it unchanged. Table 4 reports the predicted and observed means of travel cost separately for all commuters (third column) and for car commuters (fourth column). The mean travel costs for all commuters were higher in the Baseline and Railway Improved conditions (247.81 versus 240 and 204.81 versus 200) and lower in the Road Improved condition than the equilibrium travel costs (223.72 versus 240). The data from the Road improved condition appeared to contradict the D-T hypothesis.

To formally test whether our data support the D-T hypothesis, we estimated the following linear mixed-effects model with the Baseline condition as the reference group:[12]

$$\Delta\text{TC}_{it} = \beta_0 + \beta_1 \text{Railway_Improved} + \beta_2 \text{Road_Improved} + r_i + r_{j[i]} + \epsilon \tag{1}$$

where ΔTC_{it} represents participant i's travel cost in round t minus the equilibrium travel cost for the reference group, namely 240. Railway_Improved is a dummy variable taking the value of 1 if the condition is Railway Improved and 0 otherwise, and Road_Improved is a dummy variable taking the value of 1 if the condition is Road Improved and 0 otherwise. r_i is the random effect for player i, $r_{j[i]}$ is the random effect for session j in which player i participates, and ϵ is a general random term.[13] We expected an improvement of the railway to decrease and an improvement of the road not to change the travel cost, that is, $\beta_1 < 0$ and $\beta_2 = 0$.

Table 5 presents the estimation results for Model (1). The estimated intercept is positive but insignificant. This indicates that participants in the Baseline condition incurred travel costs close

[10] Although the deviation index provides a useful measure of how closely the observed distribution approximates the equilibrium prediction, it cannot equal zero in our experimental setup, which features 20 participants choosing among 61 pure strategies (60 discrete departure times plus the train). A larger number of participants and strategies would be required for the index to approach zero.

[11] We used the R package `trend` (version 1.1.6) by Pohlert 2023 to conduct the Mann-Kendall trend tests. All session-level, round-to-round deviation index values are reported in Section 4 of the supplementary material.

[12] We used the R packages `lme4` (version 1.1.35.5) by Bates et al. 2015 for estimation and `texreg` (version 1.39.4) by Leifeld (2013) for table output.

[13] The design of our experiment called for a fixed group of 20 participants who took part in 40 identical rounds of play with outcome feedback on the choices of all their group members in the previous round. Consequently, there are two sources of dependencies in our data that we must address. The first is individual dependencies in consecutive choices of the same participant that are unlikely to be independent. The second is group dependencies due to shared strategic and environmental influences.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

Table 4. Predicted and observed means of travel cost (rounds 31-40)

		Observed	
Condition	Predicted	All commuters	Car commuters
Baseline	240	247.81	253.12
Railway Improved	200	204.18	208.87
Road Improved	240	223.72	219.42

Table 5. Estimation results for model (1)

	Dependent variable ΔTC
(Intercept)	7.806 (6.592)
Railway_Improved	−43.625*** (9.323)
Road_Improved	−24.087** (9.323)
AIC	27356.957
BIC	27391.657
Log Likelihood	−13672.479
Number of observations	2400
Number of participants	240
Number of sessions	12
Variance: Participant (Intercept)	224.313
Variance: Session (Intercept)	137.561
Variance: Residual	5010.469

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$. Standard errors are in parentheses.

Table 6. Predicted and observed means of number of participants traveling by car (rounds 31-40)

Condition	Predicted	Observed
Baseline	10.65	11.90
Railway Improved	8.09	9.43
Road Improved	15.89	15.82

to the equilibrium value of 240. The coefficient of Railway_Improved is negative and highly significant, consistent with the D-T hypothesis. However, the coefficient of Road_Improved is also significantly negative, contradicting the hypothesis that road improvements do not affect mean travel costs. Therefore, we confirm the evidence partially supporting our D-T hypothesis.

Result 2. *Contrary to the D-T hypothesis, not only railway improvement but also road improvement reduced the mean travel cost.*

5.3. Behavior of car commuters

Table 6 presents the predicted and observed mean number of participants (out of 20) who traveled by car in each condition. Although the road was slightly overused in the Baseline and Railway Improved

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

Table 7. Predicted and observed means of departure, waiting, early arrival, and late arrival times (rounds 31-40)

		Departure time					
Condition		Mean	Earliest	Latest	Waiting time	Early arrival time	Late arrival time
Baseline	Predicted	20.05	7	47	9.21	17.55	0.81
	Observed	21.98	5	54	9.86	15.62	1.46
Railway Improved	Predicted	25.31	15	46	7.07	14.33	0.71
	Observed	27.15	10	60	7.93	12.09	1.17
Road Improved	Predicted	19.69	6	50	10.30	17.74	0.73
	Observed	23.48	5	50	8.42	16.10	0.99

conditions, the observed means closely matched the symmetric mixed-strategy equilibrium predictions. However, the estimation results of Model (1) show that the observed mean travel cost was significantly lower than predicted not only in the Railway Improved condition but also in the Road Improved condition (see Table 5). Because the cost of traveling by train was fixed (240 in the Baseline and Road Improved conditions and 200 in the Railway Improved condition), the differences between the predicted and observed mean travel costs reflect only the behavior of car commuters. The fourth column of Table 4 shows the mean travel cost only for car commuters. The value is higher in the Baseline and Railway Improved conditions and lower in the Road Improved conditions than that for all commuters. Hence, car commuters are solely responsible for the increase in mean travel cost in the former two conditions and the decrease in the Road Improved condition.

How does the observed behavior of car commuters differ from equilibrium play? We compared the predicted and observed behavior of car commuters using four time-based metrics, departure time, waiting time, early arrival time, and late arrival time, as shown in Table 7. The predicted earliest and latest departure times correspond to those in the support of the symmetric mixed-strategy equilibrium.

In all three conditions, car commuters tended to delay their departure times relative to equilibrium predictions. This pattern of delayed departures is accompanied by shifts in arrival patterns: the observed mean early and late arrival times are lower and higher than predicted, respectively, resulting in arrival times being shifted to later time slots.

Another discernible deviation concerns waiting time. In the Baseline and Railway Improved conditions, the observed mean waiting times exceeded the predicted ones, indicating that car commuters experienced more congestion than expected. In contrast, the observed mean waiting time is shorter than predicted in the Road Improved condition, suggesting that car commuters more effectively avoided queues at the bottleneck.

To formally test these patterns, we estimated the following intercept-only mixed-effects model for each time-based measure:

$$\Delta_{it} = \beta_0 + r_i + \epsilon \tag{2}$$

where Δ_{it} denotes participant i's observed time-based measure in round t minus the equilibrium value. This model evaluates whether the observed behavior systematically deviated from equilibrium predictions.[14] Table 8 presents the estimated intercepts of Model (2).[15] The intercepts for departure time are all significantly positive, confirming delayed departures in all three conditions. Therefore, we concluded the following result.

[14]In estimation of Model (2) with the individual- and session-specific random effects, we found that the session-level random intercept was estimated to have zero variance, and likelihood ratio tests confirmed that its inclusion did not improve model fit. Accordingly, we excluded the session-specific random effect, $r_{j[i]}$, from this model.

[15]The full estimation results for Model (2) are relegated to Section 5 of the supplementary material.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

Table 8. Estimated intercept of Model (2)

Time-based metrics	Baseline	Railway Improved	Road Improved
Departure time	2.321* (0.926)	2.612** (0.949)	4.230*** (0.960)
Waiting time	0.957 (0.620)	0.792 (0.517)	−1.741*** (0.443)
Early arrival time	−2.457* (1.191)	−2.316* (1.126)	−2.138 (1.143)
Late arrival time	0.853** (0.299)	1.038** (0.363)	0.387* (0.187)
Number of observations	476	377	633
Number of participants	71	59	78

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$. Standard errors are in parentheses.

Result 3. *In all three conditions, car commuters departed later than predicted.*

We also found that the intercepts for late arrival time were significantly positive in all conditions, whereas those for early arrival time were significantly negative in the Baseline and Railway Improved conditions. This pattern indicates an overall rightward shift in arrival times.

Result 4. *In all conditions, late arrivals occurred later than predicted. In the Baseline and Railway Improved conditions, early arrivals also occurred later than predicted.*

In terms of waiting time, the observed mean waiting times were higher than predicted in both the Baseline and Railway Improved conditions, whereas the Road Improved condition exhibited a significant reduction in waiting time. This pattern was confirmed by the intercept-only mixed-effects model, where the only significant negative intercept was found in the Road Improved condition. Car commuters in the Road Improved condition, where the D-T hypothesis was not supported, appeared to tacitly coordinate their departure times to avoid unnecessary waiting at the bottleneck.

Result 5. *Car commuters in the Road Improved condition reduced their waiting times, whereas those in the Baseline and Railway Improved conditions did not.*

The implications of Results 3-5 are reflected in the travel cost components summarized in Table 9. We added rows for the simulated results, which were generated from a reinforcement learning model calibrated using the experimental data (details provided in Section 6). In both the Baseline and Railway Improved conditions, car commuters experienced a decrease in early arrival cost and an increase in late arrival cost. These opposing effects offset each other, and the travel time cost remained close to the predicted values. Consequently, the observed mean travel costs in these two conditions were consistent with the equilibrium predictions.

In contrast, the Road Improved condition showed a different pattern. In this condition, car commuters were able to avoid queues at the bottleneck, leading to a significant reduction in travel time cost. Although the mean late arrival cost increased, the gain from reduced travel time outweighed this increase. Consequently, the observed mean travel cost significantly fell below the predicted value.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

Table 9. Predicted, observed, and simulated means of travel cost and its components for car commuters (rounds 31-40)

		Baseline	Railway Improved	Road Improved
Travel cost	Predicted	240	200	240
	Observed	253.12	208.87	219.42
	Simulated	251.75	215.65	221.08
Travel time cost	Predicted	132.06	110.71	133.02
	Observed	138.59	119.26	114.17
	Simulated	167.05	148.05	125.95
Early arrival cost	Predicted	87.76	71.65	88.71
	Observed	78.08	60.44	80.49
	Simulated	51.45	31.85	74.38
Late arrival cost	Predicted	20.18	17.64	18.27
	Observed	36.45	29.18	24.76
	Simulated	33.25	35.75	20.75

5.4. Discussion

In our experiment, the D-T paradox predicted that improving the railway system (by reducing its travel time) should lower the mean travel cost, whereas improving road capacity (by reducing the bottleneck service time) should not affect the mean travel cost in equilibrium. Contrary to this theoretical prediction, our findings showed that both the Railway Improved and Road Improved conditions led to significant reductions in mean travel cost. In particular, car commuters in the Road Improved condition appeared to tacitly coordinate their departure times, thereby substantially reducing their waiting time at the bottleneck.

These results suggest that the D-T paradox did not fully manifest in our experimental setting. However, when isolating the mode choice dimension from the departure time dimension, the aggregate observed behavior was well accounted for by the equilibrium predictions (see Table 6). This pattern resonates with previous experimental studies on market entry games (Rapoport, 1995, Sundali et al., 1995, Erev and Rapoport, 1998, Rapoport et al., 1998), route choice games (Anderson et al., 2007, Selten et al., 2007), and paradoxes in traffic networks (Morgan et al., 2009, Rapoport et al., 2009, Denant-Boèmont and Hammiche, 2010, Dechenaux et al., 2014), which focus on choices among a small number of strategies, such as routes, modes, and markets. In these simpler settings, participants may be better able to coordinate equilibrium strategies, thereby giving rise to the paradoxical outcomes predicted by theory. In particular, studies on market entry games by Rapoport and colleagues (for example, Sundali et al., 1995, Erev and Rapoport, 1998) have documented quick convergence toward equilibrium behavior. In contrast, although our participants' aggregate behavior gradually moved closer to equilibrium over time (Result 1), it did not fully converge. In the Road Improved condition, in particular, the observed distribution of departure times consistently positioned itself between the equilibrium and the social optimum. This suggests a partial improvement in efficiency, deviating from the fundamental prediction of the paradox.

The departure from the D-T paradox may stem from the added complexity of our setting. Unlike route choice games or market entry games, our environment introduces a richer strategy space: participants choose not only whether to travel by car or train but also, if traveling by car, when to depart. Furthermore, the congestion externality caused by car commuters extends not only to those who choose the same departure time but also to those who depart later. This intertemporal spillover of the negative externality increases the difficulty of strategic coordination and may hinder convergence to the equilibrium outcomes consistent with the D-T paradox.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

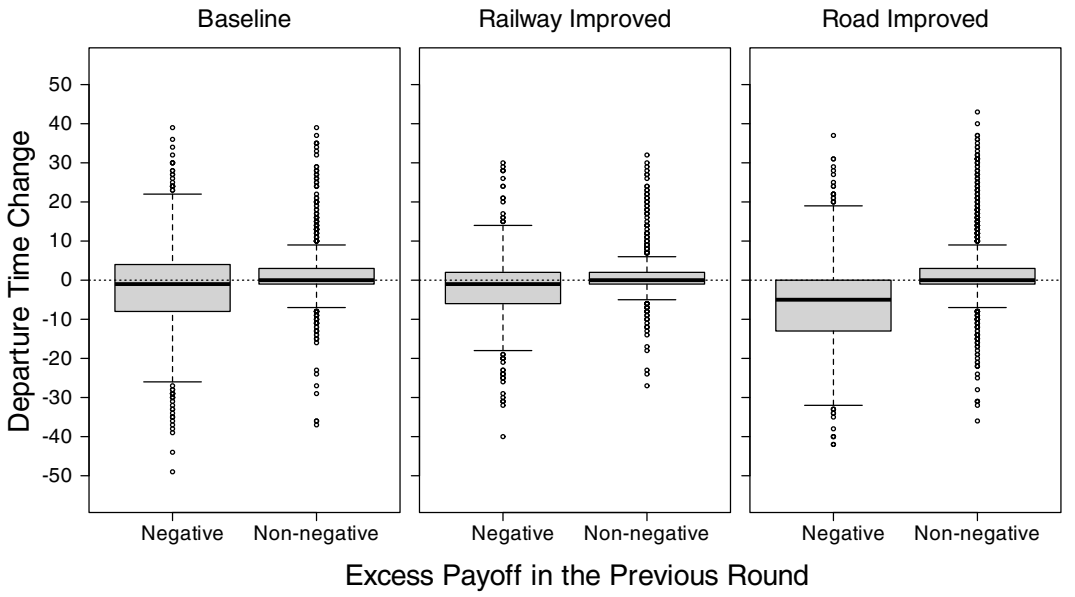


Figure 3. Effect of previous excess payoff on change in departure time

6. Learning

These findings indicate that equilibrium reasoning alone may not fully account for the observed behavior, raising the possibility that participants adapted through experience-based processes. To explore this idea, this section examined whether a simple reinforcement learning model could replicate the observed behavior at the aggregate level. Our model was a two-parameter reinforcement learning model built on the framework of Erev and Roth (1998), with an additional mechanism for experimentation. Reinforcement learning has been widely applied in economic experiments to model how individuals adapt their strategies based on past experiences. In the context of mode and departure time choices, commuters reinforced successful decisions while occasionally experimenting with alternative options.

We began by examining how participants' past outcomes shaped their subsequent decisions. Our goal was to ground the learning model in the observed individual behavior. Importantly, we did not aim to construct a general model of learning behavior. Instead, we sought to develop a simple, ad hoc reinforcement learning model that captured the key features of our participants' behavior while remaining as parsimonious as possible.

6.1. Effects of previous outcomes on subsequent decisions

How did our participants adjust their strategies across the 40 rounds of play? In our experiments, at the end of every round, each participant was informed of their choice of transportation mode, departure time, travel time, arrival time, individual payoff, and the choices of transportation mode and departure time by their group members. This feedback coupled with their personal experience may have guided participants to learn the structure of the game, adjusting their choices over time, influencing which transportation mode to use, and if traveling by car, at what time to depart, in the next round.

An excess payoff is defined as the actual payoff minus the equilibrium payoff (260 in the Baseline and Road Improved conditions and 300 in the Railway Improved condition). Figure 3 illustrates how the sign of excess payoff in the previous round affected change in departure time. To construct this

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

figure, we focused only on the group of participants traveling by car over two consecutive rounds. Since the experiment consisted of 40 rounds, excess payoffs were computed for rounds 1 to 39, and changes in departure time were calculated as the departure time in round t minus the departure time in round $t-1$ for rounds 2 to 40 ($t = 2, \ldots, 40$). Thus, each change in departure time corresponds to the excess payoff received in the previous round.

A positive (negative) change in departure time implies departing later (earlier) than the departure time in the previous round. Figure 3 shows that car commuters tended to explore departure times around the chosen one, and the degree of exploration appeared to be larger for a negative excess payoff than for a non-negative excess payoff. However, in the Road Improved condition, car commuters tended to shift toward earlier departure times than the chosen one after having a negative excess payoff. We examined how changes in departure time responded to the previous round's excess payoff by estimating the following linear mixed-effects model:[16]

$$\begin{aligned}\Delta \mathrm{DT}_{it} =& \beta_0 + \beta_1 \mathrm{Excess_Payoff}^{+}_{it-1} \\ &+ \beta_2 \mathrm{Railway_Improved} + \beta_3 \mathrm{Road_Improved} \\ &+ \beta_4 \mathrm{Excess_Payoff}^{+}_{it-1} \times \mathrm{Railway_Improved} \\ &+ \beta_5 \mathrm{Excess_Payoff}^{+}_{it-1} \times \mathrm{Road_Improved} \\ &+ r_i + \epsilon \end{aligned} \tag{3}$$

where $\Delta \mathrm{DT}_{it}$ is player i's departure time in round t minus her departure time in round $t-1$, and $\mathrm{Excess_Payoff}^{+}_{it-1}$ is a dummy variable taking the value 1 if player i's excess payoff in round $t-1$ is non-negative.

The estimation results in the second column of Table 10 indicate that car commuters adjusted their departure times in response to past payoff outcomes. After receiving a negative excess payoff, car commuters in the Baseline and Railway Improved conditions tended to depart earlier in the subsequent round. This temporal adjustment was even more pronounced in the Road Improved condition: car commuters exhibited a stronger tendency to depart earlier. In contrast, after receiving a non-negative excess payoff, car commuters generally departed later, although these temporal adjustments were relatively modest in all three conditions and did not differ significantly between the Baseline and Railway Improved conditions.

Result 6. *Car commuters adjusted their departure times based on their past payoff outcomes, departed earlier after a negative payoff, especially in the Road Improved condition, and modestly later after a non-negative payoff in all conditions.*

How did outcomes in the previous round influence modal choice in the subsequent round? If a commuter receives a lower payoff from traveling by car than the sure payoff from traveling by train, they may be more inclined to switch to the train in the next round. To examine whether the probability of choosing the train depended on the excess payoff in the previous round, we estimated the following

[16] As in Model (2), we excluded the session-specific random effect, $r_{j[i]}$, from Models (3) and (4) because of zero variance and no improvement in model fit, as confirmed by likelihood ratio tests. In Model (3), the participant-level random intercept was also estimated to have zero variant. However, we retained the participant-level random effect in both models to account for the within-subject correlation resulting from repeated decisions made by participants across rounds. Whether to include this effect or not does not affect the estimation results.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

Table 10. Estimation results for Models (3) and (4)

	Dependent variable	
	Model (3) : ΔDT	Model (4): P^{Train}
(Intercept)	-2.257*** (0.392)	-0.661*** (0.191)
Excess_Payoff^{+}	3.902*** (0.478)	-1.469*** (0.137)
Railway_Improved	0.178 (0.614)	0.663* (0.276)
Road_Improved	-3.375*** (0.545)	-0.889** (0.276)
Excess_Payoff^{+} × Railway_Improved	-0.413 (0.743)	-0.076 (0.198)
Excess_Payoff^{+} × Road_Improved	3.472*** (0.646)	-0.444* (0.209)
AIC	36342.795	4382.286
BIC	36394.992	4429.292
Log likelihood	-18163.398	-2184.143
Number of observations	5037	6095
Number of participants	224	237
Variance: Participant (Intercept)	0.000	2.145
Variance: Residual	79.428	

***$p < 0.001$; **$p < 0.01$; *$p < 0.05$. Standard errors are in parentheses.

generalized linear mixed-effects model:

$$\begin{aligned} P_{it}^{\text{Train}} = & f(\beta_0 + \beta_1 \text{Excess_Payoff}_{it-1}^{+} \\ & + \beta_2 \text{Railway_Improved} + \beta_3 \text{Road_Improved} \\ & + \beta_4 \text{Excess_Payoff}_{it-1}^{+} \times \text{Railway_Improved} \\ & + \beta_5 \text{Excess_Payoff}_{it-1}^{+} \times \text{Road_Improved} \\ & + r_i) \end{aligned} \tag{4}$$

where P_{it}^{Train} is player i's probability of traveling by train in round t, and $f(\cdot)$ denotes the standard logistic function. The third column of Table 10 summarizes the estimation results for Model (4). The coefficient of $\text{Excess_Payoff}_{it-1}^{+}$ is negative and highly significant, indicating that a non-negative excess payoff reduced the probability of traveling by train in the subsequent round. In the Road Improved conditions, this effect was even more pronounced: the interaction between $\text{Excess_Payoff}_{it-1}^{+}$ and the Road Improved dummy was significantly negative, suggesting that car commuters who earned a non-negative excess payoff in the previous round were even less likely to switch to the train.

Result 7. *Car commuters were less likely to switch to the train in the subsequent round after receiving a non-negative excess payoff. This effect was especially strong in the Road Improved condition.*

6.2. Reinforcement learning model

To specify our reinforcement learning model, we follow the notations of Camerer and Ho (1999). Let $\pi_i(s_i(t), s_{-i}(t))$ be commuter i's payoff in round t where $s_i(t)$ and $s_{-i}(t)$ are commuter i's strategy and the strategies chosen by all but i in round t, respectively. Let π^* be the equilibrium payoff.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

Assumption 1. Probabilistic Rule: Commuters select their strategy probabilistically based on accumulated attractions. Attraction is the perceived desirability of choosing a particular strategy. We denote commuter i's attraction to strategy j in round t by $A_i^j(t)$ and their probability of choosing strategy j in round t by $P_i^j(t)$. Then, each commuter selects strategy j in round t based on the following probabilistic choice rule:

$$P_i^j(t) = \frac{A_i^j(t-1)}{\sum_{k\in T} A_i^k(t-1)}.$$

Assumption 2. Initial Attractions: To model a situation in which commuters have no prior bias for specific strategies, initial attractions (i.e., attractions at $t=0$) are assumed to be uniform over all strategies: for any $j \in T$,

$$A_i^j(0) = \frac{Q}{|T|},$$

where Q is the strength of initial attraction and $|T|$ is the number of strategies (60 departure times plus the train, i.e., 61 strategies in total). This assumption ensures that learning emerges purely from reinforcement.

Assumption 3. Updating the Rule of Attractions: At the end of each round, attractions are updated according to the following rule:

$$A_i^j(t) = A_i^j(t-1) + R_i^j(t),$$

where $R_i^j(t)$ denotes the reinforcement that commuter i attaches to strategy j in round t. The reinforcement function of our model incorporates spillover reinforcement ("experimentation"), meaning that when commuters select departure times and receive payoffs, they not only reinforce the chosen strategies but also adjust the attractiveness of nearby strategies.[17] This assumption is particularly important in experimental games with a large number of strategies such as ours (61 strategies) because, as Erev and Roth (1998) argue, "players will not (quickly) become locked in to one choice in exclusion of all others" (p.863).

To determine the form of the reinforcement function, we utilized the behavior observed in our experiment. Figure 3 illustrates how commuters adjusted their departure times in response to deviations between their payoffs and the equilibrium payoff in the previous round. There are two discernible patterns that emerged consistently across all conditions. First, after experiencing a non-negative excess payoff, commuters tended to explore departure times close to their previous choices. Second, after incurring a negative excess payoff, commuters exhibited a tendency to depart earlier, and this tendency was even stronger in the Road Improved condition. In response to these observed patterns, we assume that the reinforcement function takes different forms depending on the following three cases:

(a) $s_i(t-1) \in \{1, 2, \ldots, 60\}$ and $\pi_i(s_i(t-1), s_{-i}(t-1)) \geq \pi^*$

Commuter i who traveled by car in the previous round received a higher payoff than the equilibrium payoff (i.e., a non-negative excess payoff). The reinforcement commuter i attaches to strategy j in round t is given by

$$R_i^j(t) = \begin{cases} \pi_i(s_i(t-1), s_{-i}(t-1)) \cdot e^{-g\left|j - s_i(t-1)\right|} & \text{for } j \in \{1, 2, \ldots, 60\} \\ 0 & \text{for } j = \{\text{train}\} \end{cases}$$

where g determines the sensitivity to deviations, which controls how much commuters explore different departure times. This reinforcement function ensures that commuter i attaches the

[17] Erev and Roth (1998) extend their one-parameter reinforcement learning model by adding two psychological assumptions, "forgetting" and "experimentation," the latter of which we employ in our model.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

maximum reinforcement to the chosen departure time, and the reinforcement spills over to departure times near the chosen one with diminishing intensity.

(b) $s_i(t-1) \in \{1, 2, \ldots, 60\}$ and $\pi_i(s_i(t-1), s_{-i}(t-1)) < \pi^*$

Commuter i who traveled by car in the previous round received a lower payoff than the equilibrium payoff (i.e., a negative excess payoff). The reinforcement function is given by

$$R_i^j(t) = \begin{cases} \left[\pi^* - \pi_i(s_i(t-1), s_{-i}(t-1))\right] \cdot e^{-g\left|j - s_i(t-1)\right|} \cdot I_{j<s_i(t-1)} & \text{for } j \in \{1, 2, \ldots, 60\} \\ \pi^* - \pi_i(s_i(t-1), s_{-i}(t-1)) & \text{for } j = \{\text{train}\} \end{cases}$$

where the term $\pi^* - \pi_i(s_i(t-1), s_{-i}(t-1))$ is the foregone payoff commuter i could have gained if she had chosen the train option, and $I_{j<s_i(t-1)}$ is an indicator function that takes one if departure time j is earlier than $s_i(t-1)$. Commuter i attaches the foregone payoff to the train option and also reinforces earlier departure times with diminishing intensity, in an attempt to avoid congestion.

(c) $s_i(t-1) = \{\text{train}\}$

In this case, commuter i traveled by train in the previous round. They reinforce only the train option by adding the equilibrium payoff:

$$R_i^j(t) = \begin{cases} 0 & \text{for } j \in \{1, 2, \ldots, 60\} \\ \pi^* & \text{for } j = \{\text{train}\} \end{cases}$$

6.3. Estimation and simulation

To estimate the two model parameters, Q and g, we employed the maximum likelihood estimation method. The estimation was perfomed separately for each experimental condition. Although the model was fit to the full data from rounds 1-40, the negative log-likelihood function was computed based only on the strategies observed in rounds 31-40. The parameter set that minimized the negative log-likelihood function was then identified for each condition: $Q = 610.32$ and $g = 1.30$ for the Baseline condition, $Q = 550.77$ and $g = 1.32$ for the Railway Improved condition, and $Q = 949.57$ and $g = 1.22$ for the Road Improved condition.

Following the estimation, the model was simulated for 100 iterations.[18] In each iteration, 20 simulated commuters played 40 rounds, and then we constructed a single cumulative distribution by aggregating their strategies across commuters during rounds 31-40. Figure 4 displays the 100 simulated cumulative relative frequency distributions of strategies alongside the predicted and observed distributions separately for each condition. Our learning model broadly reflects the observed patterns of excess road usage in the Baseline and Railway Improved conditions and delayed departures in the Road Improved condition.

Table 9 offers a more detailed evaluation of model performance by comparing the simulated and observed means of travel cost and their components. Although the model approximates the observed mean travel cost in each condition, it does not fully replicate the underlying behavioral patterns of car commuters. The fit is less accurate at the component level: in particular, the model tends to overestimate travel time costs and underestimate early arrival costs in the Baseline and Railway Improved conditions. Nevertheless, the model performs relatively well in the Road Improved condition, where it closely reproduces both travel cost and its components.

The fact that our model fits the Road Improved condition better than the other two conditions highlights that behavioral patterns varied meaningfully across conditions. Given these differences, it may be overly ambitious to expect a single, simple learning model to account for all observed behaviors equally well. Our goal was not to develop a general, comprehensive model of learning across all

[18]This simulation was performed using MATLAB (R2024b). The same random seed was used across all three conditions.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

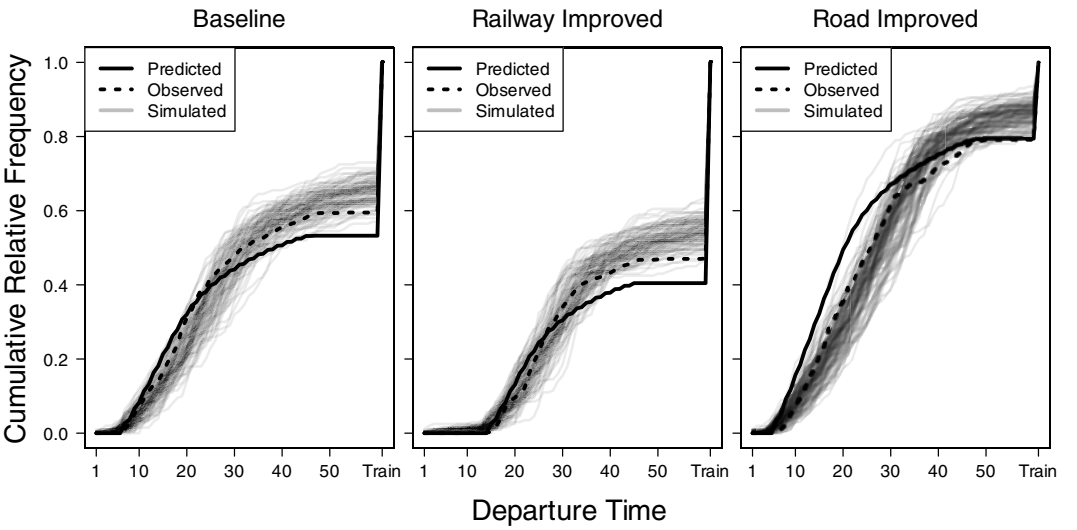


Figure 4. One hundred simulated cumulative relative frequency distributions (rounds 31-40)

environments, but rather to explore whether a parsimonious reinforcement learning model grounded in the experimental findings could replicate the key behavioral patterns observed in our experiment.

7. Conclusion

This study considers a novel scenario that combines elements of the D-T paradox with elements of queuing theory in the morning commute. Commuters face a choice of traveling by private car or train from a common origin to a common destination. Car commuters are also required to choose their departure times. They make these decisions independently and anonymously. Car travel is risky because of congestion, whereas train travel is not (non-congestible). Car commuters face a trade-off between the cost of travel time on the congested road and the cost of schedule delay at the destination. The schedule delay cost is due to the uncertainty of arrival time at the destination; commuters are charged a linearly increasing cost for arriving too early or too late. Similarly, the travel time cost for car commuters is linear in the time they wait before entering the bottleneck plus the fixed time of service. In contrast, the train departs from the origin and arrives at the destination exactly on time, and train commuters are only charged a fixed travel cost.

We considered three conditions, Baseline, Railway Improved, and Road Improved, that differed in either the duration of service time at the bottleneck or the train's travel time. For each condition, we numerically derived a symmetric mixed-strategy equilibrium in this richer strategic environment that integrates both mode and timing choices. According to the D-T hypothesis, improving the train by shortening travel time should decrease mean travel cost, whereas improving the road by reducing service time should have no effect. We experimentally tested the D-T hypothesis in a controlled laboratory experiment with a between-subjects design. Following a period of learning, participants' decisions in all three conditions approached, but did not fully match, equilibrium play.

Our study makes two main contributions to the literature. First, we provide a numerical characterization of equilibrium in a setting in which commuters must make both mode and timing decisions, offering a new environment to examine the D-T paradox. Second, we show experimentally that, although railway improvement reduces mean travel cost as predicted, road improvement also lowers mean travel cost, contrary to the paradox. These results identify a boundary condition for the D-T paradox: it may fail to emerge in environments where people face multi-dimensional decisions and

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

complex externalities that unfold over time. While this paradox may hold in simpler, cognitively less demanding decision settings, its predictive power could weaken when people must coordinate both their mode and timing decisions in the presence of intertemporal congestion effects.

The deviations from equilibrium were systematic and most evident in car commuters' departure times. In the Road Improved condition, participants delayed their departures and were better able to avoid queuing, resulting in a significant reduction in travel time costs. Our reinforcement learning model captured some of those efficiency-enhancing patterns; although the model performed relatively well in the Road Improved condition, it did not fully replicate the underlying patterns of behavior in the Baseline and Railway Improved conditions.

One possible interpretation of the observed deviation from equilibrium is that the added complexity of our environment, which makes coordination difficult, leads to more stochastic behavior. As shown in Figure 2, the observed distributions of departure times tend to deviate from equilibrium in the direction of the 45-degree line (not shown in the figure), which corresponds to random behavior. This shift toward randomness may increase efficiency in the Road Improved condition, as the observed distributions move closer to the socially optimal distribution. However, in the Baseline and Railway Improved conditions, the same pattern of deviation may reduce efficiency, as the observed distributions diverge from the social optimum. These results suggest that randomness is not universally beneficial; although it may mitigate congestion and improve outcomes in certain settings, it can also undermine efficiency when it pushes behavior away from socially optimal patterns. Future research could explore models that account for both payoff-sensitive choices and random variation.

Our scenario complements previous models of the morning commute proposed in the transportation literature in which commuters choose among multiple routes or modes. The major difference is that, in our model, only car commuters face uncertainty regarding their arrival times. Although the assumption that group members are symmetric aids analytical tractability, its empirical validity is questionable. Real-world commuters are known to differ in their attitudes toward risk; some may prefer the certainty of the train, whereas others may value the flexibility of choosing their own departure times and internalize schedule delay costs differently. In this context, future research could incorporate such individual differences into commuter behavior models.

In summary, we find that groups of symmetric commuters systematically deviated from equilibrium predictions, although their choices gradually approached equilibrium over time. This study highlights that complexity in commuters' decision environments can weaken the emergence of the paradoxical outcomes predicted by theory. The D-T paradox holds under some conditions but tends to weaken when individuals must coordinate both mode and timing decisions under congestion that unfolds dynamically over time.

Supplementary material. The supplementary material for this article can be found at https://doi.org/10.1017/eec.2025.10025.

Data availability statement. The replication materials for this study are available at https://doi.org/10.17605/OSF.IO/R678H.

Acknowledgements. We would like to thank guest editor Professor Ido Erev and an anonymous referee for their careful review of the paper and for their constructive comments and suggestions. We would also like to thank the participants of the 2023 Asia-Pacific Economic Science Association Meeting in Seoul for their valuable feedback on an earlier version of this paper.

Finally, we would like to express our heartfelt appreciation and pay tribute to our esteemed co-author and beloved dissertation advisor, Prof. Amnon Rapoport, who passed away before this paper was submitted. Therefore, we take full responsibility for any remaining errors. As his students and friends, we owe him an immeasurable debt of gratitude for his unwavering support, invaluable guidance, and lasting friendships. In loving memory and with profound gratitude, we dedicate this work to his enduring legacy.

Funding. This study was financially supported by the Max Planck Society.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

Conflict of Interest. The authors declare that they have no conflict of interest.

References

Afimeimounga, H., Solomon, W., & Ziedins, I. (2005). The downs-thomson paradox: Existence, uniqueness and stability of user equilibria. *Queueing Systems*, *49*(3), 321–334.

Anderson, L. R., Holt, C. A. & Reiley, D. (Eds.). (2007). Congestion pricing and welfare: An entry experiment. In T. L. Cherry, S. Kroll, & J. F. Shogren (Eds.), *Environmental Economics, Experimental Methods*. Routledge.

Arnott, R., de Palma, A., & Lindsey, R. (1990). Economics of a bottleneck. *Journal of Urban Economics*, *27*(1), 111–130.

Arnott, R., de Palma, A., & Lindsey, R. (1993). A structural model of peak-period congestion: A traffic bottleneck with elastic demand. *The American Economic Review*, *83*(1), 161–179.

Arnott, R., & Small, K. (1994). The economics of traffic congestion. *American Scientist*, *82*(5), 446–455.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48.

Braess, D. (1968). Über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung*, *12*(1), 258–268.

Calvert, B. (1997). The downs-thomson effect in a markov process. *Probability in the Engineering and Informational Sciences*, *11*(3), 327–340.

Camerer, C., & Ho, T. -H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, *67*(4), 827–874.

Dechenaux, E., Mago, S. D., & Razzolini, L. (2014). Traffic congestion: An experimental study of the downs-thomson paradox. *Experimental Economics*, *17*(3), 461–487.

Denant-Boèmont, L., & Hammiche, S. (2010). Flexibility of transport choice in a real-option setting: An experimental case study. *Journal of Intelligent Transportation Systems*, *14*(3), 140–153.

Ding, C., & Song, S. (2012). Traffic paradoxes and economic solutions. *Journal of Urban Management*, *1*(1), 63–76.

Downs, A. (1962). The law of peak-hour expressway congestion. *Traffic Quarterly*, *16*(3), 393–409.

Erev, I., & Rapoport, A. (1998). Coordination, "Magic," and reinforcement learning in a market entry game. *Games and Economic Behavior*, *23*(2), 146–175.

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *The American Economic Review*, *88*(4), 848–881.

Horowitz, J. L. (1984). The stability of stochastic equilibrium in a two-link transportation network. *Transportation Research Part B: Methodological*, *18*(1), 13–28.

Knight, F. H. (1924). Some fallacies in the interpretation of social cost. *The Quarterly Journal of Economics*, *38*(4), 582–606.

Koutsoupias, E. & Papadimitriou, C. (Eds.). (1999). Worst-case equilibria. In C. Meinel, & S. Tison (Eds.), *Stacs 99* (pp. 404–413). Berlin Heidelberg, Springer.

Leifeld, P. (2013). Texreg: Conversion of statistical model output in R to LATEX and HTML tables. *Journal of Statistical Software*, *55*(1), 1–24.

Mak, V., & Rapoport, A. (2013). The price of anarchy in social dilemmas: Traditional research paradigms and new network applications. *Organizational Behavior and Human Decision Processes*, *120*(2), 142–153.

Mogridge, M. J. (1990). Planning for optimum urban efficiency: The relationship between congestion on the roads and public transport. *Transportation Planning Systems*, *1*(1), 11–19.

Mogridge, M. J., Holden, D., Bird, J., & Terzis, G. (1987). The downs/thomson paradox and the transportation planning process. *International Journal of Transport Economics / Rivista Internazionale di Economia dei trasporti*, *14*(3), 283–311.

Morgan, J., Orzen, H., & Sefton, M. (2009). Network architecture and traffic flows: Experiments on the Pigou–Knight–Downs and braess paradoxes. *Games and Economic Behavior*, *66*(1), 348–372.

Murchland, J. D. (1970). Braess's paradox of traffic flow. *Transportation Research*, *4*(4), 391–394.

Papadimitriou, C. (2001). Algorithms, games, and the internet. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, STOC'01, Hersonissos, Greece, pp. 749–753. New York, NY, USA: Association for Computing Machinery.

Pigou, A. (1920). *The Economics of Welfare*. Number v. 1 in The Economics of Welfare. Macmillan and Company.

Pohlert, T. (2023). trend: Non-parametric trend tests and change-point detection.

Rapoport, A. (1995). Individual strategies in a market entry game. *Group Decision and Negotiation*, *4*(2), 117–133.

Rapoport, A., Kugler, T., Dugar, S., & Gisches, E. J. (2009). Choice of routes in congested traffic networks: Experimental tests of the braess paradox. *Games and Economic Behavior*, *65*(2), 538–571.

Rapoport, A., & Mak, V. (2018). Strategic interactions in transportation networks. In K. Donohue, E. Katok, & S. Leider (Eds.), *The Handbook of Behavioral Operations*. pp. 557–586. John Wiley & Sons, Ltd.

Rapoport, A., Seale, D. A., Erev, I, & Sundali, J. A. (1998). Equilibrium play in large group market entry games. *Management Science*, *44*(1), 119–141.

Selten, R., Chmura, T., Pitz, T., Kube, S., & Schreckenberg, M. (2007). Commuters route choice behaviour. *Games and Economic Behavior*, *58*(2), 394–406.

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press

Small, K. A. (1982). The scheduling of consumer activities: Work trips. *The American Economic Review*, *72*(3), 467–479.
Stein, W. E., Rapoport, A., Seale, D. A., Zhang, H., & Zwick, R. (2007). Batch queues with choice of arrivals: Equilibrium analysis and experimental study. *Games and Economic Behavior*, *59*(2), 345–363.
Sundali, J. A., Rapoport, A., & Seale, D. A. (1995). Coordination in market entry games with symmetric players. *Organizational Behavior and Human Decision Processes*, *64*(2), 203–218.
Tabuchi, T. (1993). Bottleneck congestion and modal split. *Journal of Urban Economics*, *34*(3), 414–431.
Thomson, J. (1978). *Great cities and their traffic*. Penguin Books.
Vickrey, W. S. (1969). Congestion theory and transport investment. *The American Economic Review*, *59*(2), 251–260.
Youn, H., Gastner, M. T., & Jeong, H. (2008). Price of anarchy in transportation networks: Efficiency and optimality control. *Physical Review Letters*, *101*(12), 128701.

Cite this article: Otsubo, H., Gisches, E. J., & Rapoport, A. (2025). The Downs-Thomson Paradox with Endogenously Determined Departure Times. *Experimental Economics*, 28(3), 600–624. https://doi.org/10.1017/eec.2025.10025

https://doi.org/10.1017/eec.2025.10025 Published online by Cambridge University Press