https://doi.org/10.1016/j.trc.2020.102623

Achieving higher taxi outflows from a drop-off lane: A simulation-based study

Fangyi Yang [a, b], Weihua Gu [b], Michael J. Cassidy [c], Xin Li [b, d], Tiezhu Li [a*]

[a] *School of Transportation, Southeast University, Nanjing, China*

[b] *Department of Electrical Engineering, Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong*

[c] *Department of Civil and Environmental Engineering, University of California, Berkeley, USA*

[d] *College of Transportation Engineering, Dalian Maritime University, Dalian, China*

Abstract

Lanes used by taxis and other shared-ride vehicles at airports and rail terminals are often congested. The present paper examines congestion-mitigating strategies for a special type of lane inside of which taxis are prohibited from overtaking each other while dropping-off patrons. Taxis must therefore often wait in first-in-first-out (FIFO) queues that form in the lane during busy periods. Patrons may be discharged from taxis upon reaching a desired area near the terminal entrance. When wait times grow long, however, some taxis discharge their patrons in advance of that desired area.

The Nanjing South Railway Station in China is selected as a case study. Its FIFO drop-off lane is presently managed by police officers who allow taxis to enter the lane in batched fashion. Inefficiencies are observed because curb space near the upstream and downstream ends of the lane often goes unused.

A microscopic simulation model is developed in-house, and is painstakingly calibrated to data measured in the study site's FIFO lane. Simulation experiments indicate that rescinding the lane's present batching strategy can increase taxi outflow by more than 25%. Further experiments show that even greater gains can be achieved by requiring taxis to discharge patrons when forced by downstream queues to stop a prescribed distance in advance of a desired drop-off area. Further gains were predicted by requiring the lead taxi in each batch to discharge its patron(s) only after travelling a prescribed distance beyond a desired location. The above findings are confirmed for scenarios calibrated with field data collected on two different days, and for hypothetical scenarios with varying input parameters. Roles that technology can play in implementing these new lane-management strategies are discussed. So are their practical implications in light of the present boom in shared-ride services.

Keywords: drop-off lane; FIFO lane; simulation; taxi queues; taxi outflow; control strategies

[*] Corresponding author.
Tel: +8613813850110; e-mail: litiezhu@seu.edu.cn.

1. Introduction

Special-use lanes for dropping-off travelers at or near terminals are common features at airports, train stations, and border crossings (de Neufville and Odoni, 2013). Present focus is on fully-separated, single-lane facilities that are reserved for vehicle drop-offs. These are found in many places in the world, and are especially common at high-speed rail terminals throughout China.

Drop-off vehicles are often made to enter a lane of this sort in batches, meaning in convoys or platoons, and to traverse the single, separated lane in first-in-first-out (FIFO) fashion. The batched outflows achieved in this way depend in part upon the locations where the lead vehicle in each batch stops to discharge its passenger(s). Other vehicles in a batch are forced to stop when their leader does. Whether or not these other vehicles drop-off their own patrons during these forced stops also affect outflows. These drop-off decisions are influenced by: a vehicle's present position relative to a desired drop-off location; the anticipated time still required of the vehicle to reach that location; and the time already spent to that end.

In short, the operation of FIFO drop-off lanes is rather complex. Importantly, this complex operation can produce sizable delays during busy periods. Passengers tend to be especially sensitive to these delays, since they may have planes or trains to catch. Improving the performance of these FIFO lanes thus becomes a worthy objective (e.g., Costa and de Neufville, 2012; Ji et al., 2016). This may be especially true in the coming era of shared-ride services, when more of these lanes will be needed to serve shared-ride vehicles at busy terminals and meeting-points (Aïvodji et al., 2016; Li et al., 2018; Lokhandwala and Cai, 2018).

To date, most of the research in this realm pertains to drop-off areas with passing lanes (e.g. Parizi and Braaksma, 1994; Chang et al., 2000; Chang, 2001), and is therefore not of present interest. Some of those studies relied upon deterministic models (Whitlock and Cleary, 1969; Neufville, 1976; Mandle et al., 1980; 1982; Shapiro, 1996; Ashford et al., 2011). The simplicity of these models is desirable, but the inherent variability in dwell times and drop-off locations are ignored as a result.

Simulation models have often been used to address these variabilities (Costa and de Neufville, 2012; Farhan, 2015). Yet, the models coded to date overlook certain other features of FIFO lanes. Some simulation models, for example, ignore that patrons prefer certain drop-off locations over others (Tillis, 1973; McCabe and Carberry, 1975; Hall, 1977). Other logic has failed to appreciate how patrons can grow impatient and opt to alight from vehicles in advance of a desired location, particularly when vehicle queues grow long (Wang, 1990; Parizi and Braaksma, 1994; Bender and Chang, 1997; Tunasar et al., 1998; Chang et al., 2000; Chang, 2001).

In light of the above, the present work has developed a microscopic simulation model that more faithfully replicates vehicle traffic in a FIFO drop-off lane. When in motion, vehicle movements are governed by the car-following model in Menendez and Daganzo (2007). Distributions of patrons' desired drop-off locations are estimated from data. So are patron tendencies to grow impatient and alight vehicles prior to reaching those desired locations.

The Nanjing South Railway Station (NSR) was selected as a case study. A FIFO drop-off lane at the NSR is reserved for taxi use only. Taxi entries to this lane are batched by policemen who are posted at the scene. Once the simulation model was calibrated to replicate observed conditions, it was used to examine alternative schemes for managing the lane's taxi operations. Removing present-day police controls and allowing taxis to enter the lane at will was found to increase taxi discharge flows by 26-

32%. Two other control strategies produced even greater gains in outflow. One requires that taxis discharge their patrons whenever stopped (by a lead taxi) at a location sufficiently close to the desired one. The other requires that a lead taxi discharge its patrons a prescribed distance beyond the desired location, to free-up desirable curb space for other taxis in the batch.

The NSR case-study site is described in the following section. The logic of our simulation model follows from observations at the site and is presented in section 3. Field data are described and used to calibrate the simulation model, as presented in section 4. The calibrated model is tested in section 5. Parametric study of the aforementioned control strategies is presented in section 6. Practical implications and roles for technology in implementing our ideas are discussed in section 7.

2. Taxi Lane Case Study

The NSR terminal and its taxi drop-off lane are described in section 2.1. Taxi operations in that lane are detailed in section 2.2.

2.1 Site

A photograph of the drop-off area at the NSR terminal is provided in Figure 1. The area's layout is illustrated in Figure 2. Note from the latter figure that the area extends for 240m and consists of 5 lanes. Four of the lanes are open to general traffic, and the fifth is reserved for taxis. The taxi lane is separated from others by means of a physical barrier, such that vehicular overtaking is not possible in the taxi lane; i.e. its operation is FIFO. A painted crosswalk that guides pedestrians to the terminal's entrance and ticketing station is also shown in Figure 2. Crossing pedestrians periodically interrupt taxi flows.[1] In an apparent effort to discourage excessive drop-off numbers near the terminal's entrance, curbside railing is installed between the 90m and 150m marks.

Figure 1. Photo of the drop-off area at the NSR terminal

2.2 Lane Operations

During uncongested periods, taxis enter their FIFO lane free of police controls, such that batch size is limited only by the lane's storage space. During congested periods, a police officer stationed approximately 65m in advance of the drop-off area admits taxis to the FIFO lane in batches. Admissions are offered whenever the lane is nearly emptied of its previous batch. The lead taxi in each batch can choose its drop-off location within the 240m area. Data measured from videos show that almost all lead taxis choose locations spanning the 90m and 170m marks. The police officer

[1] A second crosswalk exists downstream of the first (at the 186m mark). Pedestrian flows in that downstream crosswalk are small, and seldom interrupt taxi flows.

virtually always releases secondary, smaller-sized batches whenever a present batch is stopped by a leader (these stoppages occur whenever leaders discharge their patrons). Admitting secondary batches helps to fill the FIFO lane.

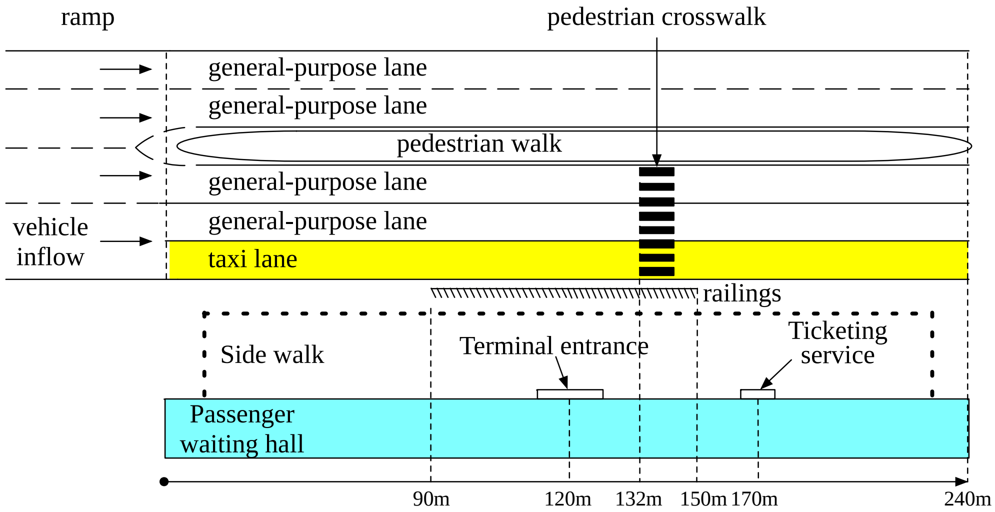

Figure 2. Layout of the drop-off area at the NSR terminal, Nanjing, China

Other taxis in a batch can discharge their patrons while the leader dwells at its drop-off location, or they can wait until advancing closer to a desired spot. A taxi that defers discharging its patron(s) tends to retard outflow from the FIFO lane by virtue of making the deferred stop. Further outflow reductions occur because each fresh batch of taxis is admitted to the lane only after its upstream portion had been empty for some time. Oftentimes even the admission of secondary batches leaves upstream space in the lane unused. Yet further outflow losses occur because drop-offs almost never occur at the downstream-most portion of the FIFO lane, beyond the 170m mark.

3. Simulation Model

Once in the FIFO drop-off lane, taxis move forward as per the car-following model in Menendez and Daganzo (2007). The logic is based on a vehicle's bounded acceleration capabilities, with added considerations for safety and driver/passenger comfort. The model has been a popular choice to simulate vehicle queues owing to its physical realism and parsimony in parameters (e.g., Cassidy et al., 2015a, b). Details of this car-following model are furnished in Appendix A.

Other features of the simulation program are original. These were developed to emulate observations taken at the NSR's taxi drop-off lane, and are described below.

The driver of the lead taxi in a batch selects a drop-off location in the lane as per distributions estimated from field data; see section 4.2. The selection process used by all other taxis to select drop-off locations is modelled as per the state transition diagram in Figure 3. A summary is given below.

Consider a taxi forced by its leader to stop in advance of the desired drop-off location. The taxi drops-off its patron(s) during that stop, if or when the elapsed time at the stop exceeds an underlying limit value, which we term the patron's *patience*. The taxi moves forward once the batch leader enables this. If a patron is still onboard, she alights: (i) at the first forced stop to occur when the time limit (i.e. patron patience) has elapsed; or (ii) upon reaching the desired drop-off location, should that occur first.[2] Distributions used in executing the above logic were estimated from field data, as described in section 4.2.

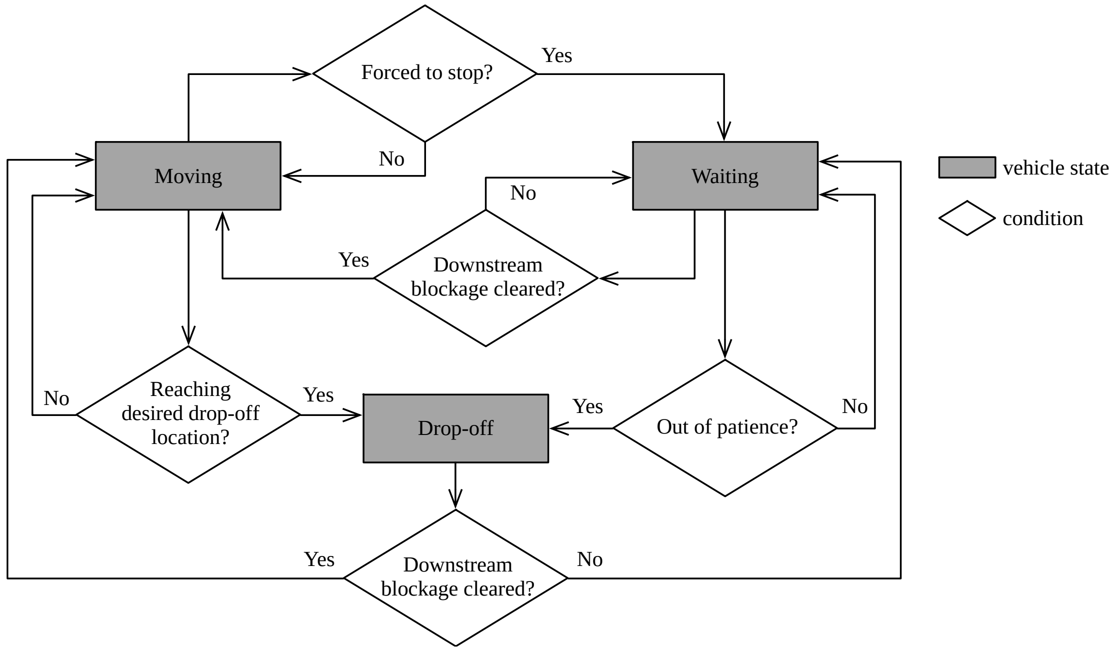

Figure 3. State transition diagram for a vehicle that is not the leading one in a batch

4. Parameter Estimation

Taxi movements were recorded by four video cameras placed in series along the FIFO lane. Recordings were made over 90-min periods (approximately) on the mornings of April 25 and July 13, 2017. Long taxi queues and police batching operations occurred throughout these periods. Trajectories of more than 1,000 taxis were constructed using methods described in Zhang (2000) and Yang et al. (2019).

The trajectories were used to estimate parameters in our simulation model. Estimations were performed in three parts. First, the FIFO lane was partitioned into multiple contiguous segments to reflect the varying locations where taxis were forced (by their leaders) to stop. Second, distributions were estimated for patron patience, as previously described in section 3. The third part of the process entailed the estimation of all other parameters used in our simulation model.

The three parts of the process are described in sections 4.1-4.3. Importantly, estimates from each part were separately generated for the two observation days. This is because data from each day appear to reflect seasonal differences in travel behavior. More will be said on this matter in due course.

4.1 Lane partitioning

The FIFO lane was partitioned with consideration of the varying location where taxis stopped due to stoppages downstream. A taxi's resulting forced wait time is defined as: (i) the elapsed time from when the taxi was forced to stop, until it began to drop-off its patron(s); or (ii) the stop's entire duration, whenever patron drop-off did not occur in that instance.

[2] The logic does not consider the taxi's distance from its desired drop-off location when forced to stop.

These wait times tended to be larger when forced stops were spent toward the FIFO lane's upstream end. This is evident in Figure 4. Its data were collected on April 25 and correspond to taxis' first instances of being forced to stop in the lane. (Second and later instances were less common and usually of shorter duration.) Note how the trend of the best-fit line is downward, though not strictly downward for reasons to be explained momentarily. The trend unveils the longer forced wait times that often occurred upstream, and reflects patron reluctance to alight taxis when still far from the terminal entrance.

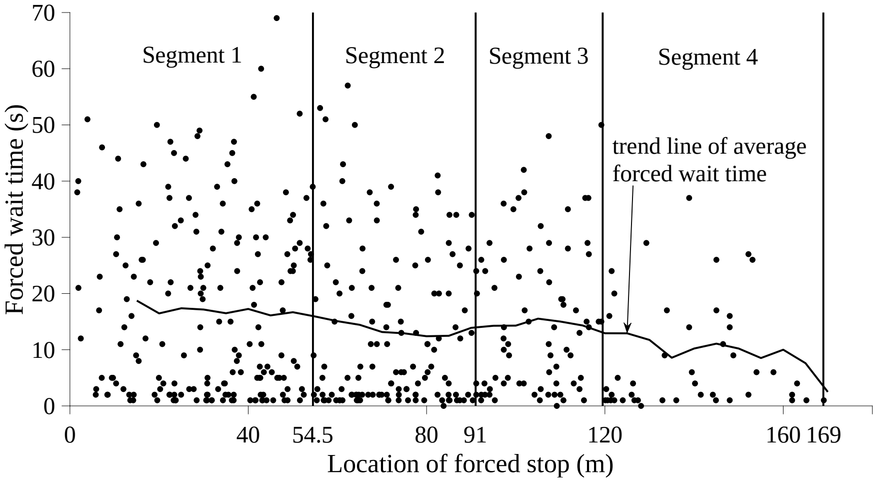

Figure 4. Lane partition for the April 25 dataset ($k = 4$)

The k-means clustering algorithm of Hartigan and Wong (1979) was used to partition the lane into four contiguous segments, also as shown in Figure 4 and detailed in Appendix B.[3] Outcomes of segment-specific analyses are presented for each observation day in Table 1.

Note from the table that the clustering algorithm selected segment lengths that were distinct across days. Further note how average forced wait times were greater in Segment 3 than in upstream Segment 2. This is because the lane's curbside railing coincides with Segment 3 (see again Figure 2), which seems to discourage patrons from alighting there.

Table 1. Four segments and average forced wait times in each segment

		Segment 1	**Segment 2**	**Segment 3**	**Segment 4**
April 25	Range (m)	0-54.5	54.5-91	91-119.5	119.5-169
	Average forced wait time (s)	17.5	13.2	15.5	8.4
July 13	Range (m)	0-38	38-74.5	74.5-110	110-167
	Average forced wait time (s)	23.4	14.9	20.8	12.7

4.2 Distributions of patron patience

While undergoing a forced stop, a patron will alight from her taxi when her patience reaches its end. Distributions of patron patience corresponding to first instances of forced stops were estimated separately for each segment of the FIFO lane and for each observation day. Second and higher

[3] The clustering algorithm was also separately run for cases in which the lane was partitioned into 3 and 5 segments. The R^2 exceeded 0.999 in all cases. A 4-segment partition was ultimately selected because it better aligned with the lane's spatially-varying physical features; e.g. the terminal entrance and curbside railing locations.

instances of forced stops were fewer in number. For this reason, data corresponding to these higher instances were combined across all segments, and a single distribution for these instances was estimated for each day.

All distributions estimated in the above fashion seem to follow a mixed-distribution pattern. The pattern is exemplified for one case via the empirical CDF shown in Figure 5. The function reveals that the patience of a good many patrons was exhausted soon after being forced to stop. In contrast, the patience of the remaining patrons was spread over a wide range.

The patron patience distributions were thus estimated as mixtures of two gamma distributions. Values for all parameters were obtained via maximum likelihood estimation, as described in Appendix C. The estimate for the example case in Figure 5 is shown with a dashed line. Note how the estimate nicely fits the empirical data.

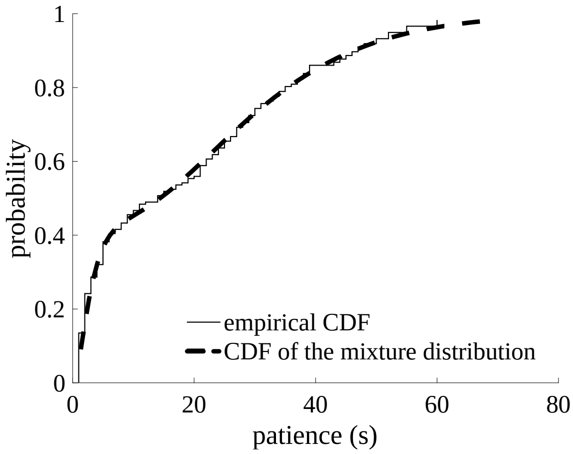

Figure 5. Patron patience distribution for first-instance forced stops in Segment 1 (the April 25 dataset)

4.3 Remaining model parameters

Additional parameters to model the taxi-batching strategy were estimated as described in Appendix D.1. As already noted, taxi motion along the FIFO lane was modeled as per the logic in Menendez and Daganzo (2007), with patron drop-off locations and durations that were fit to empirical distributions; see Appendix D.2.

Finally, a virtual (i.e. simulated) demand-responsive traffic signal was placed at the crosswalk previously shown in Figure 2. Red and green phases were varied across cycles to emulate periods when taxi movements were and were not interrupted by crossing pedestrians. Empirical distributions for these two phases were fit to each day's data.

The above-cited parameter estimates are provided in Table 2 for the two datasets. Values for the July 13 dataset are enclosed in parentheses, if they are different from those for the April 25 dataset. Further details on the parameter estimation can be found in a technical report (Yang et al., 2019).

5. Model Testing

Once calibrated with the input parameters estimated for a given day, model outputs for that day nicely matched those measured in the FIFO lane. The outcomes reported below are averages of 500 simulation runs. The duration of each run equaled the observation period used for its respective day.

Demand for the FIFO lane was always set to 400 taxis/h. This ensured that taxi queues persisted upstream of the lane, consistent with observations on both days. Importantly, the presence of upstream queues guaranteed that resulting outflows were maximum (i.e., capacity) rates.[4]

Table 2. Other parameter values

Parameter		**Value**
Vehicle motion model		
Reaction time		1 s
Jam spacing		7.5 m
Maximum acceleration		2.12 m/s^2 (2.39 m/s^2)
Maximum deceleration		-2.86 m/s^2 (-2.37 m/s^2)
Cruise speed in Segment 1		6.13 m/s (5.91 m/s)
Cruise speed in Segment 2		4.94 m/s (5.11 m/s)
Cruise speed in Segment 3		3.30 m/s (4.20 m/s)
Cruise speed in Segment 4		6.07 m/s (5.89 m/s)
Cruise speed in [169,240] m		5.72 m/s (5.82 m/s)
Initial speed		4.54 m/s (4.48 m/s)
Initial headway		2 s
Drop-off behavior		
Desired drop-off location distribution		An empirical distribution fitted by data
Drop-off duration distribution for lead taxis		An empirical distribution fitted by data
Drop-off duration distribution for other taxis		An empirical distribution fitted by data
Proportion of taxis that drop-off patrons in the lane		0.83 (0.85)[#]
Signal representing the crosswalk		
Red period distribution		An empirical distribution fitted by data
Green period distribution		An empirical distribution fitted by data
Batching control		
L_{m1}	Primary batches	122.4 m (199.4 m)
	Secondary batches	70.2 m (70.4 m)
L_{m2}	Primary batches	66.5 m (72.2 m)
	Secondary batches	15.4 m (27.7 m)
T_m	Primary batches	120 s (93.5 s)
	Secondary batches	14.5 s (16.9 s)
L_{left}	Primary batches	42.3 m (53.4 m)
	Secondary batches	12.3 m (18.3 m)

[#] The remaining vehicles might have dropped off patrons before entering the FIFO lane.

5.1 Outflows

Measured and simulated taxi outflows from the FIFO lane are shown for each day in Table 3. Each day's measured and simulated rates agree to within 6%.[5] Simulation consistently underpredicted outflows, in part because of its use of a traffic signal to describe interruptions by pedestrians; see again section 4.3. In reality, drivers sometimes squeezed their taxis between neighboring groups of crossing pedestrians. This heightened outflows in a way not considered in the model.

[4] Maximizing outflows in turn minimizes delays in the lane.

[5] Two-sided T-tests were performed to determine if simulated travel times are not statistically different from those estimated in the field (Sun and Elefteriadou, 2010; Sun et al., 2013). The p-values (0.09 for the April 25 data and 0.77 for the July 13 data) indicate that simulated and measured travel times are not statistically different at 95% confidence level.

5.2 Travel Times

Displayed in Figures 6a and b are measured and simulated probability density functions of taxi travel times in the FIFO lane. Note again the good fit each day between measured and simulated values.[6]

Table 3. Measured and simulated outflows

Parameter	Dataset	Type	Value	Relative error
Outflow (taxis/h)	April 25	field data	361.5	-5.81%
		simulation	340.5	
	July 13	field data	338.3	-2.16%
		simulation	331	

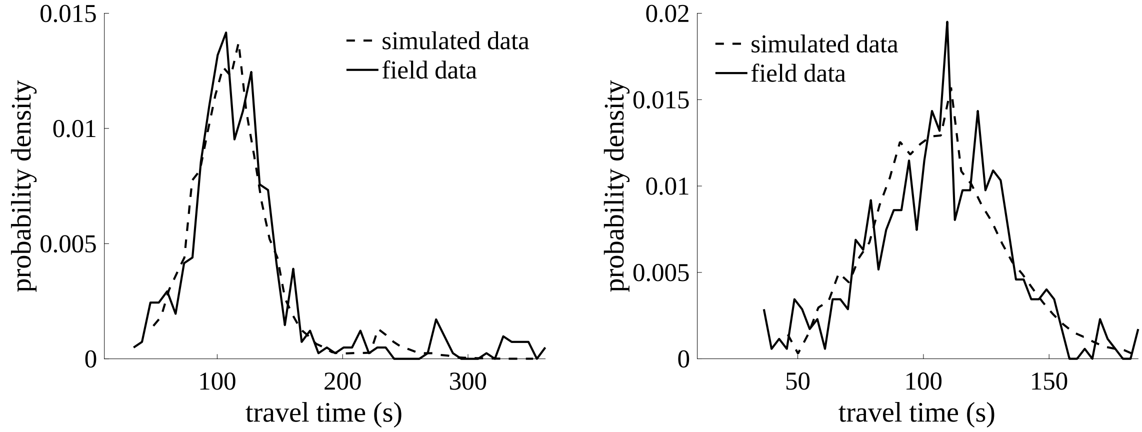

(a) the April 25 dataset (b) the July 13 dataset

Figure 6. Comparison between the PDFs of simulated and measured travel time

5.3 Forced Stops and their Durations

Predicted and measured numbers of forced stops are displayed in Table 4. Each day's predictions match observed tallies to within 7%. The model over-predicted the number of first-instance stops occurring in upstream-most segment 1 by modest amounts. It therefore tended to under-predict first-instance numbers in downstream segments. The table also shows that differences between simulated and measured stop durations were less than 10%.

5.4 Drop-off locations

The PDFs of simulated and measured drop-off locations are shown for each day in Figures 7a and b. The difference between each day's simulated and measured average location is within 3%.

6. Experiments

Having shown that the simulation model can replicate each day's taxi operations, the model was next used to evaluate three alternative management strategies. The three alternatives promote greater curbside utilization in the FIFO lane by dispensing with present-day batching controls. All three are found as a result to produce greater taxi outflows than those presently achieved.

[6] T-tests were again performed for this metric. The high p-values (0.10 for the April 25 data and 0.19 for the July 13 data) again indicate that simulated and measured outflows are not statistically different at 95% confidence level.

Table 4. Measured and simulated numbers of forced stops and mean forced wait times

		1st-instance forced stops				2nd~4th-instance forced stops	Total
		Zone 1	**Zone 2**	**Zone 3**	**Zone 4**		
The April 25 dataset							
Number of forced stops	field data	183	135	79	47	39	483
	simulation	198	128	41	22	64	453
Average forced wait time (s)	field data	17.5	13.2	15.5	8.4	12.2	16.9
	simulation	17.9	15.1	16.2	11	16.8	16.5
The July 13 dataset							
Number of forced stops	field data	87	126	96	53	34	396
	simulation	100	149	60	23	80	412
Average forced wait time (s)	field data	23.4	14.9	20.8	12.7	11.5	17.7
	simulation	20.9	14	16.1	11.6	14.7	16

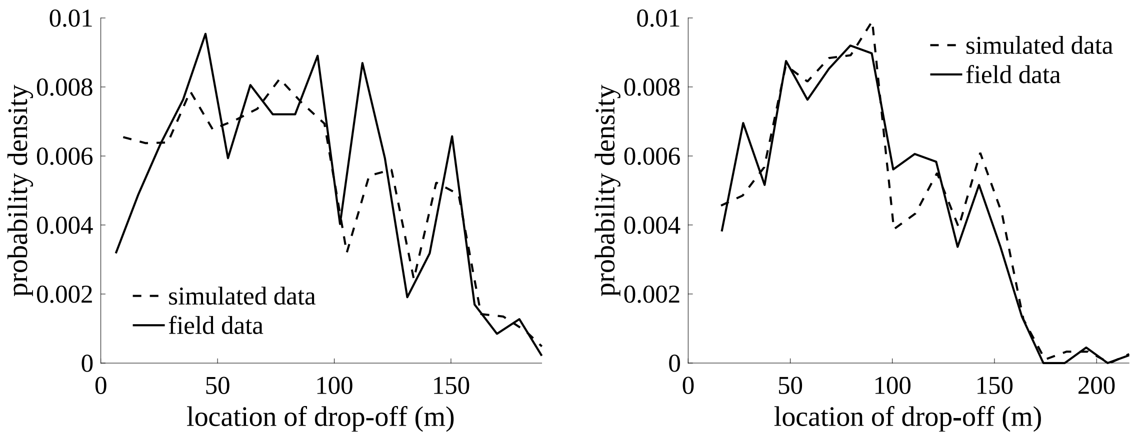

(a) the April 25 dataset (b) the July 13 dataset

Figure 7. PDFs of taxi drop-off locations

The alternatives are described in section 6.1-6.3, and taxi outflows from each are compared against simulated values produced under present-day batching. Parametric analysis presented in section 6.4 indicates that the alternatives are robust to variations in patron drop-off patterns. Maximum outflows (i.e. capacities) are assessed by setting demand for the FIFO lane at 700 taxis/h. In this way, taxi queues were always present at the lane's entry. All outcomes presented below are again averages of 500 simulations.

6.1 No-control alternative

Under the first alternative, un-batched taxis enter the FIFO lane and drop-off patrons wherever they wish. Outcomes from this no-control alternative and comparisons with present-day batching control are presented in Table 5.

The table shows that each day's outflow under the no-control alternative grew by more than 25% over present-day rates. Improvements occurred thanks to greater curbside utilization in the FIFO lane's upstream portion; i.e. note from the table that each day's average drop-off location under the alternative moved upstream by more than 10m. Note too from the table that this produced more forced stops to boot.

Table 5. Comparison of simulated outcomes between the present-day control and the no-control alternative

Parameter	Dataset	Type	Value
Outflow (taxis/h)	April 25	present day	340.5
		no control	429.3 (+26%)
	July 13	present day	331
		no control	438.4 (+32%)
Average drop-off location (m)	April 25	present day	79.5
		no control	66.2 (-13.3)
	July 13	present day	83.8
		no control	67.2 (-16.6)
Number of forced stops in Segment 1	April 25	present day	198
		no control	277 (+40%)
	July 13	present day	100
		no control	193 (+93%)

6.2 No-wait alternative

Under the second alternative, taxis having travelled a distance L_0 inside the FIFO lane must discharge their patrons upon next being forced to stop by conditions downstream. Should no forced stop occur, a taxi may discharge its patron(s) at any location desired in the lane.

The minimum-distance location for drop-offs, L_0, was examined parametrically. The full length of the FIFO lane was considered, such that $0 \leq L_0 \leq 240\text{m}$. Setting L_0 to the full length of the lane (240m), in effect, makes the no-wait alternative equivalent to the no-control alternative.

Each day's taxi outflow is plotted in Figure 8 as a function of L_0. Note how outflow is maximum when $L_0 = 0$ for both datasets. That choice of L_0 makes best use of the lane's upstream segments.

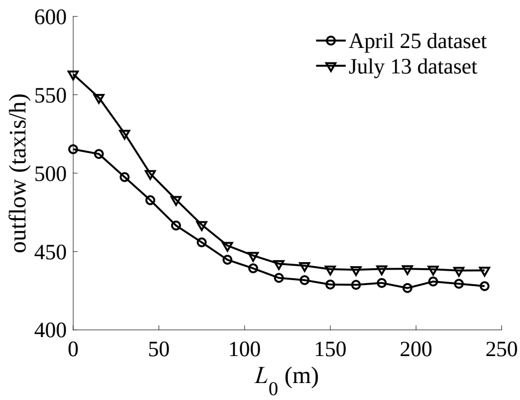

Figure 8. Effect of no-wait policy on the FIFO lane's outflow

Two related points emerge from the figure as well. First, setting $L_0 = 0$ increases taxi outflow by more than 20% over the no-control alternative (with $L_0 = 240\text{m}$), or by more than 50% over the present-day batching control. Second, benefits of this second alternative disappear when L_0 grows sufficiently large. In the present case, the curves trend horizontal when $L_0 > 120\text{m}$. Of course, setting L_0 at a small value may be objectionable to some patrons who find themselves walking long distances from their taxis to the transport terminal. This matter will be taken up in section 7.

Visual comparisons of the two curves in Figure 8 reveals that, as L_0 decrease toward 0, the outflow on July 13 increases faster than does the one on April 25. This is because on July 13, patrons were seldom observed to alight at the upstream end of the FIFO lane.

We next turn our attention to the third and final management alternative.

6.3 Promoting downstream drop-offs

The third alternative is like the second in that the parameter L_0 remains in force. Additionally, every taxi not encountering a forced stop from downstream can now drop-off patrons only upon reaching a location $L_H > L_0$.

Percentages of outflow increase as compared against the present-day control are plotted in Figures 9a and b for $L_H \in [90\text{m}, 240\text{m}]$ and $L_0 = 0$, 30, 60 and 90m.[7] The lower bound of L_H was set to 90m because, in reality, almost all lead taxis dropped patrons off beyond the 90m mark. The curves' vertical intercepts thus approximate the percentages of outflow increase for the no-wait alternative. For comparison, the percentage of outflow increase for the no-control policy is also shown as the horizontal line in each figure.

The figures show that outflows increase with large L_H, no doubt by promoting better use of the lane's downstream segments. Thus, for example, comparing the right end of each curve in Figures 9a and b against the same curve's vertical intercept unveils that introducing $L_H = 240$m typically increases taxi outflow by over 20% as compared to the no-wait alternative.

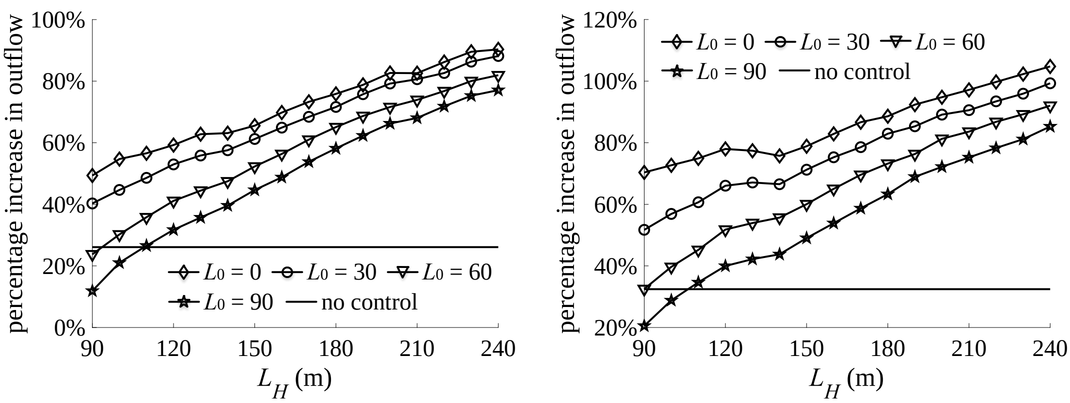

(a) the April 25 dataset (b) the July 13 dataset

Figure 9. Effect of promoting downstream drop-offs

Greatest outflows were therefore achieved by $L_0 = 0$ and $L_H = 240$m. Outflows in this extreme case were more than 90% higher than what is presently achieved via batching. These extremal thresholds, moreover, increase outflows by over 50% compared to the no-control alternative. And less restrictive thresholds of $L_0 = 90$m and $L_H = 160$m still enhance taxi outflows; e.g. by over 50% compared to present-day rates.

Taxi patrons might, of course, object to high values of L_H, as well as to low values of L_0. Matters of this kind are discussed in section 7.

[7] Larger values of L_0 were not tested in light of the findings reported in section 6.2.

6.4 Sensitivity Analysis

To verify the robustness of the benefits brought by the three alternatives, we conduct sensitivity analyses of taxi outflow gains with respect to the distributions of: (i) desired drop-off locations; (ii) drop-off durations; and (iii) patience.

For the first round of analyses, we examine two instances that differ only in the distribution of desired drop-off locations. The first instance features patrons that are "more selfish", such that their desired drop-off locations are closer to the terminal entrance. Those locations are assumed to be uniformly distributed between 100m and 140m. The second instance features patrons that are "more selfless" in that they are willing to alight further downstream in the FIFO lane; i.e., their desired drop-off locations are assumed to follow a uniform distribution between 140m and 180m. In both instances, all the other parameters take the same values as in the April 25 dataset. Figures 10a and b plot the percentages of outflow gain against the present-day control for $L_H \in [90\text{m}, 240\text{m}]$ and $L_0 = 0$, 30, 60, 90m under the two instances, respectively. Note that the curves' vertical intercepts again approximate the percent gains for the no-wait strategy for various L_0. Gains for the no-control alternative are again plotted as the horizontal lines in each figure.

In both figures, the curves exhibit similar trends as those in Figure 9a. Specifically, rescinding the present-day batching control still increases the outflow by around 20%. Enforcing the no-wait policy with $L_0 = 0$ can produce another outflow gain of over 20%. Promoting downstream drop-offs with $L_H = 240$m will bring yet another 20% or more. Comparison between Figures 10a and b unveils that outflow gains brought by the alternative policies are smaller for "selfless" patrons who are willing to alight taxis further downstream of the FIFO lane, even if they are not forced to do so. This is because their selfless behavior increases the utilization of the lane's curbside space, and thus the policies would have smaller effects.

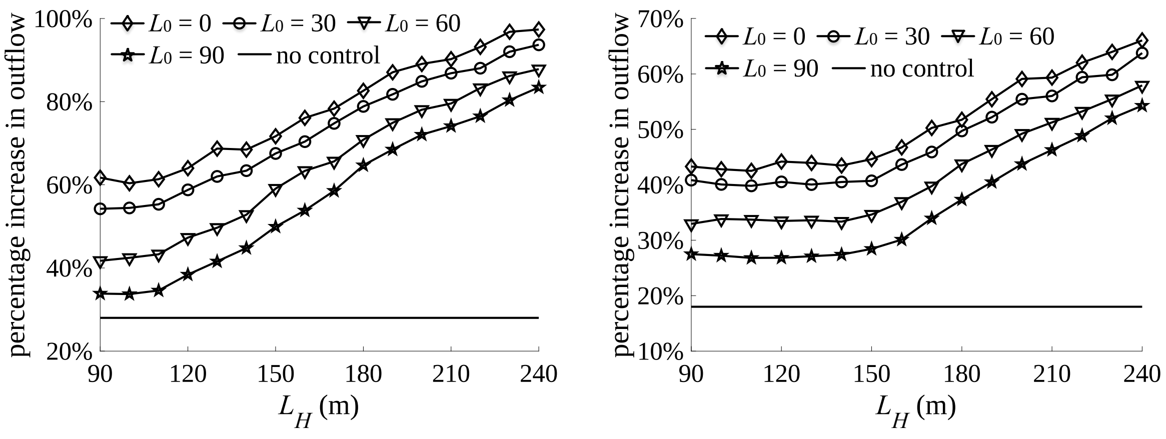

(a) More "selfish" patrons (b) More "selfless" patrons

Figure 10. Sensitivity analysis with respect to the desired drop-off locations

Our second round of sensitivity analyses entailed comparisons of two instances: (i) where the taxi drop-off durations are short and less varied, with a mean of 6s and a standard deviation of 4.2s; and (ii) where the drop-off durations are long and more varied, with a mean of 40s and a standard deviation of 20s. The outflow gains for the three alternative policies under various parameter values are plotted in Figures 11a and b for the two instances, respectively. The other parameter values are again the same as in the April 25 dataset. The curves in both figures are again similar to those in Figure 9a. The outflow gains are greater when the drop-off durations are less varied. This is as expected, because a taxi will be blocked by downstream taxis that are still dwelling in the lane. This blockage between the taxis becomes more severe when the number of taxis dwelling simultaneously in the lane increases (Gu et al., 2011; 2015; Shen et al., 2019). When the drop-off durations are less varied, this blockage is modest, and thus policies that promote better utilization of the FIFO lane's curbside space will be more effective.

The last round of analyses pertains to patron patience. Figures 12a and b plot the outflow gains for two instances: (i) less patient patrons, with a mean patience of 11s and a standard deviation of 12.7s; and (ii) patron patience has a higher mean (22.5s) and standard deviation (22.6s). The other parameter values are yet again the same as in the April 25 dataset. Note again the similarity between the two figures and Figure 9a. The no-wait and downstream drop-off policies are more effective when applied to patient patrons (Figure 12b). This is because under the two policies, patient patrons alight immediately after their taxis are forced to stop. Many of those patrons would have stayed in the taxi and waited until they could move forward again, if the policies were not enforced.

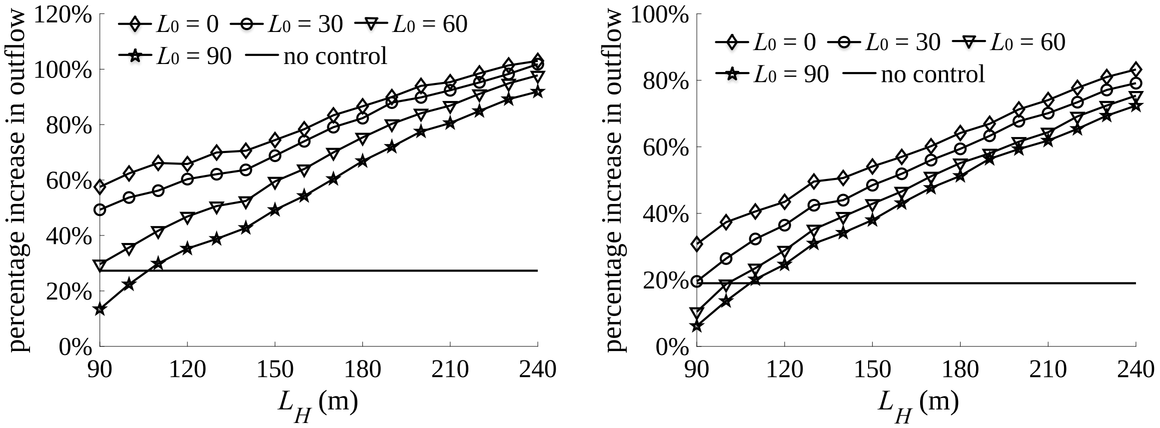

(a) Shorter and less-varied drop-off durations (b) Longer and more-varied drop-off durations

Figure 11. Sensitivity analysis with respect to the drop-off durations

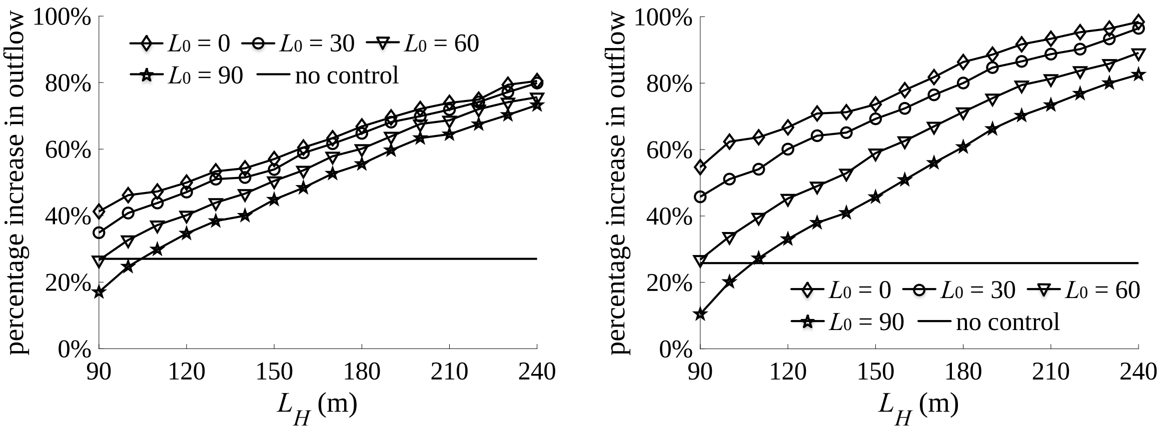

(a) Less patient patrons (b) More patient patrons

Figure 12. Sensitivity analysis with respect to the patience distribution

7. Conclusions

Simulations of a busy FIFO drop-off lane unveil the value of managing taxi operations in efficient fashion. The simulation model itself was developed in-house to emulate taxi movements in the lane. Parameters were estimated from data measured over two days. Once separately calibrated to each day's data, the model replicated the day's movements quite well. Outputs were thus used as baselines against which alternative lane-management strategies were compared.

Comparisons show that rescinding the present-day batching strategy can increase the maximum rates that taxis discharge from the FIFO lane, and thus diminish delays and queueing. Instituting a "no-control" alternative alone increased taxi outflows by more than 25%. Also tested was a distinct alternative that requires drop-offs whenever downstream conditions force a taxi to stop a distance greater than L_0 inside the lane. By promoting greater use of curb space in upstream portions of the FIFO lane, this latter alternative improved taxi outflows by up to an additional 20%. Coupling this with another requirement that taxis discharge patrons at a lengthy distance L_H inside the lane promotes greater use of downstream curb space. Instituting requirements in terms of both L_0 and L_H thus further improved outflows by as much as 20%. The alternative strategies continued to generate higher outflows when patron behavior varied from observed patterns. This underscores the robustness of the alternatives to changing inputs, or even to errors in their estimates.

The above predictions are compelling, but are not without errors. The model's failure to consider a patron's accrued delay in choosing her drop-off location is a likely source of error. The coarse method used to partition the FIFO lane (see again section 4.1) is another. Further sources may stem from unique features of taxi motion as drivers search for drop-off locations. These features are not captured in the car-following model selected for the present work. Such is the nature of simulation. Our inability to calibrate a single model to replicate operations in any given day may be a further concern, though in fairness the data suggest that taxi outflows are influenced by factors that vary day to day. These factors include train schedules and whether patrons are likely tourists or business travelers.

All these considerations motivate need for field tests. The inexorable growth in ride-sharing and ride-sourcing adds further motivation for these tests (Zha et al., 2016; Lokhandwala and Cai, 2018). They would require certain accommodations. These could be met through careful thinking, and suitable application of technologies.

In particular, the restrictive nature of our proposed drop-off rules means that some travelers would walk greater distances from their taxis to a station entrance. The onerousness of this might be lessened in simple, common-sense ways, say by providing luggage carts and human baggage handlers. Moving walkways and other commonplace technologies could play roles as well.

It would also help if stipulated drop-off distances, L_0 and L_H, were allowed to vary (e.g. over a day) based on time-varying input conditions. Stipulated distances could grow more restrictive in peak periods when taxi queues at the lane entry grow long. This sort of traffic-responsive approach would benefit from vehicle sensors, perhaps like those often used for dynamic traffic-signal and ramp-metering control (e.g. Vigos et al., 2008). Video-based surveillance could play a role here as well (Wan et al., 2014). Apprising taxi drivers of time-varying drop-off rules could rely on roadside changeable message signs (Li et al., 2016), or on-board information systems (Golob and Regan, 2005).

Surveillance, particularly of the video-based variety (Wan et al., 2014) would be needed for enforcement. The emergence of vehicle automation (Chen et al., 2016) would lessen the concern here, since the docking locations of automated taxis could be readily controlled.

Acknowledgements

The research was supported by a General Research Fund (No. 15217415) provided by the Research Grants Council of Hong Kong, and a project funded by National Natural Science Foundation of China (No. 51178111). The authors thank Liang-peng Gao, Kui-sheng Xu, Qian Yu, Yang Yang, Hong-fei Hu, Mengmiao Liu, Rui Liu, Wan-yu Yang and Zhi-peng Liu of the School of Transportation at Southeast University (China) for their help with data collection.

Appendix A. Vehicle motion model

Taxis are numbered from downstream to upstream. Their positions are updated every time interval Δt. The Δt is set to a constant coefficient termed the *reaction time*, which represents the time needed for the backward shockwave to propagate across one vehicle in queue (Daganzo, 2006; Menendez and Daganzo, 2007). Specifically, taxi n's location at $t+\Delta t$, $x^n(t+\Delta t)$, is given by:

$$x^n(t+\Delta t) = max\{l^n(t+\Delta t)\} \tag{A1}$$

subject to:

$$\Delta x_L^n(t+\Delta t) \le l^n(t+\Delta t) - x^n(t) \le \min\{\Delta x_U^n(t+\Delta t), \Delta x_S^n(t+\Delta t), \Delta x_C^n(t+\Delta t)\} \tag{A2}$$

$$s^n(t) \ge s_{jam}, \tag{A3}$$

where $\Delta x_L^n(t+\Delta t)$ and $\Delta x_U^n(t+\Delta t)$ are the minimum and maximum distances that taxi n can travel in time interval $[t, t+\Delta t]$ given the maximum deceleration and acceleration, respectively; $\Delta x_S^n(t+\Delta t)$ is the maximum distance that taxi n can travel in $[t, t+\Delta t]$ without crashing into its leader, numbered $n-1$; $\Delta x_C^n(t+\Delta t)$ is the maximum distance that taxi n can travel in $[t, t+\Delta t]$ subject to driver comfort; $s^n(t) = x^{n-1}(t) - x^n(t)$ is the spacing between taxis n and $n-1$ at time t; and s_{jam} is the jam spacing of taxis.

The $\Delta x_L^n(t+\Delta t)$, $\Delta x_U^n(t+\Delta t)$, $\Delta x_S^n(t+\Delta t)$, and $\Delta x_C^n(t+\Delta t)$ are defined by:

$$\Delta x_L^n(t+\Delta t) = \max\{0, v^n(t)\cdot\Delta t + a_L\cdot\Delta t^2\} \tag{A4}$$

$$\Delta x_U^n(t+\Delta t) = \min\{u\cdot\Delta t, v^n(t)\cdot\Delta t + a_U\cdot\Delta t^2\} \tag{A5}$$

$$\Delta x_S^n(t+\Delta t) = \max\left\{0, \frac{a_L\cdot\Delta t^2}{2} + \Delta t\cdot\sqrt{-2a_L\cdot\left[s^n(t) - s_{jam} + d^{n-1}(t)\right]}\right\} \tag{A6}$$

$$\Delta x_C^n(t+\Delta t) = s^n(t) - s_{jam}, \tag{A7}$$

where $v^n(t)$ denotes taxi n's average speed in $[t-\Delta t, t]$, given by $v^n(t) = \frac{x^n(t) - x^n(t-\Delta t)}{\Delta t}$; a_L and a_U are the minimum acceleration (i.e. the opposite of maximum deceleration) and maximum acceleration of the taxi, respectively; u is the desired travel speed; and $d^{n-1}(t)$ is the minimum stopping distance of taxi $n-1$ at time t. The $d^{n-1}(t)$ is given by:

$$d^{n-1}(t) = max\left\{0, -\frac{\left[v^{n-1}(t)\right]^2}{2a_L} - \frac{v^{n-1}(t)\cdot\Delta t}{2}\right\}. \tag{A9}$$

Derivation of (A1)-(A9) can be found in Menendez and Daganzo (2007) and Menendez (2006), and is omitted here for brevity.

Appendix B. The k-mean clustering method

The FIFO lane was partitioned by clustering taxis' forced wait times at their first instances of forced stops. Taxis' second, third and fourth instances of forced stops were excluded because they were of much shorter durations. For a given number of segments k, we seek a partition that minimizes the sum of total squared errors of forced wait times in each segment, ε:

$$\min_{C\triangleq\{C_1,C_2,\dots,C_k\}} \varepsilon = \sum_{i=1}^{k} \sum_{n:y^n\in C_i} (t^n - u_i)^2, \tag{B1}$$

where C is a lane partition, with each C_i $(i = 1,2, \dots, k)$ defining a continuous space interval (i.e. a segment) in [0,240m], $\bigcup_{i=1}^{k} C_i = [0{,}240\text{m}]$; y^n is the location of taxi n's first forced stop (given that the taxi is not leading a batch); t^n is taxi n's forced wait time during that stop; and $u_i = E[t^n | y^n \in C_i]$.

Appendix C. Estimation of patience distribution

The probability density function (PDF) of a mixture distribution for patrons' patience is given as:

$$f(p) = \gamma f_1(p) + (1-\gamma) f_2(p), \tag{C1}$$

where $f_1(p)$ and $f_2(p)$ are the PDFs of patience distributions for impatient patrons (i.e., those who alighted almost immediately after being forced to stop) and the remaining, patient ones, respectively; and γ is the probability that a taxi's patron(s) were impatient. When $f_1(p)$ and $f_2(p)$ are gamma PDFs, we have:

$$f(p;\gamma,k_1,\theta_1,k_2,\theta_2) = \gamma f_1(p;k_1,\theta_1) + (1-\gamma) f_2(p;k_2,\theta_2)$$

$$= \gamma \frac{p^{k_1-1} e^{-p/\theta_1}}{\theta_1^{k_1}\Gamma(k_1)} + (1-\gamma) \frac{p^{k_2-1} e^{-p/\theta_2}}{\theta_2^{k_2}\Gamma(k_2)}, \tag{C2}$$

where k_1 and k_2 are the shape parameters, and θ_1 and θ_2 are the scale parameters for $f_1(p)$ and $f_2(p)$, respectively; and $\Gamma(\cdot)$ is the gamma function.

To estimate the values of k_1, k_2, θ_1, θ_2 and γ, we formulate the log-likelihood function for the forced wait times as:

$$\Psi(\gamma,k_1,\theta_1,k_2,\theta_2) = \sum_{n\in\mathcal{P}} \ln f(t^n,\gamma,k_1,\theta_1,k_2,\theta_2) + \sum_{n\in Q} \ln[1 - F(t^n,\gamma,k_1,\theta_1,k_2,\theta_2)], \tag{C3}$$

where $\mathcal{P}$ denotes the index set of taxis that dropped-off patrons at the present forced stop (i.e., the taxis whose forced waits equaled their patience); Q denotes the index set of taxis that did not drop-off patrons at the forced stop (i.e., those whose forced waits were less than their patience); and $F(\cdot)$ is the CDF of the mixture distribution.

The MLE problem is then formulated as:

$$\max_{\gamma,k_1,\theta_1,k_2,\theta_2} \Psi(\gamma,k_1,\theta_1,k_2,\theta_2). \tag{C4}$$

This problem was solved by the nonlinear program solver "fminsearch" in Matlab R2017b.

The distribution parameter estimates for the two datasets are presented in Table C1. Note in each day's data that, for all the five distributions, the mean and variance for the impatient patrons are much smaller than those for the remaining, patient ones; i.e., $k_1\theta_1 \ll k_2\theta_2$ and $k_1\theta_1^2 \ll k_2\theta_2^2$ for all the five rows of each dataset. Also, the table shows that the probability of impatient patrons, γ, increases from Segment 1 to Segment 4. This is consistent with intuition, since patrons were observed to become less patient as they moved downstream.

The above distributions are coarse estimates of patron patience due to the limited data. Better estimates can be obtained by using more sophisticated methods (e.g., the one developed in Sun and Elefteriadou, 2014), should larger, more detailed datasets be available.

Appendix D. Estimation of other parameters

D.1 Taxi-batching parameters

We assume that a taxi batch is admitted to the FIFO lane whenever either: (i) the lane is vacant for a distance L_{m1} in its upstream-most portion; or (ii) the lane is vacant for at least a distance $L_{m2} < L_{m1}$ in its upstream portion, and the last taxi in the previous batch has dwelled for a duration of at least T_m. The admission of secondary batches of taxis follows the same logic, but with distinct values for parameters L_{m1}, L_{m2} and T_m. We further denote L_{left} as the lane space upstream of a batch that is left unoccupied when the batch stops. The number of taxis in a batch is thus determined by dividing the length of the batch (e.g., $L_{m1} - L_{left}$) by the jam, or stopped-vehicle spacing. Parameters L_{m1}, L_{m2} and T_m were estimated for each day's data via the k-means clustering algorithm (Hartigan and Wong, 1979). Parameter L_{left} was set to the average lane space upstream of the primary and secondary batches, respectively, again for each day's data.

Table C1. Optimal parameters for patience distribution

	k_1	θ_1	$k_1\theta_1$	$k_1\theta_1^2$	k_2	θ_2	$k_2\theta_2$	$k_2\theta_2^2$	γ
The April 25 dataset									
1st-instance forced stops in Segment 1	2.13	1.42	3.02	4.29	3.62	8.77	31.8	278.4	0.43
1st-instance forced stops in Segment 2	1.81	1.25	2.26	2.83	3.63	7.29	26.5	192.9	0.48
1st-instance forced stops in Segment 3	0.97	6.80	6.60	44.85	5.71	4.78	27.3	130.5	0.55
1st-instance forced stops in Segment 4	1.10	2.15	2.37	5.08	6.25	3.17	19.8	62.8	0.64
2nd~4th-instance forced stops	3.48	0.70	2.44	1.71	3.75	8.72	32.7	285.1	0.49
The July 13 dataset									
1st-instance forced stops in Segment 1	17.09	0.17	2.91	0.49	2.71	14.74	39.95	588.80	0.25
1st-instance forced stops in Segment 2	14.67	0.16	2.35	0.38	1.84	14.57	26.81	390.60	0.40
1st-instance forced stops in Segment 3	3.92	1.05	4.12	4.32	4.44	7.23	32.10	232.09	0.36
1st-instance forced stops in Segment 4	6.27	0.47	2.95	1.39	5.93	4.58	27.16	124.39	0.56
2nd~4th-instance forced stops	20.02	0.13	2.60	0.34	1.70	13.01	22.12	287.74	0.31

D.2 Distributions for drop-off locations and durations

We assume that the lead taxis of each batch dropped off patrons at their desired locations, and fit an empirical distribution (van der Vaart, 2000) to those locations measured from the videos. The desired drop-off locations of the other taxis were assumed to follow the same distribution, since a taxi's desired drop-off location should be irrespective of whether or not it is a batch leader. Two other empirical distributions were fit to the drop-off durations of lead taxis and other taxis. For the latter, the drop-off duration is defined as the time between the taxi door opening and closing, plus a fixed time spent on necessary drop-off activities that occur before door opening and after door closing (e.g., payment collection and receipt preparation). This fixed time was estimated by subtracting the average time between door openings and closings from the average dwell time for lead taxis.

References

Aïvodji, U.M., Gambs, S., Huguet, M.J., Killijian, M.O., 2016. Meeting points in ridesharing: A privacy-preserving approach. Transportation Research Part C: Emerging Technologies, 72, 239-253.

Ashford, N., Wright, P., Mumayiz, S., 2011. Airport Engineering: Planning, Design, and Development of 21st Century Airports. John Wiley & Sons, Hoboken, New Jersey.

Bender, G., Chang, K., 1997. Simulating roadway and curbside traffic at Las Vegas McCarran. IIE Solutions, 29(11), 26-30.

Cassidy, M.J., Kim, K., Ni, W., Gu, W., 2015a. A problem of limited-access special lanes. Part I: Spatiotemporal studies of real freeway traffic. Transportation Research Part A: Policy and Practice, 80, 307-319.

Cassidy, M.J., Kim, K., Ni, W., Gu, W., 2015b. A problem of limited-access special lanes. Part II: Exploring remedies via simulation. Transportation Research Part A: Policy and Practice, 80, 320-329.

Chang, K., Haghani, A., Bender, G., 2000. Traffic simulation at airport terminal roadway and curbside. Proceedings of the International Air Transportation Conference, 165-176.

Chang, K., 2001. A simulation model for analyzing airport terminal roadway and traffic and curbside parking. Ph.D. Thesis. University of Maryland.

Chen, Z., He, F., Zhang, L., Yin, Y., 2016. Optimal deployment of autonomous vehicle lanes with endogenous market penetration. Transportation Research Part C: Emerging Technologies, 72, 143-156.

Costa, D., de Neufville, R., 2012. Designing Efficient Taxi Pick-up Operations at Airports. Transportation Research Record, 2300, 91-99.

Daganzo, C., 2006. In traffic flow, cellular automata=kinematic waves. Transportation Research Part B: Methodological, 40(5), 396-403.

de Neufville, R., 1976. Airport Systems Planning: A Critical Look at the Methods and Experience. Macmillan, London.

de Neufville, R., Odoni, A., 2013. Airport Systems Planning, Design, and Management, 2nd edition. Mc-Graw-Hill.

Farhan, J., 2015. An agent-based multimodal simulation model for capacity planning of a cross-border transit facility. Transportation Research Part C: Emerging Technologies, 60, 189-210.

Golob, T.F., Regan, A.C., 2005. Trucking industry preferences for traveler information for drivers using wireless Internet-enabled devices. Transportation Research Part C: Emerging Technologies, 13(3), 235-250.

Gu, W., Cassidy, M.J., Li, Y., 2014. Models of bus queueing at curbside stops. Transportation Science, 49(2), 204-212.

Gu, W., Li, Y., Cassidy, M.J., Griswold, J.B., 2011. On the capacity of isolated, curbside bus stops. Transportation Research Part B: Methodological, 45(4), 714-723.

Hall, C., 1977. A simulation model for an enplaning-passenger-vehicle curbside at high-volume airports. Ph.D. Thesis. University of Colorado.

Hartigan, J., Wong, M., 1979. A k-means clustering algorithm. Journal of the Royal Statistical Society, Series C, 28(1), 100-108.

Ji, Y., Cao, Y., Du, Y., Zhang, H.M., 2016. Comparative analyses of taxi operations at the airport. World Conference on Transport Research, Shanghai, July 2016.

Li, M., Lin, X., He, F., Jiang, H., 2016. Optimal locations and travel time display for variable message signs. Transportation Research Part C: Emerging Technologies, 69, 418-435.

Li, R., Liu, Z., Zhang, R., 2018. Studying the benefits of carpooling in an urban area using automatic vehicle identification data. Transportation Research Part C: Emerging Technologies, 93, 367-380.

Lokhandwala, M., Cai, H., 2018. Dynamic ride sharing using traditional taxis and shared autonomous taxis: A case study of NYC. Transportation Research Part C: Emerging Technologies, 97, 45-60.

Mandle, P., Lamagna, F., Whitlock, E., 1980. Collection of calibration and validation data for an airport landside dynamic simulation model. Report EM-80-2, Federal Aviation Administration.

Mandle, P., Whitlock, E., Lamagna, F., 1982. Airport curbside planning and design. Transportation Research Record, 840, 1-6.

McCabe, L., Carberry, T., 1975. Simulation methods for airport facilities. Transportation Research Board Special Report, 159.

Menendez, M., 2006. An analysis of HOV lanes: their impact on traffic. PhD thesis, Department of Civil and Environmental Engineering, University of California, Berkeley, CA.

Menendez, M., Daganzo, C.F., 2007. Effects of HOV lanes on freeway bottlenecks. Transportation Research Part B: Methodological, 41(8), 809-822.

Parizi, M., Braaksma, J., 1994. Optimum design of airport enplaning curbside areas. Journal of Transportation Engineering, 120(4), 536-551.

Shapiro, P., 1996. Intermodal ground access to airports: a planning guide - a good start. TRB Conference on the Application of Transportation Planning Methods, 114-121.

Shen, M., Gu, W., Hu, S., Cheng, H., 2019. Capacity approximations for near-and far-side bus stops in dedicated bus lanes. Transportation Research Part B: Methodological, 125, 94-120.

Sun, D., Elefteriadou, L., 2010. Research and implementation of lane-changing model based on driver behavior. Transportation Research Record, 2161, 1-10.

Sun, D., Elefteriadou, L., 2014. A driver behavior-based lane-changing model for urban arterial streets. Transportation Science, 48(2), 187-205.

Sun, D.J., Zhang, L., Chen, F., 2013. Comparative study on simulation performances of CORSIM and VISSIM for urban street network. Simulation Modelling Practice and Theory, 37, 18-29.

Tillis, R., 1973. Curb space at airport terminals. Traffic Quarterly, 27(4), 563-582.

Tunasar, C., Bender, G., Yung, H., 1998. Modeling curbside vehicular traffic at airports. The 30th Conference on Winter Simulation, 1113-1118.

van der Vaart, A.W., 2000. Asymptotic Statistics, 3. Cambridge University Press.

Vigos, G., Papageorgiou, M., Wang, Y., 2008. Real-time estimation of vehicle-count within signalized links. Transportation Research Part C: Emerging Technologies, 16(1), 18-35.

Wan, Y., Huang, Y., Buckles, B., 2014. Camera calibration and vehicle tracking: Highway traffic video analytics. Transportation Research Part C: Emerging Technologies, 44, 202-213.

Wang, L., 1990. Simulation of airport curbside operations. Master Thesis. University of Maryland.

Whitlock, E., Cleary, E., 1969. Planning ground transportation facilities for airports. Highway Research Record, 274.

Yang, F., Gu, W., Cassidy, M.J., Li, X., Li, T., 2019. Achieving higher taxi outflows from a congested drop-off lane: A simulation-based policy study. Working paper. http://arxiv.org/abs/1910.02484.

Zha, L., Yin, Y., Yang, H., 2016. Economic analysis of ride-sourcing markets. Transportation Research Part C: Emerging Technologies, 71, 249-266.

Zhang, Z., 2000. A flexible new technique for camera calibration. IEEE Transactions on Pattern Analysis and Machine Intelligence, 22(11), 1330-1334.