Transit Serving Communities Optimally, Responsively, and Efficiently Center

Final Report - Project M2

January 2023

Multi-Agent Simulation

Dr. Einat Tenenboim | Georgia Institute of Technology

Yufei Xu, Georgia Institute of Technology

Dr. Gregory Erhardt | University of Kentucky

Dr. Srinivas Peeta | Georgia Institute of Technology

Dr. Gregory Macfarlane | Brigham Young University

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
TSCORE-M2	N/A	N/A

4. Title and Subtitle
T-SCORE Project M2: Multi-Agent Simulation
A Model of Ride-Hailing Driver Participation

5. Report Date
March 2023

6. Performing Organization Code
N/A

7. Author(s)
Einat Tenenboim, PhD, https://orcid.org/0000-0001-6018-2182
Yufei Xu
Gregory Macfarlane, PhD, https://orcid.org/0000-0003-3999-7584
Gregory Erhardt, PhD, https://orcid.org/0000-0001-8133-3381
Srinivas Peeta, PhD, https://orcid.org/0000-0002-4146-6793

8. Performing Organization Report No.
N/A

9. Performing Organization Name and Address
Georgia Institute of Technology
School of Civil and Environmental Engineering
790 Atlantic Dr, Atlanta, GA 30332-0355

10. Work Unit No.
N/A

11. Contract or Grant No.
USDOT Grant 69A3552047141

12. Sponsoring Agency Name and Address
U.S. Department of Transportation
Office of the Assistant Secretary for Research and Technology
1200 New Jersey Avenue, SE, Washington, DC 20590

13. Type of Report and Period Covered
Final Report
Aug 2020-Jan 2023

14. Sponsoring Agency Code
USDOT OST-R159667

15. Supplementary Notes
N/A

16. Abstract
The upsurge in ride-hailing services and their rapidly growing demand have recently ignited considerable research interest. However, present research on ride-hailing is mainly focused on the demand for trips, while less attention is given to the supply enabling them, i.e., the drivers. To obtain a comprehensive understanding of ride-hailing demand, a good understanding of driver participation should be achieved. The present study aims to obtain a realistic representation of driver participation that will later be embedded in a multi-agent simulation. To date, ride-hailing research has been hampered by a lack of data, yet here we leverage a unique dataset of Lyft and Uber vehicle traces collected in San Francisco. Using a choice modeling approach, we model driver participation within four steps, estimating: (1) number of shifts on the working day, (2) shift duration, (3) shift start time, and (4) shift start location. Driver type (full-time, part-time, occasional) was found to be a strong determinant of driver participation. Full-time, but not occasional, drivers were found to work both more shifts and longer shifts. A higher number of shifts started in the downtown area (where population and employment are higher) and in higher income areas, while less shifts started in areas of high student density. Based on these modeling results, a driver fleet was generated for a typical weekday. This is one of the first studies to estimate ride-hailing driver behavior, offering insights that can potentially support transit agencies in effectively planning and regulating multi-modal transportation.

17. Key Words	18. Distribution Statement
Ride-Hailing, TNC, Driver Behavior	No restrictions.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	26	N/A

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Disclaimer

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, by a grant from the U.S. Department of Transportation's University Transportation Centers Program. However, the U.S. Government assumes no liability for the contents or use thereof.

Acknowledgments

This study was funded, partially or entirely, by a grant from the Transit – Serving Communities Optimally, Responsively and Efficiently (T-SCORE) center, supported by the U.S. Department of Transportation (USDOT) through the University Transportation Centers program. The authors would like to thank T-SCORE and the USDOT for their support of university-based research in transportation, and especially for the funding provided in support of this project. The research team would like to also acknowledge our partners at the San Francisco County Transportation Authority and the Metropolitan Transportation Commission for providing the data necessary to support this study and for their ongoing engagement with the research team in completing the models.

Publication

A description of the data gathered, analyses performed, and results achieved was presented at the 2022 International Conference on Advanced Systems in Public Transport (CASPT 2022):

> Tenenboim, Einat, Y. Xu, G.D. Erhardt, G.S. Macfarlane, S. Peeta. (2022). A Generative Model of Ride-Hail Driver Shifts: Time, Duration, and Location. *Proceedings of the 2022 International Conference on Advanced Systems in Public Transport* (CASPT 2022), Tel Aviv, Israel, November 2022. https://easychair.org/smart-program/CASPT2022/index.html

We have also prepared a draft manuscript on this topic which we anticipate publishing as an open-access journal article. The expected authors and title are:

> Tenenboim, Einat, Y. Xu, G.D. Erhardt, G.S. Macfarlane, J. Castiglione, D. Cooper, S. Peeta (forthcoming). A Model of Ride-Hailing Driver Participation: Shift Duration, Start Time, and Start Location.

Introduction

The Tier 1 University Transportation Center known as Transit – Serving Communities Optimally Responsively and Efficiently (T-SCORE) was a consortium from 2020 to 2023 led by Georgia Tech (GT) that included research partners at University of Kentucky (UK), Brigham Young University (BYU) and University of Tennessee, Knoxville (UTK). The investigators from each university are:

1. **Georgia Tech:** Dr. Kari Watkins (Center Director, now at University of California, Davis), Dr. Michael Hunter, Dr. Pascal Van Hentenryck, and Dr. Srinivas Peeta
2. **University of Kentucky**: Dr. Gregory Erhardt
3. **Brigham Young University**: Dr. Gregory Macfarlane
4. **University of Tennessee**, Knoxville: Dr. Candace Brakewood, and Dr. Christopher Cherry

The overarching goal of the T-SCORE research center was to define a set strategic visions that will guide public transportation into a sustainable and resilient future, and to equip local planners with the tools needed to translate their chosen vision into their own community. The research approach for the T-SCORE center is shown in Figure 1. The research began with a strategy generation stage, which generated qualitative descriptions of strategic directions that transit agencies and their partners can take for further evaluation. These strategic visions fed into a two-track research assessment that includes a Community Analysis Track (led by Dr. Candace Brakewood at University of Tennessee) and a Multi-Modal Optimization and Simulation (MMOS) track (led by Dr. Greg Erhardt at University of Kentucky). Both of these tracks aimed to identify the potential feasibility, benefits, costs and implications of each strategic vision, such as on-demand transit services or new fare policies. These tracks came together in the final strategy evaluation stage, in which the findings were again considered in the context of expert advice, as shown in Figure 1. More information about the various research activities conducted as part of the UTC Tier 1 center can be found on the T-SCORE website hosted by Georgia Tech: https://tscore.gatech.edu/

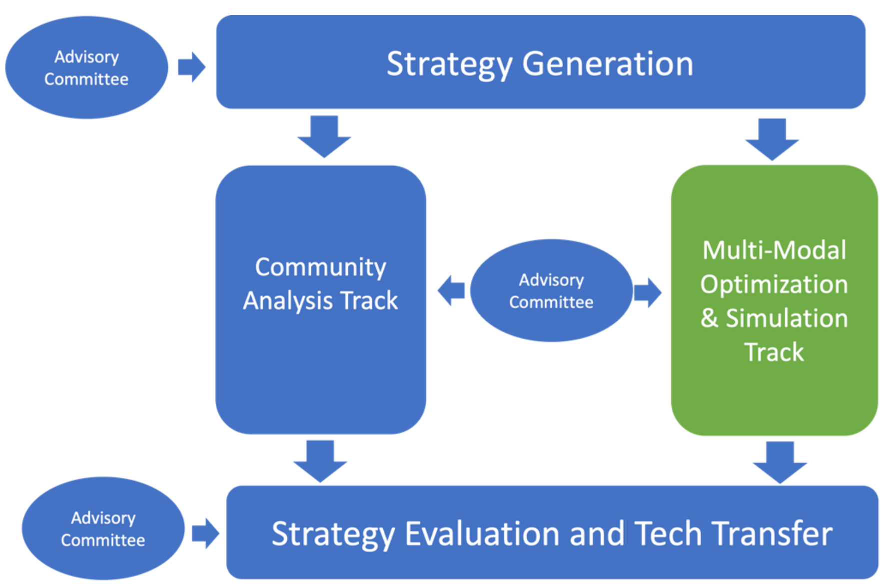

Figure 1: Overarching Research Approach for the T-SCORE Center

The focus of this Final Report is the MMOS research track (highlighted in green in Figure 1). The MMOS research track implemented cutting-edge simulation and optimization techniques to model, forecast, and understand the relationship between traditional and novel public transportation services. The goal of this research was to develop tools that researchers and public agencies can use to develop policies related to new transit modes including agency-organized microtransit services and privately run ride-hail services, and to apply those models to gain insight into both types of services.

The MMOS track's research approach was divided into four separate research projects on these key topics. These four projects (numbered M1-M4) are briefly described in Figure 2 below. These four research projects are strongly related and were completed by an integrated research team consisting of T-SCORE researchers and our partners. For a more holistic description of the relationship between these projects, see Report M4.

M1: Multi-Modal Optimization

- Project M1 focused on developing software tools and optimization models for the design of publicly-operated on-demand multi-modal transit systems.

M2: Multi-Agent Simulation *(This report)*

- Project M2 focused on developing and applying agent-based simulation models to better understand the behavior of private ride-hailing and its effect on the transportation system.

M3: MMOS Integration

- Project M3 focused on better understanding the relationship between private ride-hailing and public transit through improved models and data.

M4: Implementation and Quantitative Evaluation

- Project M4 focused on applying the above tools to realistic scenarios in multiple cities to gain insights into the effective design of transit systems.

Figure 2: MMOS Track Research Projects

Project M2: Multi-Agent Simulation

This Final Report presents the outcomes of MMOS track project M2. Whereas project M1 focused on the tools needed to design publicly-operated transit systems, including on-demand microtransit, project M2 considers their private counterpart: ride-hailing. Microtransit and ride-hailing often look similar from a user perspective—they both allow customers to book on-demand rides served in small vehicles—but they are operated very differently. Microtransit is typically operated to serve the larger public good: providing rides to those who otherwise lack access, reducing vehicle travel, and feeding into the fixed-route transit system. In contrast, ride-hailing is operated to serve the business interests of the ride-hailing companies or drivers. Whereas microtransit is designed specifically to complement fixed-route transit, ride-hailing often competes with public transit, reducing ridership and adding to congestion.

Our team's previous research found that competition with ride-hailing was the biggest contributor to pre-COVID transit ridership declines (Erhardt et al., 2022; Watkins et al., 2022). It was this previous research that led to the formation of the T-SCORE Center, with the aim of helping public transit agencies respond to these changing conditions. In order to respond effectively, transit agencies must understand not only their own operations, but also these new competing modes. Therefore, Project M3 focuses specifically on developing models to predict ride-hail demand, and analyses to measure the effects and applications of ride-hailing. Also, because microtransit and ride-hailing share a parallel user experience, developing the tools to analyze the demand for one enhances our ability to evaluate the other. Specifically, both are well-suited to be modeled using agent-based frameworks because the waiting time for individual users varies based on where the available vehicles are at the point in time when a ride request is made.

The project was focuses specifically on modeling ride-hail driver behavior.

Abstract

The upsurge in ride-hailing services and their rapidly growing demand have recently ignited considerable research interest. However, present research on ride-hailing is mainly focused on the demand for trips, while less attention is given to the supply enabling them, i.e., the drivers. To obtain a comprehensive understanding of ride-hailing demand, a good understanding of driver participation should be achieved. The present study aims to obtain a realistic representation of driver participation that will later be embedded in a multi-agent simulation. To date, ride-hailing research has been hampered by a lack of data, yet here we leverage a unique dataset of Lyft and Uber vehicle traces collected in San Francisco. Using a choice modeling approach, we model driver participation within four steps, estimating: (1) number of shifts on the working day, (2) shift duration, (3) shift start time, and (4) shift start location. Driver type (full-time, part-time, occasional) was found to be a strong determinant of driver participation. Full-time, but not occasional, drivers were found to work both more shifts and longer shifts. A higher number of shifts started in the downtown area (where population and employment are higher) and in higher income areas, while less shifts started in areas of high student density. Based on these modeling results, a driver fleet was generated for a typical weekday. This is one of the first studies to estimate ride-hailing driver behavior, offering insights that can potentially support transit agencies in effectively planning and regulating multi-modal transportation.

1. Introduction

Ride-hailing (or ride-sourcing) services have revolutionized urban transportation systems, showing rapid growth in ridership in recent years, both in and outside the US. Today, transportation network companies (TNCs), such as Uber and Lyft, operate in more than 250 urban areas in the US and 700 globally (Erhardt et al., 2022). In many metropolitan areas, ride-hailing has a significant market share of urban trips and has become an integral part of travelers' mode choice set (Graehler et al., 2019; Habib, 2019). In San Francisco, CA, for example, 15% of all intra-city vehicle trips in 2016 were made by TNCs (Millard-Ball et al., 2022; Roy et al., 2020).

Ride-hailing entails many benefits to the user and the transportation system. Mainly, it improves rider comfort and security and increases mobility for low income/car-free households and for people with disabilities (Tirachini, 2020). Ride-hailing also has a positive effect on parking availability and requirements, especially in busy urban areas where parking demand exceeds the supply (Khavarian-Garmsir et al., 2021; Millard-Ball et al., 2022). However, over time it became apparent that ride-hailing might also have negative consequences on the transportation system. The upsurge in ride-hailing services and the increasing demand have ignited considerable research interest in these. One major line of research seeks to understand whether ride-hailing services are competing with transit services. Ample evidence reveals that ride-hailing has indeed shifted a considerable number of trips away from transit (Erhardt et al., 2022; Jin et al. 2019; Schaller, 2021). However, there is also evidence indicating that ride-hailing can complement transit, offering travelers improved first mile/last mile access and filling-in gaps of scheduled bus/rail services (Habib, 2019; Hall et al. 2018; Khavarian-Garmsir et al., 2021). Another major line of research seeks to determine the role of ride-hailing in urban congestion. Recent studies have shown that ride-hailing significantly increases vehicle miles traveled (VMT) and consequently contributes to congestion (Henao & Marshall, 2019; Roy et al., 2020; Schaller, 2021; Tirachini, 2020) with a clear environmental impact (Khavarian-Garmsir et al., 2021; Ward et al., 2021). Indeed, several studies have established a link between deadheading (empty) trips performed by TNCs and congestion (Cramer and Krueger, 2016; Erhardt et al., 2022; Nair et al., 2020; Roy et al., 2020). In both lines of research, however, the investigation is trip-focused, rather than driver-focused. In other words, present research on ride-hailing is largely motivated by the demand for trips while less attention is given to the supply enabling them. However, to better understand and predict ride-hailing demand, one must also study ride-hailing supply, among other factors.

1.1 Driver participation

Obtaining a good understanding of the demand for trips and a comprehensive picture or travelers' mode choice is not likely without first achieving a fundamental understanding of driver participation in ride-hailing services. The operating force behind ride-hailing services is the individual drivers who make up TNC fleet. TNCs' performance is determined first and foremost by the collective choices of drivers who are, in essence, private fleet providers (Ashkrof et al., 2022). Ride-hailing drivers are free to decide when to work, how long to work, and whether to accept or decline ride requests offered to them at any given moment. Thus, individual driver participation choices can influence system performance, including level of service (a product of waiting time) and environmental effects (a product of idle time). Driver participation may refer to who the drivers are, when and where they choose to work, and the strategies they employ to make various decisions towards maximizing their own profit and/or utility. Estimation and prediction of ride-hailing driver participation is critical for generating reliable forecasts of TNC operation. Such a systematic understanding of ride-hailing supply is essential for effectively regulating, planning, and managing urban multi-modal transportation (Ghaffar et al., 2020). To date, only a few studies have examined ride-hailing supply (e.g., Ashkrof et al., 2020; Ashkrof et al., 2022; de Ruijter et al., 2022; Millard-Ball et al., 2022).

One of the few studies which focused on ride-hailing drivers was Millard-Ball et al. (2022), who presented an analysis of drivers' activities in-between paid trips in San

Francisco. They showed that ride-hailing drivers mostly use the time between trips to reposition themselves to high-demand locations. The second most frequent activity in-between trips was cruising while waiting for a trip request with no specific purpose. In this study, Millard-Ball et al. showed that such passenger-less trips have a significant role in urban congestion, as well as implications for the environment (see also Ward et al., 2021). Further, they demonstrated how the investigation of driver choices can shed light on the overall impacts of ride-hailing.

Research on ride-hailing participation has been limited not only by the focus on trips rather than drivers, but also by the lack of data. Only limited amount of empirical research has been performed on ride-hailing thus far, given that TNCs have not been subject to data sharing requirements until recently (Cooper et al., 2018; Ghaffar et al., 2020). Some of the fruitful insights on ride-hailing travel comes from work performed in major cities like Chicago (Ghaffar et al., 2020), New York City (Jin et al., 2019), Toronto (Habib, 2019), and San Francisco (Cooper et al., 2018; Roy et al., 2020). Major insights from these studies include: (i) a high concentration of ride-hailing trips is typically observed in city centers and other busy areas (Cooper et al., 2018), and (ii) two distinct increases in ride-hailing activity are apparent on weekdays during morning and evening peak periods (Gehrke et al., 2019). One can assume that with the availability of TNC data more insightful ride-hailing research is forthcoming.

1.2 The present study

In the present study, we leverage a unique dataset to investigate ride-hailing driver behavior. The dataset includes vehicle traces scraped from Uber and Lyft application programming interfaces (APIs) operating in San Francisco. This study is part of a project aiming to develop a multi-modal optimization model with a multi-agent simulation. Thus, the objective of the present study is to obtain a realistic representation of ride-hailing driver participation, which will then be embedded in this simulation. Using a choice modeling approach, we model ride-hailing driver participation throughout a 24-hour workday in four steps, each of which includes a set of models and is aimed at estimating a different aspect of driver participation. In the first step, we estimate a model for predicting the number of driver shifts on a typical working day (i.e., how much). In the second and third steps, we model driver shift(s) duration(s) and start time(s), respectively (i.e., for how long and when). And in the fourth and final step, we model driver shift(s) start location (i.e., where). These estimations will then be used to generate a driver fleet that will be used in the multi-modal simulation.

2. Method

2.1 Dataset

In this study, we used a unique dataset of Uber and Lyft vehicle traces from San Francisco. This dataset includes data scraped from the APIs of Uber and Lyft during a 31-day period in November and December 2016. Given our aim of modeling driver participation on a typical weekday, we used Monday through Thursday data only, excluding holidays (19 days; henceforth 'weekday data'). For this period, Uber data includes 3,960,728 trips, and Lyft data includes 1,344,319 trips. Using a choice modeling approach, we modeled ride-hailing driver participation throughout a 24-hour workday.

We encountered two main challenges using this dataset for modeling. First, Uber vehicle traces do not include unique driver identifiers (Cooper et al., 2018). With the absence of unique driver identifiers, we can only study independent trip characteristics, not driver participation. To address this challenge, we used Lyft data only to model driver participation, and then scaled it up to also account for Uber trips. Second, the dataset lacks driver socio-demographic characteristics, impeding our ability to account for variability in driver participation. A preliminary data exploration revealed that drivers differ considerably in terms of the amount of time they work. To account for some of the driver heterogeneity observed, we constructed a driver type variable, corresponding to the number of weekly hours worked by the driver. This led to three driver type categories: occasional drivers (working less than 5 hours a week; 7,081 drivers), part-time drivers (working 5-35 hours a week; 7,771 drivers), and full-time drivers (working over 35 hours a week; 818 drivers). In total, the dataset includes 15,670 Lyft drivers working on weekdays.

2.2 Driver shifts

Preliminary data exploration revealed that many drivers work more than once a day. In some cases, the time gap between one “shift” and a subsequent “shift” is negligible and may result from a short break for meal/errands or from accepting a trip through a another TNC app. Thus, we defined a *shift* as a continuous working duration that does not include breaks longer than 60 minutes. Based on this definition, Lyft drivers in the dataset performed 117,579 shifts on weekdays.

We differentiated between different shifts of one driver on a single working day based on shift length. We considered the longest shift on the working day to be their primary shift, assuming that the longest shift is the first determined by the driver. Respectively, we considered the second longest shift as the secondary shift, and all remaining shifts on the working day as tertiary shifts (only 3.3% of drivers worked four or more shifts on a working day). Note that this shift categorization may not correspond to the chronological order of shifts, as the tertiary shift, for example, may be the first shift performed on the working day.

2.3 Discrete choice modeling

We used multinomial logit models to account for driver participation in four steps targeting distinct features of their participation. The first step was set up to establish the number of shifts the driver is working on the simulated day. Steps 2-4 each included a set of models, as shown in Fig. 1. The second step focused on shift duration and included a set of three models estimating shift duration for the primary, secondary, and tertiary shifts. The third step focused on the time of day in which the shift started, and like the second step, it included a set of three models. Finally, the fourth step focused on the location, or the traffic analysis zone (TAZ), in which the shift started. This step included a set of three models estimating driver start location for the first/second/third shift of the working day, given that chorological shift order was critical in this step, as explained in section 3.4.

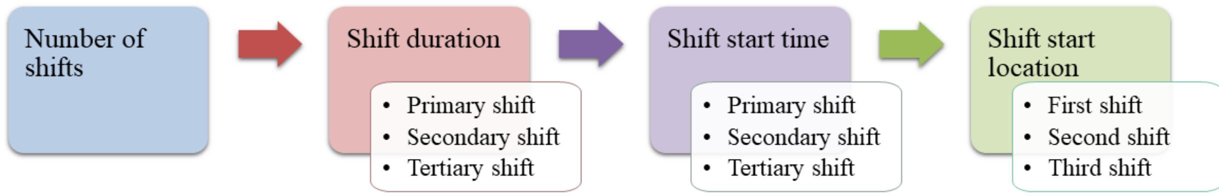

Fig. 1. Modeling driver participation in four steps.

3. Results

3.1 Number of shifts modeling

In the first modeling step, we estimated the number of shifts the drivers worked on the simulated day (i.e., typical weekday). Four choice alternatives were defined: 0, 1, 2, 3+ shifts. The 0 shifts alternative (i.e., no shifts on the working day) was set as the baseline in this estimation and was therefore null. The utility functions for the other alternatives are shown in Eq. 1. The driver type variable served as the sole explanatory variable in this step and was included as two 9out of three) dummy variables corresponding to full-time and occasional drivers.

(Eq. 1) $U_{1/2/3}$ = constant + β_1 * Full-time + β_2 * Occasional

Modeling results are shown in Table 1. The positive coefficients for full-time drivers indicate that these drivers are likely to work more shifts during the working day. The negative coefficients for occasional drivers, on the other hand, indicate that these drivers are less likely to do so. All coefficients are significant (p<0.01). Overall, model fit shows a good result, with an adjusted R^2 of 0.449, implying that driver type is a strong predictor of drivers' number of shifts choice.

Table 1. Number of shifts modeling results including β coefficients (and t-statistics in parentheses).

Utility function	No shifts	1 shift	2 shifts	3+ shifts
Constant	- (fixed)	-0.635 (22.10)	-1.58 (38.64)	-3.34 (36.56)
Full-time driver*	- (fixed)	1.26 (11.48)	1.95 (16.43)	2.42 (13.06)
Occasional driver*	- (fixed)	-1.35 (27.05)	-2.65 (21.09)	-4.07 (6.95)
Number of observations	12,535			
Adjusted R^2	0.449			
Initial log-likelihood	-17,377			
Final log-likelihood	-9,567			

* Dummy variable

3.2 Shifts duration modeling

In the second modeling step, we estimated the duration of driver shifts on a typical working day. Although drivers are the ones to determine shift duration, they must comply with TNC regulations according to which shifts cannot exceed 12 hours. We estimated three separate models to account for drivers' primary, secondary, and tertiary shift durations (see section 2.2). For the primary shift, choice alternatives were defined as 0-1, 1-2, and so forth, with the last alternative being 10+ hours (less than 2% of the shifts were longer than 11 hours). For the secondary and tertiary shifts, choice alternatives were similarly defined, with the last alternative being 4+ hours (less than 1.5% of secondary shifts and 0% of tertiary shifts were longer than 5 hours). The utility functions for the other alternatives are shown in Eq. 2. The utility function for 0-1 shift duration alternative served as baseline in all models and was therefore equal to zero. Explanatory variables included driver type (i.e., two dummy variables corresponding to full-time and occasional drivers), number of additional shifts in the working day (see section 3.1), and previous shift duration (for secondary and tertiary shift estimation).

(Eq. 2) $U_{1\text{-}2/\ldots/10+}$ = constant + β_1 * Full-time + β_2 * Occasional + β_3 * AddiShifts

Modeling results for primary, secondary, and tertiary shift duration are shown in Tables 2, 3, and 4, respectively. Note that in the primary shift model we combined the occasional coefficients for 8-9, 9-10, and 10+ alternatives, given that the coefficients were close in value and not significantly different. Similar coefficient combinations were performed in the secondary and tertiary models. As we expected, the positive coefficients for the full-time variable (in all 3 models) and their increasing values (in the primary and secondary models) indicate that these drivers have a higher utility from working longer shifts. Occasional drivers show the opposite trend, with negatively increasing coefficient values, indicating a lower utility from working longer shifts.

The coefficients of number of additional shifts were largely positive, indicating that drivers who have additional shift/s on the working day have a higher utility of working longer shifts (yet no longer than 8 hours, as can be seen in Table 2). This finding is likely to be related to the driver type variable, capturing drivers who generally work more, i.e., have more shifts and work longer shifts (probably because it is their main occupation), whereas other drivers work less. Note that the occasional and number of additional shifts variables were not significant in the tertiary shift model (most likely due to the lower number of observations) and were therefore left out.

In the secondary and tertiary shift models, we included explanatory variables corresponding to the duration of the previous shift, i.e., the primary and the secondary shifts, respectively. These variables largely yielded positive coefficients, indicating that the longer the previous shift was (or is about to be), the longer the current shift. This finding is another indication of the inherent difference between drivers who work more/less.

To summarize, driver type variable is once again a strong predictor of driver participation. Full-time drivers are more likely to work longer shifts, whereas occasional drivers are less likely to do so. Further, drivers who work more shifts are likely to also work longer shifts.

Table 2. Primary shift duration model results including β coefficients (and t-statistics in parentheses).

Utility function	0-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10	10+
Constant	- (fixed)	0.043 (2.36)	-0.012 (0.63)	-0.185 (9.67)	-0.437 (21.30)	-0.772 (33.96)	-1.11 (43.04)	-1.50 (49.84)	-1.93 (53.41)	-2.41 (53.35)	-2.89 (51.11)
Full-time driver*	- (fixed)	0.006 (0.08)	0.535 (7.38)	0.861 (12.09)	1.37 (19.45)	1.77 (24.84)	2.07 (28.12)	2.39 (30.98)	2.68 (32.54)	2.91 (31.72)	3.14 (30.72)
Occasional driver*	- (fixed)	-0.784 (28.50)	-1.40 (43.76)	-2.00 (48.90)	-2.39 (44.87)	-2.91 (37.09)	-3.05 (30.33)	-3.17 (24.21)	-3.71 (24.40)		
Number of additional shifts (AddiShifts)	- (fixed)	0.979 (30.38)	1.22 (38.24)	1.29 (39.72)	1.18 (34.85)	1.07 (29.55)	0.766 (18.50)	0.484 (9.85)	0.160 (2.60)	-0.283 (3.28)	-0.464 (4.28)
Number of observations	68,181										
Adjusted R^2	0.184										
Initial log-likelihood	-163,490										
Final log-likelihood	-133,426										

* Dummy variable

Table 3. Secondary shift duration model results including β coefficients (and t-statistics in parentheses).

<table>
<tr><th>Utility function</th><th>0-1</th><th>1-2</th><th>2-3</th><th>3-4</th><th>4+</th></tr>
<tr><td>Constant</td><td>-
(fixed)</td><td>-0.517
(12.26)</td><td>-0.914
(15.74)</td><td>-1.54
(18.45)</td><td>-2.30
(22.16)</td></tr>
<tr><td>Full-time driver*</td><td>-
(fixed)</td><td>0.346
(6.57)</td><td>0.629
(11.35)</td><td>1.02
(16.07)</td><td>1.35
(15.65)</td></tr>
<tr><td>Occasional driver*</td><td>-
(fixed)</td><td>-0.609
(8.28)</td><td colspan="3">-0.904
(8.42)</td></tr>
<tr><td>Number of additional shifts</td><td>-
(fixed)</td><td colspan="2">0.696
(14.64)</td><td colspan="2">0.601
(11.56)</td></tr>
<tr><td>Primary shift duration</td><td>-
(fixed)</td><td>0.127
(11.37)</td><td>0.168
(12.39)</td><td colspan="2">0.188
(11.17)</td></tr>
<tr><td>Number of observations</td><td colspan="5">20,609</td></tr>
<tr><td>Adjusted R^2</td><td colspan="5">0.042</td></tr>
<tr><td>Initial log-likelihood</td><td colspan="5">-24,926</td></tr>
<tr><td>Final log-likelihood</td><td colspan="5">-23,873</td></tr>
</table>

* Dummy variable

Table 4. Tertiary shift duration model results including β coefficients (and t-statistics in parentheses).

<table>
<tr><th>Utility function</th><th>0-1</th><th>1-2</th><th>2-3</th><th>3-4</th><th>4+</th></tr>
<tr><td>Constant</td><td>-
(fixed)</td><td>-1.32
(10.01)</td><td>-2.53
(13.29)</td><td>-4.36
(2.66)</td><td>-5.63
(2.56)</td></tr>
<tr><td>Full-time driver*</td><td>-
(fixed)</td><td>0.353
(3.70)</td><td colspan="3">0.314
(1.66)</td></tr>
<tr><td>Primary shift duration</td><td>-
(fixed)</td><td>-0.0891
(2.90)</td><td colspan="3">0.0194
(0.11)</td></tr>
<tr><td>Secondary shift duration</td><td>-
(fixed)</td><td>0.427
(7.26)</td><td colspan="3">0.430
(0.93)</td></tr>
<tr><td>Number of observations</td><td colspan="5">3,180</td></tr>
<tr><td>Adjusted R^2</td><td colspan="5">0.163</td></tr>
<tr><td>Initial log-likelihood</td><td colspan="5">-2,283</td></tr>
<tr><td>Final log-likelihood</td><td colspan="5">-1,900</td></tr>
</table>

* Dummy variable

3.3 Shift start time modeling

In the third step, we estimated the time of day at which drivers start their shifts. We included 24 choice alternatives corresponding to the time of day. As in the second modeling step, we separately estimated models for the primary, secondary, and tertiary shifts. The utility functions for the other alternatives are shown in Eq. 3. The utility function for the 1AM alternative served as baseline in all models and was therefore equal to zero. Explanatory variables included driver time (as two dummy variables corresponding to full-time and part-time drivers), number of shifts on the working day (see section 3.1), and shift duration (see section 3.2).

(Eq. 3) $U_{2/\ldots/24}$ = constant + β_1 * Full-time + β_2 * Part-time + β_3 * NumShifts + β_4 * Duration

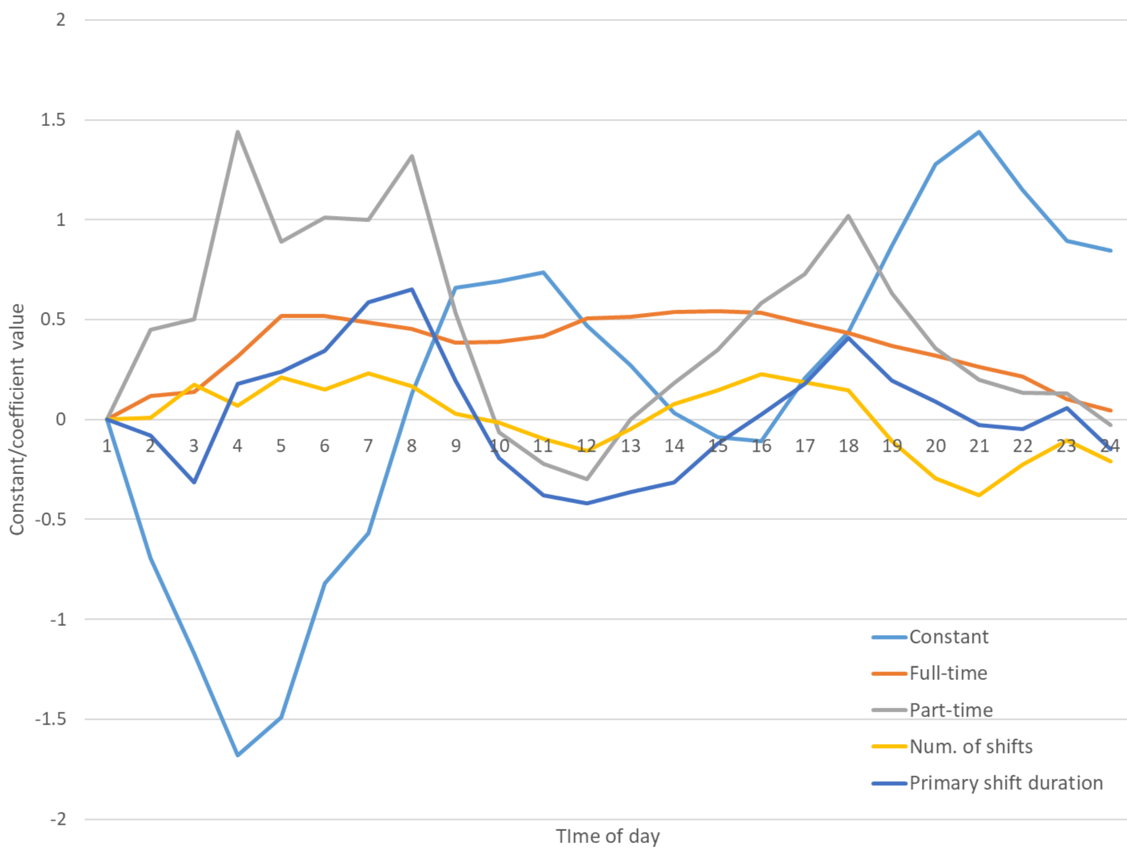

Fig. 1. Primary shift start time modeling results including constant- and full-time, part-time, number of shifts and shift duration coefficient- values, as a function of time of day (1-24).

The results for the primary shift start time model are depicted in Fig. 1. As seen in the figure, full-time drivers have a higher utility in starting shifts between 5AM and 6PM, displaying a somewhat unified distribution in between these times of day. In contrast, part-time drivers have a higher utility in starting shifts between 2-9AM and 4-8PM. These findings are reasonable given that TNC driving is the primary/sole occupation of full-time, but not part-time, drivers. The number of shifts and the primary shift duration variables yielded significant and positive coefficients around the AM/PM peak periods, indicating that the longer the drivers work (in terms of number of shifts/hours) the greater the utility they have in starting shifts at around the peak periods. This finding is likely to arise from full-time drivers who both work longer shifts and prefer to work during the peak periods when demand is high (especially because they are relatively flexible with little to no work-related obligations). The results for the secondary and tertiary shift start time models yielded similar insights. See Tables A1 to A3 in Appendix A for full lists of constant and coefficient values.

3.4 Shift start location modeling

In the fourth and final step, we estimated where drivers start their shifts, i.e., in which of San Francisco's 981 traffic analysis zones (TAZs). In contract to the shift duration and start time models, in this step we estimated separate models for the first, second, and third shifts (and beyond) of the working day, corresponding to the chronical order of the shifts. This decision was motivated by our interest in the relation between the end location of one shift and start location of the next shift. The technical approach for this step included two separate stages: (1) distributing shifts *between* San Francisco's 14 districts, using a choice modeling approach, and (2) distributing shifts *within* the districts, proportionally to the population of each TAZ within the district.

In the first, choice modeling, stage, 14 choice alternatives were indicated, corresponding to San Francisco's 14 districts. The utility functions are shown in Eq. 4.

(Eq. 4) $U_{1/\ldots/14} = \log(Pop + \beta_1 * Emp) + \beta_2 * PopEmpDensity + \beta_3 * LowIncome + \beta_4 * HighIncome + \beta_5 * College$

Three main components were incorporated in the utility functions: (a) *a size term*, for measuring the quantity of opportunities to users, including coefficients for population and employment counts, (b) *a utility term*, for qualitative measures related to district characteristics, including population and employment, income, and college enrolment densities (which were all calculated within a radius of 0.5 mile and were divided by the land area), and (c) for the second and third shifts, a dummy variable denoting a shared border between the district in which the shift started and the district in which the previous shift ended. Note that the same coefficients were estimated for the variables in all utility functions.

Table 8. Shift start location modeling results for first, second, and third shifts in San Francisco's 14 districts (t-statistics appear in parentheses).

	First Shift	Second Shift	Third Shift
Shared border*	-	2.560 (9.65)	1.1 (3.98)
Size term (count)			
Population	1.000 (fixed)	1.000 (fixed)	1.000 (fixed)
Employment	4.00 (16.2)	2.32 (8.8)	7.61 (3.51)
Utility term (density)			
Population and employment density	0.0173 (57.1)	0.0204 (36.6)	0.0103 (7.76)
Low-income*	-6.37E-05 (34.4)	-0.000071 (21.3)	-3.76E-05 (4.36)
High-income*	0.000213 (58.6)	0.000223 (34.4)	0.000108 (7.34)
College enrollment	-0.000243 (60.6)	-0.000255 (35.175)	-0.000157 (8.72)
Number of observations	69,480	21,124	3,459
Adjusted R^2	0.133	0.104	0.137
Initial log-likelihood	-159,002	-53,913	-8,654
Final log-likelihood	-158,940	-48,308	-7,464

* Dummy variable

Modeling results of the start location models for the first, second and third shifts are shown in Table 8. For the size terms, population count was fixed whereas employment count was estimated. The estimation yielded significant and positive coefficients for these variables in all models. The positive coefficients for population and employment density, as part of the utility terms, are consistent with this finding, suggesting that more driver shifts start in the downtown districts of San Francisco, where population and employment are higher. The negative low-income and the positive high-income coefficients (all significant) suggest that a greater number of shifts start in districts of higher, as opposed to lower, income. Further, the significant negative coefficients for college enrolment for all models indicate that a lower number of shifts start in districts of higher college student density. This outcome is presumably due to students typically living in close proximity to campus, commuting daily by walking or micro-mobility modes rather than by ride-hailing services. Finally, the shared border variable yielded positive coefficients in the second and third shift models, confirming a link between the end location of one shift and the start location of the next shift[1].

Based on these models, we distributed driver first, second, and third shifts between San Francisco's 14 districts. Then, within each district, we further distributed the start location of each shift between the TAZs of that district. This was performed in proportion to the population count of each TAZ. Overall, the start location of all shifts was distributed between San Francisco's 981 TAZs. The right column of Figure 2 contains heat maps reflecting driver start location according to the modeling results, whereas the left column contains the corresponding heat maps based on the observed start location, for comparison. The TAZs' colors in the figure reflect the number of shifts starting in each TAZ on the working day, such that darker colors reflect more shifts. Two main conclusions can be drawn from the heat maps in this figure. First, as also indicated by the coefficients, a higher number of shifts start in San Francisco's downtown TAZs. This finding is consistent with findings from other studies (Cooper et al., 2018; Erhardt et al., 2022; Ghaffar et al., 2020; Schaller 2021). For example, Ghaffar et al. (2020) obtained evidence indicating that restaurants, population and employment density are positively associated with high ride-hailing demand. Second, the modeled and observed data generated similar geographical patterns suggesting an overall good fit for the data.

[1] A variable capturing the distance in miles between these two locations was initially included in the models, yet it was later omitted due to negligible contribution to model estimation.

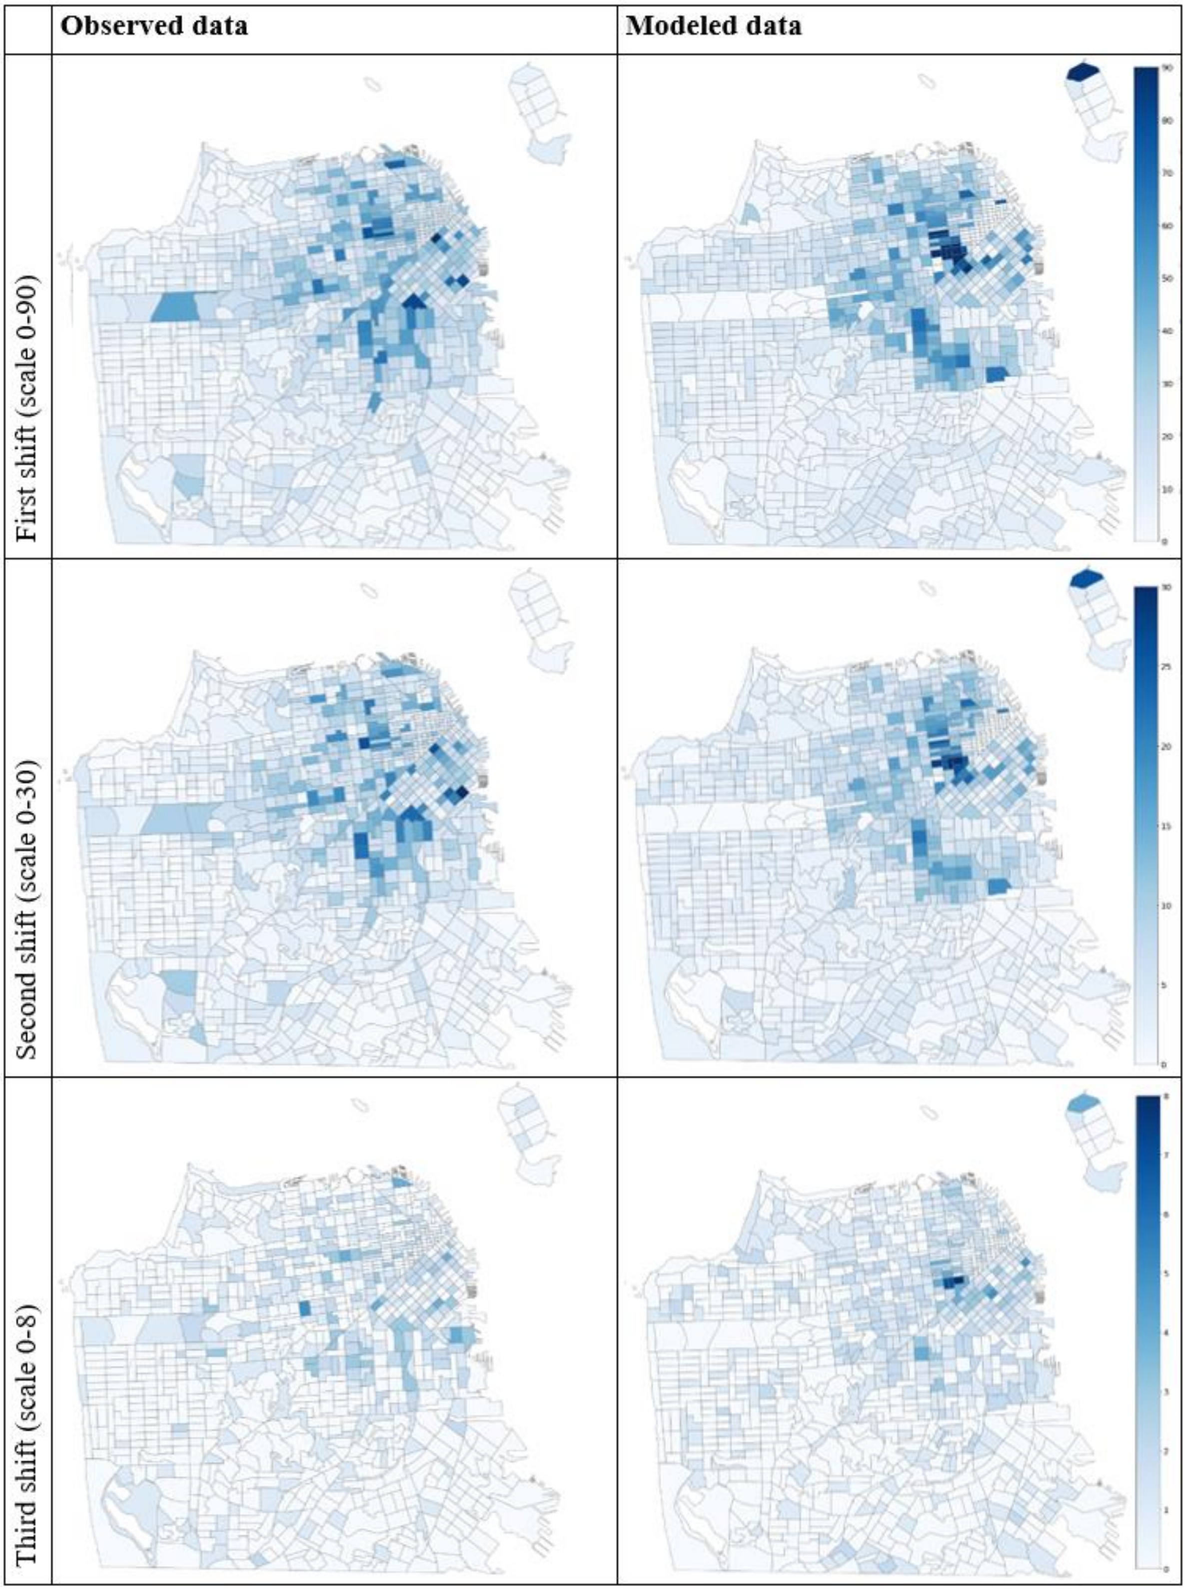

Fig. 2. Heat maps including driver start location modeling results for San Francisco's 981 TAZs – observed vs. modeled data, for first, second, and third shifts.

3.5 Driver fleet generation

Based on the models reported in sections 3.1-3.4, a ride-hailing driver fleet was generated for a typical weekday. As aforementioned, this fleet was generated to realistically represent TNC driver participation in a multi-modal multi-agent simulation of San Francisco. However, due to the absence of unique driver identifiers in the Uber data, the modeling results are based on Lyft data only. To generate the driver fleet for San Francisco, we first scaled up the number of drivers to also account for Uber's fleet. Then, we filtered out drivers who in the first modeling step (number of shifts estimation) chose not to work on the simulated day, i.e., selected the choice alternative of 0 shifts. Figure 3 depicts the steps in this process.

Fig. 3. Driver fleet generation process including scaling up the number of drivers and filtering out drivers not working on the simulated day.

As can be seen in the figure, modeling results are based on data capturing the behavior of 11,732 Lyft drivers, yet after scaling up the fleet to also include Uber drivers, we came up with a total of 46,224 drivers for the San Francisco fleet, 14,679 of which are estimated to work on the simulated day.

Due to gaps between modeled and observed data in shift start time, we calibrated the modeled data based on the observed data. Specifically, we calculated new constants for the utility functions of the start time models, as shown in Eq. 5.

$$\text{(Eq. 5) } Constant_{new} = Constant_{present} + \ln(\%X_{obs} / \%X_{mod})$$

where X_{obs} is the observed number of shifts and X_{mod} is the modeled number of shifts starting at each time alternative. Model calibration included three iterations, after which we obtained the results shown in Fig. 4.

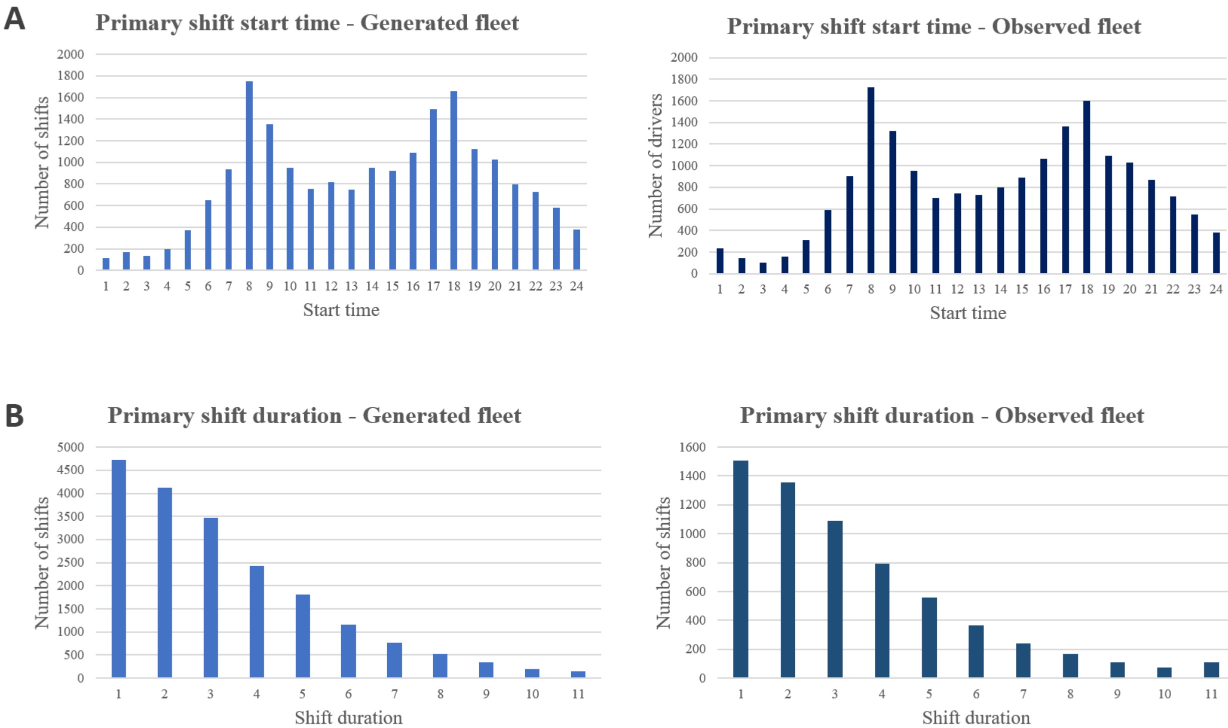

Fig. 4. Generated fleet (left) and observed fleet (right) comparisons in terms of shift start time (panel A) and shift duration (panel B) for primary shift.

4. Discussion

While most contemporary research on ride-hailing focuses on the demand for trips, in the present study we focused on the supply side. To this aim, we took on a data-centric approach and investigated driver participation in San Francisco based on a unique dataset of Lyft and Uber vehicle traces. We estimated a set of four discrete choice models to study how drivers determine if, when, where, how many times, and for how long they will work on a given day. We then generated a ride-hailing driver fleet to obtain a realistic representation of ride-hailing driver participation that could later be embedded in a multi-modal multi-agent simulation of San Francisco transportation. Differentiation between driver shifts on the working day was made based on shift length (primary, secondary, and tertiary shifts), and proved useful in data analysis and modeling.

A distinction between full-time, part-time, and occasional drivers (based on the number of weekly working hours) has significantly contributed to the explanation of driver participation (see also Ashkrof et al., 2020; Ramezani et al., 2022). Full-time, but not occasional, drivers were found to not only work more shifts per day, but to also work longer shifts. Further, full-time drivers were found to begin their shifts between 5AM and 6PM with a close-to-unified distribution in between. In contrast, part-time drivers were found to begin their shifts between 2-9AM and 4-8PM, most likely because driving is not their primary/sole occupation. Overall, drivers were found to prefer starting their shifts during the AM and PM peak periods. Consistent with previous findings (Cooper et al., 2018; Erhardt et al., 2022; Ghaffar et al., 2020; Schaller 2021), TNC drivers were found to frequently prioritize the downtown area as a starting point for their shifts, where population and employment densities are higher and consequently demand is greater (see also Millard-Ball et al., 2022). Moreover, drivers were also more likely to start their shifts in higher income districts, and less likely to start them in areas of high student density. Finally, as we expected, drivers tended to start their second or third shifts of the working day in close geographical proximity to the end location of the previous shift. Overall, the models showed a high level of goodness of fit.

4.1 Contribution

The present study is one of the first to estimate ride-hailing driver behavior and participation, given the scarcity of objective data. The findings of the present study may benefit the generation of reliable forecasts of TNC supply, that could then be pooled together and modeled with ride-hailing demand data, towards improving our understanding of this new travel mode and its impacts on the transportation system. The ride-hailing fleet we generated will be used in a multi-modal simulation of San Francisco transportation, thereby allowing local transit agencies to be nimbler in adapting to dynamic conditions.

The present findings are also likely to promote our understanding of the documented decline in transit ridership over the past decade, in light of empirical evidence linking it to the onset of ride-hailing services (Erhardt et al., 2022; Graehler et al., 2019). As part of the continuous efforts made to account for this negative trend, a comprehensive approach to multi-modal investigation has been gaining support in recent years. According to this approach, the decline in transit ridership should not be investigated independently, considering only transit user perception, trust, and acceptance, but also those of competing modes, such as TNCs. This work may also be beneficial in supporting effective planning and regulating of coordinated on-demand multi-modal transit.

4.2 Future research

Several lines of research naturally evolve from the present study. First, depending on data availability, incorporating socio-demographic variables of drivers (such as age, gender, income level, driving experience, and duration of ride-hailing participation) is expected to significantly improve estimation accuracy and consequently our ability to account for and predict driver participation. In addition, incorporating transit accessibility data in driver's shift start location modeling may generate valuable insights. On the one hand, one could expect ride-hailing demand to be higher where transit accessibility is lower, yet there is evidence that areas with a higher density of bus stops are linked to greater ride-hailing demand (Ghaffar et al., 2020). Thus, driver start location modeling is likely to benefit from the inclusion of this, and other, transit-related variables.

Another interesting line of research relates to the similarities and differences between ride-hailing driver participation characteristics (especially those who are identified as full-time drivers) and that of taxi drivers. Some previous research has looked into this question (e.g., Millard-Ball et al., 2022), yet based on the present findings, further inquiry is warranted. Such an investigation would be fruitful in indicating the extent to which we can employ insights obtained from taxi driver behavior research to study ride-hailing driver behavior. An example for one potentially important insight is the extent of similarity between ride-hailing driver cruising and repositioning behavior and that of taxi drivers, along with the impacts on congestion and the environment (e.g., Yang et al., 2005).

Finally, taking on a complementary approach to study driver decision-making processes, specifically by conducting interviews, focus groups, and surveys with drivers, may reveal important insights that could not be extracted from objective datasets. Some attempts in this direction were previously made, including focus groups (e.g., Ashkrof et al., 2020) and surveys (e.g., Ashkrof et al., 2022), yet these have not entirely been focused on driver participation and behavior.

4.3 Conclusions

In contrast to drivers in other modes of travel (e.g., taxis), TNC drivers naturally exhibit large and inherent heterogeneity. Combined with the lack of driver socio-demographic data, our ability to effectively study and model TNC driver participation is substantially hindered. In the present study, we took this on-going investigation one step further, offering insights on daily participation decisions made by TNC drivers. We showed the advantages of separately investigating different types of drivers to effectively model driver participation. To continue this endeavor, future research could leverage these insights, investigating both fundamental and novel factors associated with ride-hailing driver decision-making processes.

4.4 Recommendations

As ride-hailing and other new modes are very different than traditional private travel because they involve an explicit interaction between the behavior of the driver and the behavior of the traveler. Therefore, we must recognize that traditional models of travel behavior do not capture the full set of behavior we may wish to capture when understanding the transportation system. If we wish to understand the impacts of these new modes, we must consider how these two groups interact. We recommend that travel demand models be updated to reflect this interaction.

References

Ashkrof, P., de Almeida Correia, G. H., Cats, O., & Van Arem, B. (2020). Understanding ride-sourcing drivers' behaviour and preferences: Insights from focus groups analysis. *Research in Transportation Business & Management*, 37, 100516.

Ashkrof, P., de Almeida Correia, G. H., Cats, O., & van Arem, B. (2022). Ride acceptance behaviour of ride-sourcing drivers. *Transportation Research Part C: Emerging Technologies*, 142, 103783.

Chalermpong, S., Kato, H., Thaithatkul, P., Ratanawaraha, A., Fillone, A., Hoang-Tung, N., & Jittrapirom, P. (2022). Ride-hailing applications in Southeast Asia: A literature review. *International Journal of Sustainable Transportation*, 1-21.

Cooper, D., Castiglione, J., Mislove, A., & Wilson, C. (2018). Profiling TNC activity using big data. In Transportation Research Board Annual Meeting, Washington, DC, January.

Cramer, J., & Krueger, A. B. (2016). Disruptive change in the taxi business: The case of Uber. *American Economic Review*, 106(5), 177-182.

de Ruijter, A., Cats, O., Kucharski, R., & van Lint, H. (2022). Evolution of labour supply in ridesourcing. *Transportmetrica B: Transport Dynamics*, 10(1), 599-626.

Erhardt, G. D., Mucci, R. A., Cooper, D., Sana, B., Chen, M., & Castiglione, J. (2022). Do transportation network companies increase or decrease transit ridership? Empirical evidence from San Francisco. *Transportation*, 49(2), 313-342.

Gehrke, S. R., Felix, A., & Reardon, T. G. (2019). Substitution of ride-hailing services for more sustainable travel options in the greater Boston region. *Transportation Research Record*, 2673(1), 438-446.

Ghaffar, A., Mitra, S., & Hyland, M. (2020). Modeling determinants of ridesourcing usage: A census tract-level analysis of Chicago. *Transportation Research Part C: Emerging Technologies*, 119, 102769.

Graehler, M., Mucci, R. A., & Erhardt, G. D. (2019). Understanding the recent transit ridership decline in major US cities: Service cuts or emerging modes. In Transportation Research Board 98th Annual Meeting, Washington, DC, January.

Habib, K. N. (2019). Mode choice modelling for hailable rides: An investigation of the competition of Uber with other modes by using an integrated non-compensatory choice model with probabilistic choice set formation. *Transportation Research Part A: Policy and Practice*, 129, 205-216.

Hall, J. D., Palsson, C., & Price, J. (2018). Is Uber a substitute or complement for public transit?. *Journal of Urban Economics*, 108, 36-50.

Henao, A., & Marshall, W. E. (2019). The impact of ride-hailing on vehicle miles traveled. *Transportation*, 46(6), 2173-2194.

Jin, S. T., Kong, H., & Sui, D. Z. (2019). Uber, public transit, and urban transportation equity: A case study in New York City. *The Professional Geographer*, 71(2), 315-330.

Khavarian-Garmsir, A. R., Sharifi, A., & Hajian Hossein Abadi, M. (2021). The social, economic, and environmental impacts of ridesourcing services: A literature review. *Future transportation*, 1(2), 268-289.

Millard-Ball, A., Liu, L., Hansen, W., Cooper, D., & Castiglione, J. (2022). Where ridehail drivers go between trips. *Transportation*, 1-23.

Nair, G. S., Bhat, C. R., Batur, I., Pendyala, R. M., & Lam, W. H. (2020). A model of deadheading trips and pick-up locations for ride-hailing service vehicles. *Transportation Research Part A: Policy and Practice*, 135, 289-308.

Ramezani, M., Yang, Y., Elmasry, J., & Tang, P. (2022). An empirical study on characteristics of supply in e-hailing markets: A clustering approach. *Transportation Letters*, 1-14.

Roy, S., Cooper, D., Mucci, A., Sana, B., Chen, M., Castiglione, J., & Erhardt, G. D. (2020). Why is traffic congestion getting worse? A decomposition of the contributors to growing congestion in San Francisco-Determining the Role of TNCs. *Case Studies on Transport Policy*, 8(4), 1371-1382.

Schaller, B. (2021). Can sharing a ride make for less traffic? Evidence from Uber and Lyft and implications for cities. *Transport Policy*, 102, 1-10.

Tirachini, A. (2020). Ride-hailing, travel behaviour and sustainable mobility: An international review. *Transportation*, 47(4), 2011-2047.

Ward, J. W., Michalek, J. J., & Samaras, C. (2021). Air pollution, greenhouse gas, and traffic externality benefits and costs of shifting private vehicle travel to ridesourcing services. *Environmental Science & Technology*, 55(19), 13174-13185.

Yang, H., Ye, M., Tang, W. H., & Wong, S. C. (2005). Regulating taxi services in the presence of congestion externality. *Transportation Research Part A: Policy and Practice*, 39(1), 17-40.

Appendix A

Table A1. Start time modeling results for the primary shift.

Utility function	0-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10	10-11	11-12	12-13	13-14	14-15	15-16	16-17	17-18	18-19	19-20	20-21	21-22	22-23	23-24
Constant	- (fixed)	-0.695 (4.16)	-1.17 (6.4)	-1,68 (9.01)	-1.49 (10.11)	-0.818 (6.35)	-0.57 (4.67)	0.132 (1.16)	0.659 (5.79)	0.694 (5.98)	0.738 (6.28)	0.469 (3.9)	0.271 (2.26)	0.0335 (0.28)	-0.0895 (0.76)	-0.11 (0.95)	0.206 (1.81)	0.438 (3.89)	0.871 (7.64)	1.28 (11.18)	1.44 (12.43)	1.15 (9.74)	0.896 (7.37)	0.847 (6.49)
Full-time driver*	- (fixed)	0.119 (2.9)	0.138 (3.05)	0.315 (8.61)	0.52 (17.48)	0.518 (18.07)	0.486 (17.12)	0.452 (16.11)	0.387 (13.65)	0.391 (13.58)	0.418 (14.48)	0.508 (17.73)	0.514 (18.00)	0.539 (18.94)	0.544 (19.23)	0.534 (18.94)	0.481 (17.10)	0.434 (15.46)	0.368 (12.94)	0.322 (11.24)	0.263 (9.06)	0.214 (7.18)	0.102 (3.28)	0.0472 (1.4)
Part-time driver*	- (fixed)	0.449 (1.51)	0.503 (1.63)	1.44 (5.16)	0.891 (3.69)	1.01 (4.44)	1 (4.52)	1.32 (6.12)	0.533 (2.44)	-0.0654 (0.29)	-0.223 (0.99)	-0.299 (-1.33)	0.000198 (0)	0.183 (0.83)	0.349 (1.58)	0.585 (2.68)	0.728 (3.36)	1.02 (4.75)	0.631 (2.88)	0.358 (1.62)	0.2 (0.89)	0.133 (0.58)	0.132 (0.55)	-0.0276 (0.11)
Number of shifts	- (fixed)	0.0076 6 (0.07)	0.173 (1.4)	0.0689 (0.61)	0.211 (2.4)	0.15 (1.84)	0.231 (2.94)	0.166 (2.17)	0.0278 (0.36)	-0.0131 (0.16)	-0.0972 (1.18)	-0.158 (1.89)	-0.0457 (0.56)	0.0774 (0.96)	0.148 (1.87)	0.227 (2.91)	0.189 (2.44)	0.148 (1.93)	-0.108 (1.36)	-0.292 (3.63)	-0.38 (4.59)	-0.226 (2.69)	-0.105 (1.22)	-0.209 (2.19)
Primary shift duration in hours	- (fixed)	-0.082 (0.59)	-0.315 (2.03)	0.181 (1.1)	0.238 (1.87)	0.344 (3.14)	0.587 (5.7)	0.65 (6.82)	0.197 (2.1)	-0.194 (2.03)	-0.38 (3.96)	-0.42 (4.3)	-0.363 (3.7)	-0.314 (3.19)	-0.122 (1.24)	0.0266 (0.27)	0.181 (1.92)	0.409 (4.38)	0.196 (2.1)	0.0891 (0.96)	-0.0253 (0.27)	-0.0492 (0.51)	0.0586 (0.59)	-0.149 (1.44)
Number of observations	69,480																							
Adjusted R^2	0.072																							
Initial log-likelihood	-220811.18																							
Final log-likelihood	-204809.06																							

Table A2. Start time modeling results for the secondary shift.

Utility function	**0-1**	**1-2**	**2-3**	**3-4**	**4-5**	**5-6**	**6-7**	**7-8**	**8-9**	**9-10**	**10-11**	**11-12**	**12-13**	**13-14**	**14-15**	**15-16**	**16-17**	**17-18**	**18-19**	**19-20**	**20-21**	**21-22**	**22-23**	**23-24**
Constant	- (fixed)	1.34 (2.08)	1 (1.38)	2.36 (3.35)	4.14 (7.37)	5.53 (10.3)	6.68 (13)	7.65 (15.5)	8.71 (17.3)	9.08 (17.5)	9.14 (17.2)	9.5 (17.7)	9.6 (17.9)	9.87 (18.3)	9.69 (18)	10.3 (19)	10.3 (19.1)	9.99 (18.7)	9.36 (17.3)	9.2 (16.8)	8.04 (14.5)	7.84 (14.3)	8.67 (15.5)	8.99 (15.3)
Number of shifts	- (fixed)	-0.359 (1.62)	-0.238 (0.96)	-0.053 (0.22)	-0.0055 (0.03)	-0.204 (1.13)	-0.465 (2.71)	-0.459 (2.83)	-0.42 (2.52)	-0.237 (1.35)	-0.237 (1.3)	-0.261 (1.4)	-0.239 (1.29)	-0.34 (1.81)	-0.323 (1.72)	-0.531 (2.81)	-0.59 (3.15)	-0.518 (2.81)	-0.37 (1.96)	-0.365 (1.89)	-0.303 (1.56)	-0.168 (0.88)	-0.424 (2.11)	-0.56 (2.58)
Secondary shift duration in hours	- (fixed)	-0.264 (2.07)	-0.184 (1.27)	-0.062 (0.46)	0.105 (1.01)	0.218 (2.21)	0.252 (2.63)	0.231 (2.46)	-0.161 (1.68)	-0.467 (4.69)	-0.396 (3.87)	-0.411 (3.92)	-0.307 (2.94)	-0.204 (1.96)	-0.01 (0.09)	0.049 (0.48)	0.081 (0.79)	-0.029 (0.28)	-0.288 (2.78)	-0.534 (5.05)	-0.409 (3.84)	-0.49 (4.59)	-0.677 (6.21)	-0.824 (7.22)
Gap between primary shift and secondary shift in hours	- (fixed)	0.0395 (1.66)	0.086 (3.1)	0.275 (8.9)	0.409 (15.69)	0.494 (19.91)	0.493 (21.51)	0.517 (23.65)	0.621 (27.26)	0.77 (31.61)	0.865 (34.44)	0.946 (37.79)	0.984 (39.74)	1.02 (41.16)	1.05 (41.96)	1.09 (43.18)	1.18 (46.32)	1.3 (50.54)	1.34 (51.09)	1.37 (52.05)	1.47 (55.33)	1.48 (55.69)	1.45 (54.19)	1.42 (51.33)
Number of observations	19,428																							
Adjusted R^2	0.242																							
Initial log-likelihood	-54990.5																							
Final log-likelihood	-41615.1																							

Table A3. Start time modeling results for the tertiary shift.

Utility function	0-1	1-2	2-3	3-4	4-5	5-6	6-7	7-8	8-9	9-10	10-11	11-12	12-13	13-14	14-15	15-16	16-17	17-18	18-19	19-20	20-21	21-22	22-23	23-24
Constant	- (fixed)	-0.797 (0.67)	2.71 (1.95)	5.41 (4.28)	7.18 (6.16)	7.37 (6.45)	9.34 (8.36)	10.2 (9.2)	11 (9.83)	12.1 (10.88)	12.9 (11.45)	13.2 (11.73)	13.7 (12.19)	13.5 (11.93)	13.3 (11.72)	13.2 (11.57)	12.7 (11.02)	12.7 (11.1)	12 (10.31)	11.4 (9.66)	8.62 (7.13)	7.86 (6.52)	7.25 (5.99)	7.07 (5.72)
Secondary shift duration in hours	- (fixed)	0.529 (1.99)	0.311 (1.05)	0.349 (1.2)	0.0014 (0)	0.215 (0.78)	0.0987 (0.36)	0.0271 (0.1)	-0.115 (0.39)	0.052 (0.18)	-0.361 (1.2)	-0.237 (0.81)	-0.025 (0.08)	0.0465 (0.15)	0.176 (0.57)	0.349 (1.12)	0.349 (1.08)	0.57 (1.78)	0.243 (0.71)	0.0837 (0.24)	0.721 (2.11)	0.635 (1.87)	0.918 (2.73)	0.578 (1.68)
Tertiary shift duration in hours	- (fixed)	-0.251 (0.58)	-2.2 (2.87)	-0.62 (1.28)	-0.455 (0.99)	-0.07 (0.16)	-0.311 (0.69)	0.209 (0.48)	0.0139 (0.03)	-0.234 (0.52)	-0.507 (1.04)	-0.515 (1.07)	-0.708 (1.47)	-0.494 (1.01)	-0.622 (1.24)	-0.676 (1.32)	-0.77 (1.45)	-0.661 (1.27)	-0.521 (0.96)	-0.67 (1.22)	-0.858 (1.57)	-0.524 (0.97)	-0.826 (1.53)	-1.46 (2.55)
Gap between primary shift and secondary shift in hours	- (fixed)	0.0235 (0.51)	0.127 (2.15)	0.407 (6.46)	0.385 (6.72)	0.394 (7.03)	0.484 (8.75)	0.577 (10.27)	0.55 (10)	0.701 (12.35)	0.765 (13.03)	0.863 (14.32)	1.01 (16.19)	1.14 (17.33)	1.27 (18.53)	1.28 (18.38)	1.46 (20.29)	1.46 (20.37)	1.58 (21.44)	1.63 (21.94)	1.81 (23.55)	1.84 (24.33)	1.85 (24.31)	1.92 (24.71)
Gap between primary secondary and tertiary shift in hours	- (fixed)	-0.011 (0.27)	0.137 (2.75)	0.208 (4.12)	0.309 (6.49)	0.336 (7.19)	0.366 (8.01)	0.396 (8.83)	0.434 (9.56)	0.488 (10.78)	0.531 (11.32)	0.606 (12.61)	0.734 (14.54)	0.869 (15.95)	0.962 (17.05)	0.992 (17.29)	1.1 (18.88)	1.08 (18.7)	1.18 (20.01)	1.25 (20.87)	1.29 (20.91)	1.38 (22.55)	1.41 (23.01)	1.45 (23.08)
Number of observations	2,797																							
Adjusted R^2	0.343																							
Initial log-likelihood	-7328.0																							
Final log-likelihood	-4700.5																							