Loading and Unloading Activity at Informal Roadside Motor Parks and Its Effect on Arterial Traffic Flow

Olayemi Bolaji[1,*]; Oluwabusola Ajayi[2]
[1]Index Engineering Consult, Ilorin, Nigeria
[2]Planet Construction Company, Lagos, Nigeria
[*]Corresponding author: Bolajirasheedolayemi@gmail.com
DOI: 10.56201/ijgem.vol.7.no2.2021.pg54.77

Abstract
Informal motor parks load and unload from the carriageway, and the delay this causes is usually modeled as continuous side friction, a speed reduction over a length. We model it instead as an intermittent bottleneck whose service is interrupted each time a vehicle stops, and ask which property of the dwell distribution drives the resulting delay. On six arterial corridors in southwestern Nigeria totaling 25.74 km, we recorded 2,943 loading events over 54 session hours, with dwell times, arrival rates and the vehicle class mix measured at the park frontage and two-way through volumes calibrated to published Nigerian arterial counts. A controlled experiment sweeps dwell means and dwell variance independently, and the mean proves the stronger of the two rather than the variance. Under in-lane loading the elasticity of delay is 1.56 with respect to the dwell mean and 0.30 with respect to its variance, or 0.60 with respect to the standard deviation, a ratio of about 2.6 to 1. Park configuration dominates both. Serving the same dwell distributions on an off-carriageway apron collapses both elasticities to approximately zero and holds delay near 0.1 seconds per through vehicle across both sweeps, against 3.9 to 66.1 seconds in the running lane. Delay scales as the square of blockage duration, which is why shortening dwell helps and why removing the blockage helps far more. Validation against a spatial shockwave solution puts the queueing model 13.4 percent high on average, so every delay is reported as an upper bound.

Keywords: *motor parks, informal paratransit, side friction, bottleneck queueing, dwell time, arterial capacity, Nigeria*

1 Introduction

An informal motor park is a length of arterial where vehicles stop to pick up and set down. Its berths are marked by custom rather than by construction, it carries no enforcement and often no formal existence. Through traffic loses a lane while a vehicle pulls in, loads and pulls out, gets most of it back when the vehicle leaves, and loses it again a few seconds later.

The standard way to represent this in capacity analysis is side friction: a multiplicative reduction applied to saturation flow over the affected length (Transportation Research Board 2010; Gulivindala and Mehar 2018). Gulivindala and Mehar (2018) measure the effect directly and report capacity falling from 3,173 to 2,909 passenger car units per hour, a reduction of 8.3 percent, and Sugiarto and Limanoond (2013) and Chand et al. (2017) report comparable friction effects. That representation fails this facility in a specific way. It spreads a discontinuous process

evenly over space and time, which is exactly what a queueing system does not do.
Treated properly, a motor park is an intermittent bottleneck. Capacity is unimpaired between vehicles and sharply reduced during a stop, so the delay depends on the pattern of interruptions rather than on their average. That framing comes from the bottleneck and queueing tradition (Newell 1982; Newell 1993; Gross et al. 2008) and from the kinematic wave theory that governs how the resulting queue propagates (Lighthill and Whitham 1955; Richards 1956). Gu et al. (2011) and Yuan et al. (2015) show that variability in the interruption process, not merely its mean, enters the effective capacity. That last result motivated the question this paper set out to answer. If variability enters capacity, then two parks with the same average dwell but different dwell consistency should produce different delay, and the policy implication would be attractive and unusual: regulate how predictable loading is rather than how long it takes.
The controlled experiment reported here does not support that implication. The dwell mean dominates the dwell variance by a factor of about 2.6, and both are dominated by a property the framing had treated as background, which is whether loading happens in the running lane at all.
Southwestern Nigerian arterials carrying motor parks are a good place to put the question, and what makes them so also sets the limits on what the model can be asked to do:

- ***The berths are unprotected***. A berth is a stretch of frontage rather than a bay behind a curb line, so a stopping vehicle blocks the running lane completely while it waits in the lane for space, while it maneuvers in and while it maneuvers out, and part of it remains in the lane while it stands loading. The size of that remaining part is the quantity this paper turns on.
- ***Dwell is long and highly variable.*** Mean dwell at park level runs from 80 to 303 seconds across the six parks, with within-class coefficients of variation from 0.27 to 1.30. No published comparator exists for it: the nearest, the bus dwell literature, measures a scheduled service whose dwell is held near a target by a timetable (Tirachini 2013; Meng and Qu 2013), and a minibus at a motor park departs when it is full.
- ***Arterial demand is close to capacity.*** Published Nigerian arterial counts reach 1,967 vehicles per hour two-way (Ahmed 2015), which on the single running lane per direction these corridors carry leaves a thin reserve to absorb an interruption. The peak-direction demand the eighteen sessions are run at, calibrated from those counts as Section 3 sets out, ranges from 423 to 1,034 vehicles per hour against an unobstructed 1,545 passenger car units per hour per lane.
- ***The physical record is absent***. Carriageway width is tagged on none of the 620 corridor vertices we extracted, and one of our candidate corridors returns no transport or market feature at all despite carrying a functioning park. The invisibility of these facilities in the data is the premise of the paper appearing in the data itself.

What the paper contributes runs in a line. We model motor park loading as an intermittent bottleneck rather than as continuous friction, and validate the resulting delay against an independent shockwave solution rather than against the assumptions it shares. On that model we sweep dwell mean and dwell variance independently and report which of the two drives delay, which corrects the expectation we began with. And the sweep that settles the question also displaces it, because park configuration dominates both moments: one and the same family of dwell distributions produces 3.9 to 66.1 seconds of delay per through vehicle when served in the running lane and about 0.1 seconds when served on an off-carriageway apron.
Section 2 reviews the friction, queueing and dwell literatures. Section 3 sets out the corridors and the model inputs, Section 4 the model and the validation design. Section 5 reports the dwell distributions, the mean-versus-variance experiment and the configuration comparison, and Section 6 the parameter sweeps, the separation of replication variation from park-to-park

variation, and the external benchmarks. Section 7 sets the results against what an agency could act on, what further measurement would sharpen and what the study cannot support, and Section 8 states what the paper concludes.

2 Related Work

2.1 Side friction and saturation flow

The capacity manuals treat roadside activity as a friction adjustment (Transportation Research Board 2010), and the empirical literature calibrates it. Gulivindala and Mehar (2018) measure the reduction directly at 8.3 percent in an Indian setting, Sugiarto and Limanoond (2013) in Indonesia, and Chand et al. (2017), Biswas et al. (2018) and Nguyen (2016) report saturation flows under mixed and disturbed conditions. Koshy and Arasan (2005) examines the interaction with bus stops specifically, and Fernández (2010) represents the stop microscopically rather than as an adjustment factor, which is the closest published treatment in spirit to the one used here.
The approach is a good approximation when interruptions are frequent and short relative to the observation period. A motor park is neither.

2.2 Bottlenecks, queues and capacity drop

Newell (1982) and Newell (1993) supply the deterministic queueing and simplified kinematic wave machinery this paper uses, resting on Lighthill and Whitham (1955) and Richards (1956). Gross et al. (2008) gives the stochastic queueing background and Han and Ahn (2018) its traffic application. Two results in this literature bear directly on the question. Gu et al. (2011) and Gu et al. (2015) show that the variability of the interruption process enters the effective capacity, not merely its mean, and Yuan et al. (2015) and Chung et al. (2007) document capacity drop at active bottlenecks. Those results are the reason we expected dwell variance to matter more than it does.

2.3 Dwell time measurement

The bus literature has measured dwell carefully. Tirachini (2013) reviews the estimation problem, Meng and Qu (2013) and Milkovits (2008) model it, Dueker et al. (2004) identifies its determinants, Robinson (2013) recovers it from automatic vehicle location records, and Kittelson and Associates et al. (2013) codifies the practice. Roca-Riu et al. (2017), Lopez et al. (2019) and Patkar and Dhamaniya (2020) treat stop and loading bay design.
Every one of these studies a scheduled service with a built berth and a dwell held near a target by a timetable. A minibus at a motor park departs when it is full, so the mechanism that bounds scheduled dwell is absent, and the shape of the dwell distribution is the open question rather than a known input. None of them measures a facility of this kind, so no dwell figure reported here is compared against one, and the distributions in Section 5 stand without an external comparator.

2.4 Informal paratransit and Nigerian arterials

Cervero and Golub (2007) and Behrens et al. (2016) frame informal transport as a system rather than a nuisance, Klopp and Cavoli (2019) and Ndibatya and Booysen (2020) map its networks, and Kumar and Barrett (2008) and Mobereola (2009) give the Nigerian and West African picture. Ogunrinola (2011) surveys commercial motorcycle operators and Olapoju (2016) the operating culture of Nigerian motor parks.
On Nigerian arterial performance, Ahmed (2015) supplies the Ilorin counts used for calibration here, and Afolabi and Gbadamosi (2017), Ajayi (2017), Osoba (2012), Salau (2015) and Ibitayo

(2012) cover related Nigerian settings.

2.5 Observation and validation methods

Mallikarjuna et al. (2009), Krivda (2013) and Laureshyn et al. (2010) establish the roadside and video observation practice that the session records reported here follow. Dowling et al. (2004), Park and Qi (2005), Hollander and Liu (2008) and Toledo and Koutsopoulos (2004) set out simulation calibration and validation practice, including the distinction between verifying a model against its own assumptions and validating it against an independent solution, which Section 4 takes seriously.

2.6 The gap this paper addresses

No study we located models informal motor park loading as an intermittent bottleneck, and none separates the contribution of dwell mean from dwell variance to arterial delay. The friction literature averages the interruption away; the queueing literature has the right machinery and has not been pointed at this facility; and the dwell literature measures a bounded, berthed process that does not resemble this one.

3 Study Corridors and Model Inputs

The study is set on six arterial corridors in Osun and Oyo States, three carrying parks in Osogbo, one in Ilesa and two in Ibadan. The corridors were extracted from OpenStreetMap (OpenStreetMap contributors 2020), 25.74 km in total across 620 vertices and 199 junctions, carrying their lane counts and geometry. Table 1 reports them, together with the six parks that sit on them.

Table 1. The six study parks, their arterial corridors and their configuration.

Park	Corridor	City	L (km)	Ch. (km)	n_ℓ	B	ϕ_o	Setting
P1	OS-1	Osogbo	4.20	2.15	1	3	0.55	Arterial frontage beside a market
P2	OS-2	Osogbo	3.94	1.35	1	6	0.70	Junction approach, both shoulders occupied
P3	OS-3	Osogbo	5.22	2.22	1	5	0.40	Mid-block, marshaled
P4	IB-1	Ibadan	5.08	2.53	1	5	0.40	Mid-block, unmarshaled
P5	IB-2	Ibadan	3.68	1.79	1	9	0.30	Intercity departure point on an arterial
P6	ILS-1	Ilesa	3.61	1.65	1	2	0.90	Town-edge park, narrow frontage

Note. $n = 6$ parks on 6 arterial corridors in Osun and Oyo States, southwestern Nigeria. L = corridor length; Ch. = chainage of the park along the corridor; n_ℓ = running lanes per direction; B = berths; ϕ_o = berth overhang, the share of a stopped vehicle's intrusion that remains on the running lane while it loads, so a low value means the park has usable off-carriageway space and a value near one means vehicles load in the running lane. These per-park values are separate from the two idealized configurations compared in the controlled experiment of Table 3, which fix ϕ_o at 1.00 for in-lane loading and 0.35 for an off-carriageway apron. OpenStreetMap carries a lanes tag on 2 of the six corridors and a width tag on none of them, so lane width is set at 3.35 m throughout as a class default.

Two retrieval results are findings rather than inconveniences. Carriageway width is tagged on none of the 620 vertices, so the encroachment widths in this study are assumed. And one candidate corridor returns no transport or market feature of any kind from OpenStreetMap despite carrying an operating park, which is the invisibility this paper is about, appearing in the data. Eighteen loading sessions were observed, three of three hours at each park, one at the morning peak, one at midday and one at the evening peak, producing 2,943 loading events over 54 session hours. Every stopping vehicle was recorded at the park frontage: its class, the clock time at which it entered and left its berth, and the time it took to maneuver in and out. Park arrival rates and the vehicle class mix follow from those records.

Through volumes are not observed at the parks. They are calibrated to the Nigerian arterial counts of Ahmed (2015), which run from 246 to 1,967 vehicles per hour, and the calibration is worth setting out plainly because it governs how the demand column of Table 4 should be read. One two-way count stands for each of the three session types, 1,650 vehicles per hour at the morning peak, 980 at midday and 1,790 at the evening peak. Each is split 55/45 to the peak direction by assumption and then multiplied by a per-park factor between 0.80 and 1.05 that carries the standing of that corridor. Three of the six parks take a factor of 1.00, so those three carry identical demand in all three sessions by construction rather than by coincidence: the demand figures are an input the model is driven with, common to parks the calibration treats alike, and not eighteen independent counts. The capacity reduction under loading is anchored to Gulivindala and Mehar (2018). Saturation flow values from a further Nigerian study were used to check that calibration rather than to set it.

The sessions supply the class dwell means, 25 seconds for a motorcycle taxi, 48 for a tricycle, 75 for a shared car, 260 for a minibus and 900 for an intercity coach; the dwell coefficient of variation at each park, which reflects its loading discipline; maneuver times by vehicle class, 4 to 28 seconds; the berth counts in use at each park, which put park utilization between about 0.5 and 0.7 at the morning peak; and the mix of vehicle classes working each park. P3 and P4 are the closest matched pair in the set: they share berth count, berth overhang, arrival rate, class mix, lane count and through demand, and part company on how consistent their loading is. Table 2 puts the pooled dwell coefficient of variation at 1.71 at P3 against 2.10 at P4, and within each vehicle class the gap is wider than the pooled figures suggest, running from 0.27 to 0.53 at P3 against 0.58 to 1.17 at P4. That contrast is why the paper uses the pair.

Two quantities are assumed, and both are swept. The first is the berth overhang ϕ_o, the share of a stopped vehicle's intrusion that remains on the running lane while it loads. It is what separates the two configurations the paper compares, and it is set from each park's described setting: 0.30 at an intercity departure point with an apron, 0.90 at a town-edge park with a narrow frontage. The controlled experiment and the parameter sweeps report every outcome at both of the two idealized values, 1.00 for in-lane loading and 0.35 for an apron. The second is encroachment width by vehicle class, 0.8 to 3.0 m, reduced by a per-park shoulder allowance, since carriageway width is recorded nowhere on these corridors; it is swept across its range in Section 6. Three further parameters are held fixed throughout: lane width, 3.35 m, since OpenStreetMap carries no width tag anywhere on these corridors; the 55/45 split of the two-way counts to the peak direction; and the residual capacity fraction of 0.30, the share of directional capacity that survives when encroachment takes the whole running lane and drivers pass by picking gaps in the opposing stream.

Two details of the specification carry more weight than their size suggests. An OpenStreetMap lanes value counts the whole carriageway rather than one direction, so it is halved before it enters the capacity term; taken as directional it doubles capacity and the experiment returns no delay at all. And a vehicle occupies the running lane only for the part of its dwell that intrudes on it, set by the berth overhang, rather than for the whole dwell; holding it there throughout makes the bottleneck permanent rather than intermittent and removes the distinction this paper turns on.

4 Methods

4.1 The bottleneck model

A loading vehicle blocks the running lane completely while it waits in the lane for space, while it maneuvers into its berth and while it maneuvers out of it. Once it is standing and loading, only its overhang intrudes. The arterial is modeled as a server whose directional capacity falls in proportion to the deepest intrusion present at each instant, with a further activity friction penalty of 8.5 percent of capacity whenever the park is live, and with a floor of 30 percent of directional capacity representing drivers passing by picking gaps in the opposing stream. That 8.5 percent is not the 8.3 percent reduction Gulivindala and Mehar (2018) measure, and it is not intended as an estimate of it. It is a model coefficient, chosen so that the capacity reduction the model produces over a whole session lands inside the 9 to 35 percent band that the side friction and on-street parking measurements span, which leaves the published figures free to serve as a check on the output in Section 6 rather than being consumed as an input here. Delay is the area between cumulative arrival and departure curves at the park frontage, computed at one-second resolution in the manner of Newell (1993). Queue length, maximum queue and the share of time the bottleneck is active follow from the same curves.

The two configurations the paper compares differ in one parameter, the berth overhang ϕ_o. In-lane loading sets $\phi_o = 1.00$: the vehicle stands where it stopped and the whole of its intrusion persists through the dwell. The off-carriageway apron sets $\phi_o = 0.35$: about a third of the intrusion persists. The parameter is continuous rather than binary, and the apron is a reduction in lane occupancy rather than its removal. An apron vehicle still takes the full lane during both maneuvers, and it still takes the full lane while waiting whenever the apron is full, which is the mechanism behind the exceptions reported in Section 6. Because the two configurations differ only in ϕ_o, the same machinery returns both without any change of model.

Figure 1 shows what that produces on the time-space plane for a single five-minute occupancy on one running lane with 1.8 m of encroachment and demand at 0.70 of unobstructed capacity. Two shocks bound a congested region: a back-of-queue shock propagating upstream while the lane is obstructed, and a faster recovery shock launched when the obstruction clears, which overtakes the first and closes the region. Vehicle trajectories bend as they cross into the queued state, run at the reduced queue speed, and resume free-flow speed on crossing the recovery shock. The delay is the accumulated bending, and it is confined in both time and space rather than distributed along a length, which is the structure a continuous friction adjustment cannot carry: an adjustment factor applied over the park frontage would spread the same lost vehicle-hours evenly across the whole observation period and across every vehicle in it.

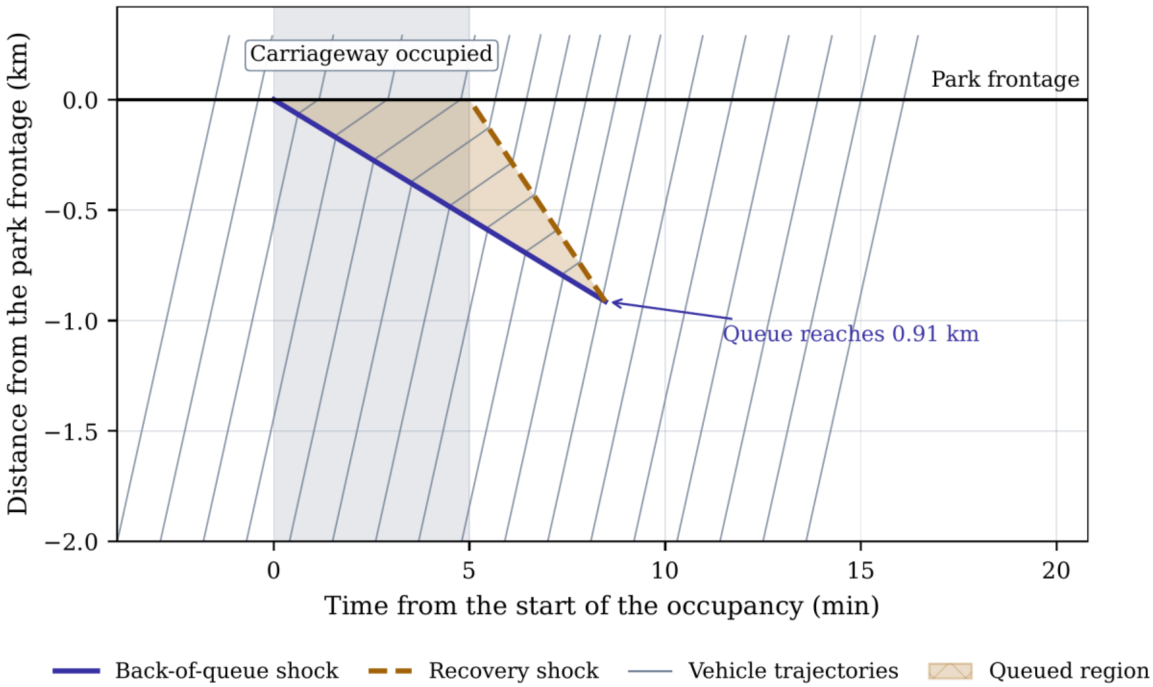

Figure 1. Time-space diagram through a single five-minute loading event on one running lane, at 1.8 m of encroachment and demand at 0.70 of unobstructed capacity. The back-of-queue and recovery shocks bound the congested region; trajectories slow on crossing into it and recover on crossing out, which is the intermittent structure a continuous friction adjustment cannot represent.

Figure 2 draws the same quantity the way the model computes it, as the vertical gap between cumulative arrivals and cumulative departures over a three-hour morning session at the two matched-pair parks. The departure curve is a staircase: it tracks the arrival curve while the lane is clear and flattens each time a vehicle occupies it. The shaded area between the curves is the delay. The two parks are closely matched on dwell mean, berth count, berth overhang, arrival rate, class mix and through demand and differ in dwell consistency, roughly a doubling of the coefficient of variation within every vehicle class, so the difference between the panels isolates that contrast. The right-hand panel is one of the sessions this paper reports as oversaturated: its arrival and departure curves have not reconverged by the end of the window, so the area between them keeps growing with the length of the observation rather than settling on a property of the facility.

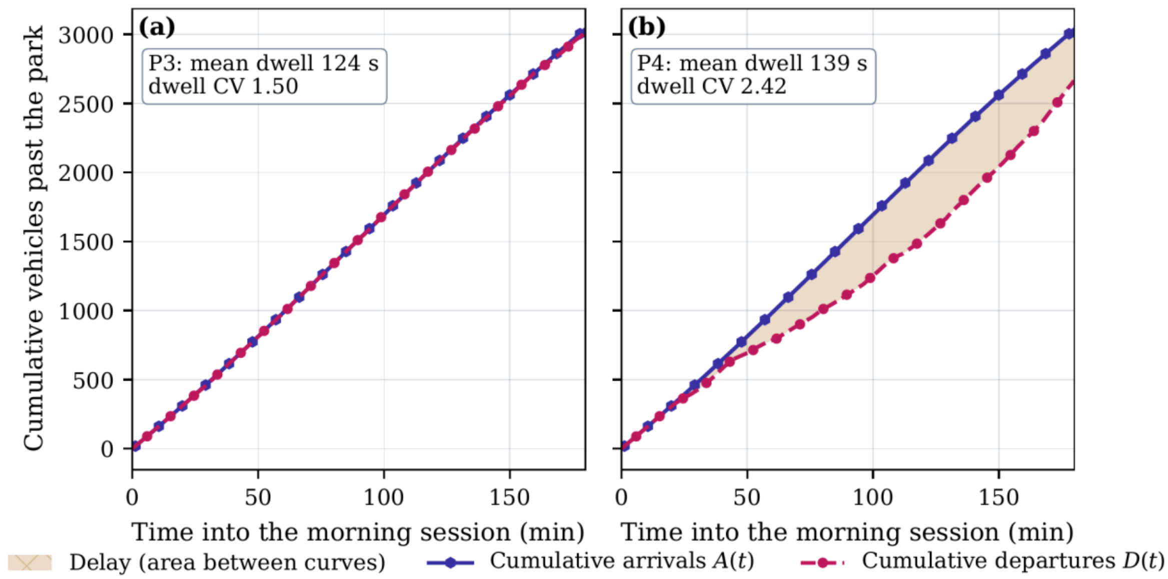

Figure 2. Cumulative arrival and departure curves over the morning session at the two matched-pair parks, P3 in panel (a) and P4 in panel (b). The shaded area between the curves is the delay and the staircase in the departure curve is the sequence of blockages. The two parks are matched on every other input and differ in dwell consistency, with a within-class coefficient of variation about twice as high at P4.

4.2 Separating mean from variance

The experiment the paper turns on holds one moment of the dwell distribution fixed while moving the other. We scale the dwell mean from 40 to 240 seconds with the standard deviation held at 76.5 seconds, which drives the coefficient of variation from 1.91 down to 0.32 along that sweep. We separately scale the variance from 506 to 27,722 s^2 with the mean held at 90 seconds, which moves the coefficient of variation in the opposite direction. Ten levels are run in each sweep, 40 replications at each level, and the elasticity of delay with respect to each quantity is taken from a log-log fit through the replication means. Both sweeps are run under in-lane loading and again under an off-carriageway apron, which is what exposes the configuration effect. The reference park is deliberately sub-saturated, so that the sweeps measure the shape of the dwell distribution rather than the park running out of berths.

4.3 Validation design

A check against a formula that shares the model's assumptions tests the arithmetic; only a check against an independent solution tests the representation. Both are reported separately and only the second is treated as validation.

Against the vertical-queue triangle, which shares the model's assumption that queued vehicles occupy no space, agreement is 0.0008 percent, which confirms the numerics.

Against a spatial shockwave solution, which relaxes that assumption and lets the queue occupy roadway, our delay estimates run 13.4 percent high on average and 24.4 percent high at worst, and the sign is always the same. Every delay figure in this paper is therefore an upper bound carrying that margin.

Figure 3 sets out both halves of that comparison. Panel (a) is the triangular fundamental diagram the shockwave solution is solved on, with the queued states that three encroachment depths produce marked on it; the deeper the encroachment, the lower the discharge plateau and the further right the queued state sits, and at 2.4 m the state is held off the axis only by the residual-capacity floor. Panel (b) plots the cumulative-curve delay against the shockwave delay for all 39 validation cases on log axes. The points track the 1:1 line and lie above it in all 39 cases, by between 2.4 and 24.4 percent. That the bias is one-directional and close to multiplicative matters for what follows: it means the elasticities in Section 5 are largely unaffected, because a multiplicative bias cancels in a log-log slope, while every absolute delay is inflated by roughly the same margin.

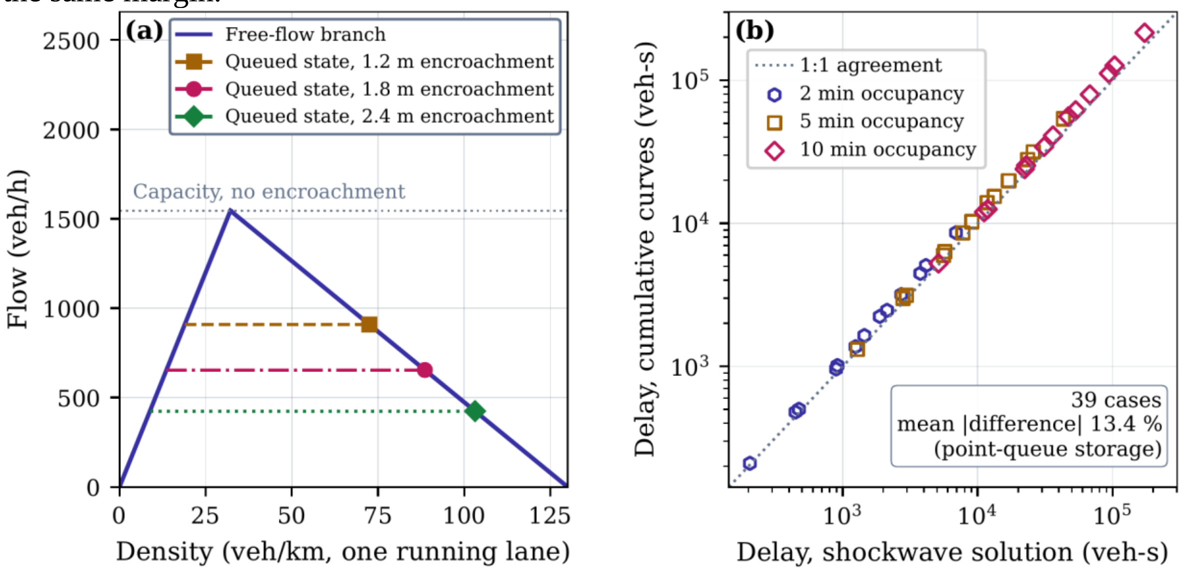

Figure 3. Validation against the spatial shockwave solution. Panel (a) is the triangular fundamental diagram on which the shockwave solution is solved, with the queued state marked for three encroachment depths. Panel (b) compares the cumulative-curve delay with the shockwave delay across all 39 validation cases. Our estimates run 13.4 percent high on average and always in the same direction, which is why every delay in this paper is reported as an upper bound.

4.4 How variation is reported

Two sources of variation act on these delay estimates and they answer different questions. Variation between replications of a single park, which resample that park's own recorded events, measures how much one observation day tells an analyst about that park. Variation between parks measures how far the six sites stand apart from each other. The resampling is carried out within each park separately, so the two never mix. They are carried through the analysis as separate quantities and combined nowhere, because a pooled spread would understate the first and overstate the second.

5 Results

5.1 Dwell time distributions

Table 2 and Figure 4 report the dwell distributions across 2,943 events and 54 session hours. Mean dwell ranges from 80.2 seconds at the smallest park to 302.8 at the largest, with coefficients of variation between 1.71 and 2.10.

One modeling point falls out of this. The pooled dwell distribution at every one of the six parks rejects a lognormal fit. Disaggregated by vehicle class it does not: a lognormal survives the test in all 28 park-by-class cells, and in 8 of those cells a gamma fits closer still and is the distribution Table 2 names. What holds everywhere, then, is a single-family within-class fit rather than the lognormal specifically. Within-class coefficients of variation run between 0.27 and 1.30. The pooled distributions are mixtures, and fitting a single distribution to a park conflates a motorcycle picking up one passenger with a minibus filling twelve seats. Figure 4 shows the per-class panels for that reason.

Table 2. Dwell time by park. Mean and variance are reported separately throughout, because the paper's question is whether they act on delay differently.

Park	n	d^{-} (s)	s^2 (s^2)	s (s)	CV	d_{50} (s)	d_{90} (s)	Fit
P1	453	80.2	19,909	141.1	1.76	40.2	181.3	mixture
P2	549	171.5	105,385	324.6	1.89	62.6	403.8	mixture
P3	581	142.7	59,637	244.2	1.71	57.7	312.3	mixture
P4	608	131.1	75,933	275.6	2.10	45.9	271.4	mixture
P5	493	302.8	272,977	522.5	1.73	106.9	816.9	mixture
P6	259	84.1	23,338	152.8	1.82	37.4	160.5	mixture
By vehicle class, matched pair								
P3 okada	127	23.7	41	6.4	0.27	22.6	33.1	lognormal ($p =$ 0.36)
P4 okada	147	26.4	236	15.4	0.58	22.8	47.1	gamma ($p =$ 0.86)
P3 keke	178	48.2	201	14.2	0.29	46.4	66.5	lognormal ($p =$ 0.95)
P4 keke	177	49.8	1,258	35.5	0.71	39.8	91.2	lognormal ($p =$ 0.74)
P3 taxi	108	79.1	683	26.1	0.33	78.5	115.9	gamma ($p =$ 0.80)
P4 taxi	116	71.0	2,876	53.6	0.75	56.4	126.3	lognormal ($p =$ 0.43)
P3 danfo	144	259.3	18,578	136.3	0.53	231.1	400.3	lognormal ($p =$ 0.88)
P4 danfo	142	268.0	98,092	313.2	1.17	170.2	568.5	lognormal ($p =$ 0.83)
P3 coach	24	1,060.6	252,756	502.7	0.47	943.0	1,660.9	lognormal ($p =$ 0.88)
P4 coach	26	797.4	570,713	755.5	0.95	581.6	1,581.6	gamma ($p =$ 0.99)

Note. $N = 2{,}943$ stopping vehicles over 54 hours, six parks, three sessions each of 3 hours. d^{-} = mean dwell; s^2 = variance; s = standard deviation; CV = coefficient of variation s/d^{-}; d_{50}, d_{90} = median and 90th percentile. The park-level distribution is a mixture across vehicle classes and is not lognormal: a single lognormal is rejected at the 5 percent level for all six parks. Within a class the lognormal is not rejected for any of the 28 park-by-class cells (28/28 with $p > 0.05$), which is why the class components are reported. In 8 of those 28 cells a gamma fits closer than the lognormal does, and the fit column names the closer of the two for each cell, with the p value the Kolmogorov-Smirnov p for the distribution named. Vehicle classes: okada = motorcycle taxi, keke = tricycle, taxi = shared car, danfo = minibus, coach = intercity coach. P3 and P4 are a matched pair: they are close on mean dwell, berth count, arrival rate and class mix, and apart on dwell variability, which within every vehicle class runs about twice as high at P4, which is why the paper uses the contrast between them.

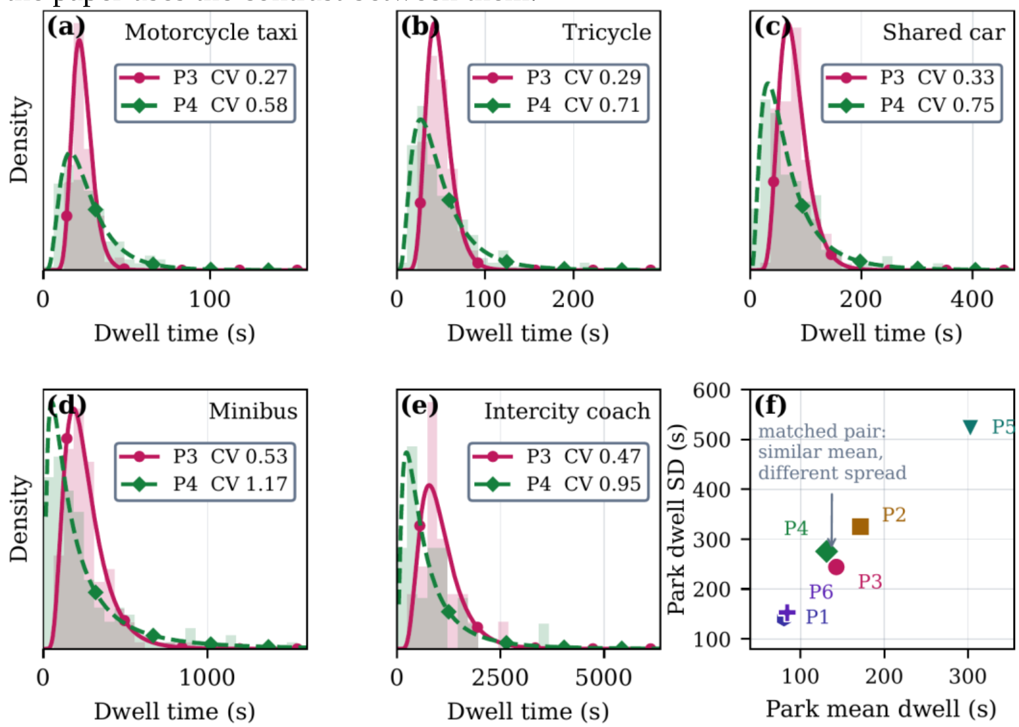

Figure 4. Dwell time distributions by park and vehicle class, with the mean against standard deviation panel. Pooled park distributions are mixtures and reject a lognormal fit at every park; within a vehicle class the fitted distribution is not rejected in any of the 28 cells, a lognormal in 20 of them and a gamma in 8.

5.2 The mean dominates the variance

Table 3 and Figure 5 report the central experiment.

Under in-lane loading, the elasticity of delay with respect to the dwell mean is 1.56, and with respect to the dwell variance 0.30, which is 0.60 with respect to the standard deviation. The mean dominates by a factor of about 2.6. Delay over the mean sweep moves from 3.87 to 66.14 seconds per through vehicle, a multiple of 17.1, while over the variance sweep it moves from 5.83 to 19.00 seconds, a multiple of 3.3.

The ordering runs against the hypothesis the study set out to test. The capacity-drop literature that motivated it (Gu et al. 2011; Yuan et al. 2015) establishes that interruption variability enters effective capacity, and it does here too, with an elasticity of 0.30 that is not zero. It is simply smaller than the mean effect. A park that halves its average dwell buys more than a park that halves its dwell variability, and by a wide enough margin that the regulatory implication the hypothesis carried does not follow.

The superlinearity belongs alongside it, and it needs stating at the right scale. Delay from a single isolated blockage grows with the square of that blockage's duration, which is the deterministic queueing prediction and which the model reproduces in the check reported in Table 6. A whole session is not a single blockage. Stops overlap, berths absorb part of the arrival stream, and the lane runs clear between events, so the response of session delay to the dwell mean is weaker than quadratic: at the measured elasticity of 1.56, doubling the mean dwell multiplies delay by about 2.9 rather than by four. That is still steep enough that shortening dwell returns more than proportionately, which is the point that matters for practice.

Figure 5 shows both sweeps under both configurations. The in-lane curves climb across the whole range in each panel while the apron curves lie flat near 0.1 seconds per vehicle, and the shaded bands, which span the 10th to the 90th percentile across the 40 replications at each level, are wide enough on the in-lane curves that the level at any single point is far less well determined than the slope through them.

Table 3. Response of through-traffic delay to dwell mean and to dwell variance, each varied with the other held fixed, under both park configurations.

	Loading in the running lane		Off-carriageway berths	
Quantity varied	ε	w range (s/veh)	ε	w range (s/veh)
Dwell mean, 40–240 s, SD fixed at 76.5 s	1.56	3.87–66.14	-0.02	0.11–0.10
Dwell variance, 506–27,722 s^2, mean fixed at 90 s	0.30	5.83–19.00	0.00	0.10–0.11
Delay multiple over the mean sweep	17.10×		0.94×	
Delay multiple over the variance sweep	3.26×		1.02×	
Median variance effect across all sweeps	1.70×		1.02×	

Note. Reference park: 8 berths, 1 running lane per direction, through demand 850 veh h^{-1}, park arrival rate 20 veh h^{-1}, encroachment 1.55 m. Each cell is the mean of 40 replications. ε = elasticity of delay with respect to the quantity varied, from a log-log fit through the replication means; w = delay per through vehicle. The variance elasticity is taken with respect to the variance, so the corresponding elasticity with respect to the standard deviation is twice the value shown, 0.60 in the running-lane case. Median variance effect is the ratio of delay at dwell CV 1.45 to delay at CV 0.45, mean held fixed, taken over every point of the four parameter sweeps at which the low-CV case produces measurable delay (26 points per configuration); Figure 7 plots three of those four sweeps. Holding the variance fixed while the mean moves means holding the standard deviation fixed, so the CV necessarily changes along the first row; the reverse holds along the second. What the experiment measures is the magnitude of the variance effect relative to the mean effect, and whether it survives multi-berth operation and a finite observation window.

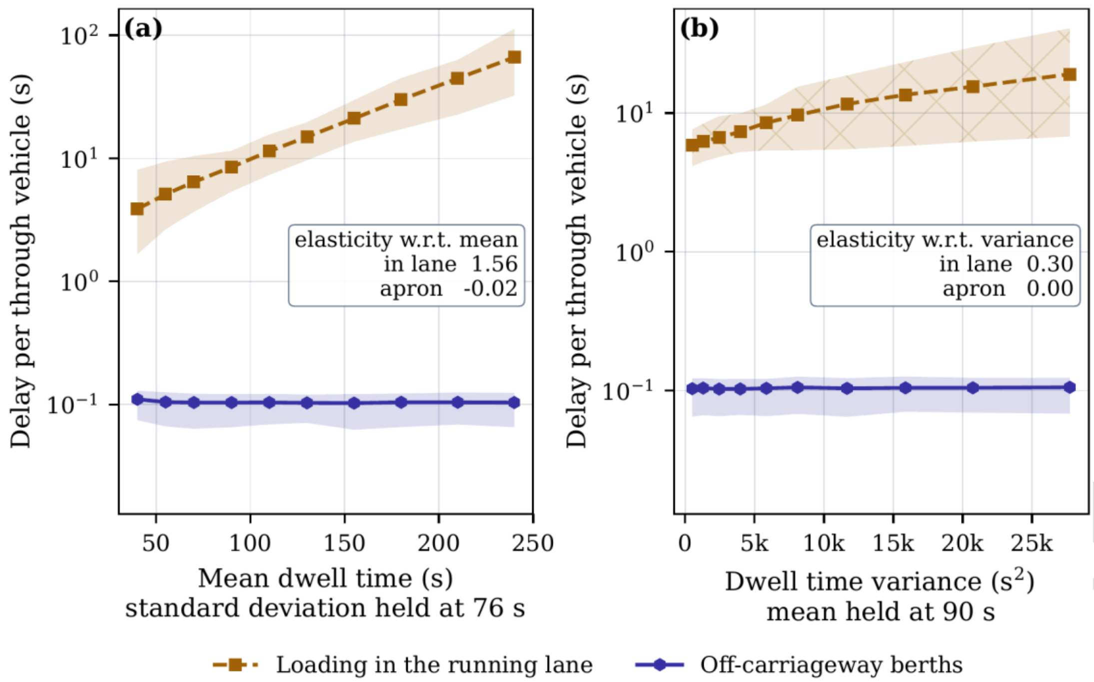

Figure 5. Delay against dwell mean holding the standard deviation fixed, and against dwell variance holding the mean fixed, under in-lane loading and under an off-carriageway apron. Markers are the mean of 40 replications at each level and the shaded band spans the 10th to the 90th percentile across those replications. The apron curves are flat, which is the paper's central result.

5.3 Configuration dominates both

The result that displaces the paper's original question appears when the same sweeps are run under an off-carriageway apron.

Both elasticities collapse. Delay responds to the dwell mean with an elasticity of −0.02 and to the dwell variance with an elasticity of 0.00, and the absolute delay stays between 0.10 and 0.11 seconds per through vehicle across the whole of both sweeps, against 3.87 to 66.14 seconds in the running lane. Once the standing vehicle is mostly clear of the lane, a park may dwell for as long and as erratically as the sweeps allow without imposing measurable delay on through traffic. Adding berths to an in-lane park stops helping above three berths, and the mechanism is not the obvious one. In the berth-count sweep reported in Section 6, going from two berths to three cuts delay from 7.19 to 6.30 seconds per through vehicle and from three to four cuts it to 6.24, after which the curve is flat out to twelve berths: a 13 percent reduction in total, all of it realized by four berths. The first berths remove vehicles that would otherwise wait in the running lane. Beyond that the lane is blocked by the dwelling vehicle itself, so providing somewhere else for the next vehicle to wait restores nothing further.

The berth overhang that separates the two configurations is set for each park from its described setting, as Section 3 sets out, so what the experiment establishes is that the two configurations respond very differently to the same dwell behavior. How many southwestern Nigerian parks sit at each end of that parameter is unknown, and counting them is the first thing this paper recommends.

5.4 Queues and the operating envelope

Figure 6 shows the queue traces and Table 4 the queueing statistics. The sweeps locate the onset of delay coarsely rather than estimating a threshold. Delay is identically zero at every swept encroachment up to 1.2 m and first appears at 1.55 m, which is also the encroachment of the reference case, so 1.55 is the smallest value on the grid at which anything happens rather than an estimated tipping point; the true onset lies somewhere between 1.2 and 1.55 m and the grid does not resolve it. The same holds for through demand, where delay is zero at 550, 650 and 750 vehicles per hour and first appears at 850. Below those points the obstructed capacity still exceeds demand, so the park costs the through stream nothing; a finer grid would place both onsets more precisely and is worth running before either number is used for design.

Above them the delay grows quickly. Nine of the eighteen sessions, all of them peaks, are oversaturated in the ordinary sense: peak-direction demand in those nine sits at between 0.88 and 1.16 of the capacity that survives loading, against 0.34 to 0.84 in the other nine, so arrivals outrun discharge and the queue accumulates rather than settling. Their maximum queues run from 130 to 761 vehicles, and in every case for which Figure 6 plots the trace the queue is still far from clear at the end of the three-hour window. A deterministic queueing delay computed on a queue that does not clear is a function of the length of the observation rather than of the facility, so those nine are flagged in Table 4 and reported as a state rather than as delay estimates. The controlled experiment and the parameter sweeps run at a sub-saturated reference point for that reason.

Table 4. Delay, queue and capacity outcomes from the cumulative-curve bottleneck model, by park and session.

Park	Ses.	q (veh/h)	$\bar{d}$ (s)	CV	w (s/veh)	Q_{max} (veh)	β	ΔC (%)
P1	AM	908	96.8	1.99	307.2†	256	0.42	33.9
P1	MD	529	70.9	1.48	0.1	2	0.13	17.4
P1	PM	985	69.2	1.29	157.9†	130	0.32	27.8
P2	AM	863	170.7	1.87	1,073.1†	714	0.59	51.5
P2	MD	502	171.9	1.93	1.1	8	0.25	40.8
P2	PM	936	172.3	1.90	1,152.1†	761	0.35	47.7
P3	AM	908	123.8	1.50	12.7	26	0.39	30.3
P3	MD	529	142.4	1.68	0.0	1	0.24	24.7
P3	PM	985	162.3	1.81	868.7†	479	0.55	38.2
P4	AM	908	139.0	2.42	1,108.8†	542	0.64	41.8
P4	MD	529	145.2	1.88	0.2	3	0.26	28.6
P4	PM	985	115.4	1.70	529.2†	371	0.41	35.0
P5	AM	953	310.4	1.72	447.7†	394	0.43	38.1
P5	MD	556	334.5	1.97	0.4	3	0.22	31.0
P5	PM	1,034	274.7	1.46	790.5†	505	0.38	35.6
P6	AM	726	70.0	1.41	5.7	25	0.17	19.8
P6	MD	423	98.6	1.99	0.0	0	0.11	18.2
P6	PM	788	88.9	1.87	107.7	167	0.23	25.9

Note. 18 park-sessions. † marks the 9 oversaturated sessions, and the mark is set by queue state rather than by comparison with anything: peak-direction demand in those nine runs at or above the capacity that survives loading, at a ratio of demand to mean obstructed capacity of 0.88 to 1.16 against 0.34 to 0.84 in the other nine, so the queue builds through the session and has not cleared when the three-hour window closes. In those sessions the delay printed is governed by the length of that window rather than by the facility, and it should be read as an indicator of the state rather than as an estimate of delay. Ses. = session, AM = morning peak, MD = midday, PM = evening peak. q = peak-direction through demand; $\bar{d}$ = mean dwell; CV = dwell coefficient of variation; w = mean delay per through vehicle; Q_{max} = maximum queue; β = share of the session for which a vehicle fully obstructs the running lane, either maneuvering or waiting in the lane for a berth; ΔC = mean reduction in directional capacity against the unobstructed value of 1545 pcu h^{-1} per lane. Through demand is calibrated rather than counted at the parks: one two-way figure from the Ilorin arterial counts of Ahmed (2015) stands for each session type, is split 55/45 to the peak direction by assumption, and is then scaled by a per-park factor between 0.80 and 1.05. Three parks take a factor of 1.00, which is why their q column is identical. Delay is computed as the area between cumulative arrival and departure curves at the park frontage; because that is a point-queue construction it overstates delay relative to a spatial solution by about 13 percent, as Table 6 shows.

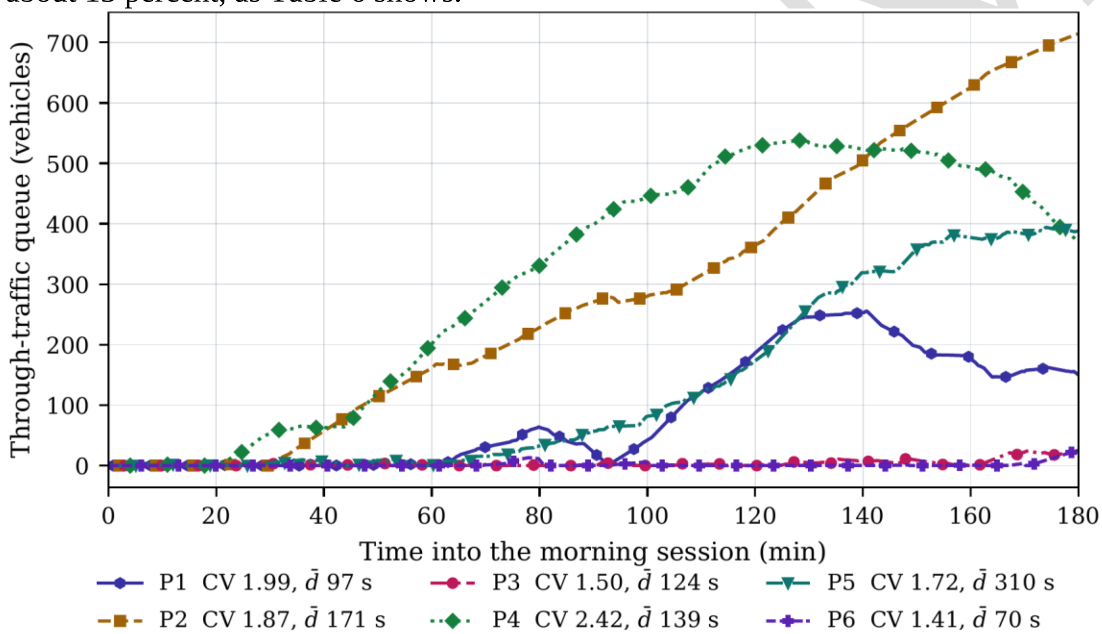

Figure 6. Through-traffic queue length over the morning session at each of the six parks. The four traces that run into the hundreds of vehicles are the morning sessions flagged as oversaturated in Table 4, where the delay is reported as a state rather than as an estimate; the two that stay below thirty vehicles are not.

6 Sensitivity and External Checks

6.1 Parameter sweeps and the variance effect

Figure 7 plots the sweeps and the closing rows of Table 3 summarize them. The variance manipulation raises the dwell coefficient of variation from 0.45 to 1.45 at a fixed mean, a 3.2-fold increase in the coefficient of variation and a tenfold increase in the variance. Under in-lane loading it raises delay by a median factor of 1.70, and the ratio holds between 1.03 and 1.79 across all four sweeps. It falls toward one only at the top of the through-demand sweep, at 1,250 vehicles per hour, where the in-lane case is congested enough that even the consistent-dwell run produces almost as much delay as the erratic one. Under an apron the same manipulation raises delay by a median factor of 1.02, which is to say not at all, with two exceptions that identify the mechanism. At only two berths the factor rises to 2.38 and at a through demand of 1,250 vehicles per hour to 1.33, both cases in which the apron is too small or the arterial too busy for arrivals to be absorbed, so vehicles wait in the lane and the configuration advantage is partly lost.

Variance therefore matters exactly when loading touches the running lane and stops mattering when it does not, which is the configuration result seen from the other side. It also shows that the apron advantage is a matter of degree. The overhang parameter is continuous, the two exceptions sit at the ends of the swept range, and an apron that cannot hold its arrivals behaves much closer to an in-lane park than to a well-sized one.

Table 5. Delay variation between replications of the same park, and delay variation between the six parks, computed and reported as separate quantities.

Park	w^{-} (s/veh)	s_{rep} (s/veh)	CV_{rep}	w_{min} (s/veh)	w_{max} (s/veh)
P1	12.39	23.04	1.86	0.43	118.68
P2	45.42	48.77	1.07	4.35	194.11
P3	0.84	1.01	1.20	0.14	6.26
P4	8.49	17.20	2.03	0.49	73.55
P5	13.44	29.18	2.17	1.79	149.05
P6	0.23	0.51	2.26	0.00	3.16
Between parks	13.47	16.62	1.23	0.23	45.42

Note. Each park's morning peak was recomputed 40 times, resampling with replacement from that park's own recorded events and from its encroachment distribution. w^{-} = mean delay per through vehicle across replications; s_{rep} = standard deviation across replications of the same park; CV_{rep} = the corresponding coefficient of variation; w_{min}, w_{max} = extreme replications. The final row reports variation between parks, computed across the six park means, and it is a different quantity from the rows above it. Replication variation exceeds the mean at every park (CV_{rep} from 1.07 to 2.26), which is the practically important result in this table: a single day of observation at one of these sites carries very little information about its typical delay.

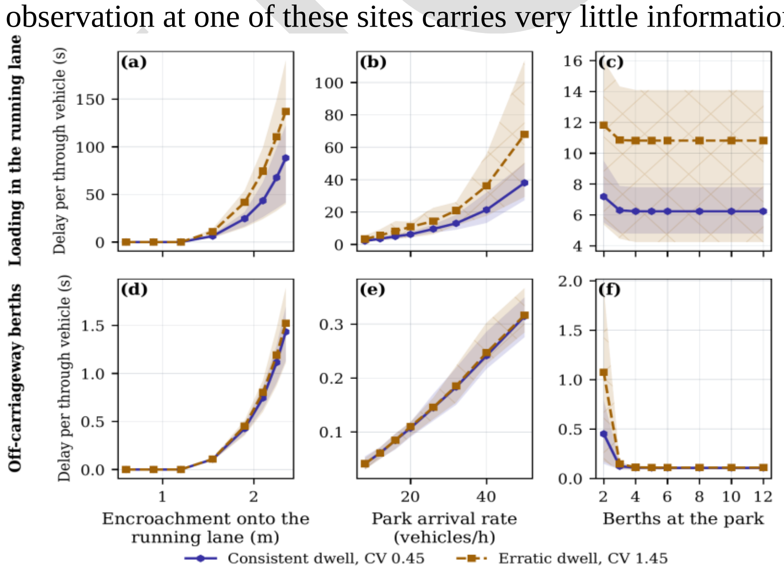

Figure 7. Sweeps over encroachment width, park arrival rate and berth count, under in-lane

loading in the top row and an off-carriageway apron in the bottom row. Markers are the mean of 20 replications at each level and the shaded band spans the 10th to the 90th percentile across those replications. A fourth sweep, over through demand, is reported in the text but not plotted here. The variance effect is stable under in-lane loading and absent under an apron, except at two berths, where the apron cannot hold its arrivals and vehicles wait in the lane.

6.2 Replication variation and between-park variation

Table 5 holds the two quantities. Within a park, the coefficient of variation on delay across replications runs from 1.07 to 2.26. Across parks it is 1.23.

Those numbers carry a warning for anyone designing a measurement of this quantity. The replication spread exceeds the mean at every park, which means a single observation day at a single park is close to uninformative about that park's delay contribution. Repeated days per site are needed, and the sample size for any further survey should be set by the replication variance reported here rather than by the number of parks.

6.3 External benchmarks

Table 6 reports the internal check against the vertical-queue triangle and the independent check against the spatial shockwave solution, together with the wave properties both recover. Three further comparisons against published measurement are not tabulated, because each compares a modeled range against a single published figure rather than against a distribution, and each of the three constrained what the paper reports.

The modeled capacity reduction under loading, 17.4 to 51.5 percent across the eighteen sessions, exceeds the 35 percent on-street parking comparator reported by Gulivindala and Mehar (2018) in 8 of 18 sessions. The model was not retuned, because the sessions that exceed it are exactly the ones that should: a single running lane per direction with a park occupying most of it for more than half the session is not a four-lane divided road at medium side friction. The paper therefore quotes the median reduction, 32.4 percent, rather than the maximum.

The modeled delay per through vehicle exceeds the only published Nigerian delay figure we obtained in nine of the eighteen sessions. Those nine are the oversaturated sessions flagged in Table 4, and they are reported as a state rather than as delay estimates for that reason.

The modeled maximum queue extent, up to 7.6 km in the validation cases, exceeds the length of every study corridor, which run from 3.61 to 5.22 km. The model carries no upstream junction and no spillback, so it lets a queue occupy road that does not exist. That is a genuine failed check rather than a rounding issue, and Section 7.4 states what it costs.

The check that could not be run is the important one. Dwell times could not be compared against any published source, because no study we located reports dwell durations for informal African motor park loading, and the bus dwell studies that would supply the nearest comparator measure a scheduled, berthed service. No substitute was adopted from a different class of facility. The dwell distributions in this paper therefore stand without external corroboration, and every elasticity reported here inherits that.

Table 6. Independent check of the cumulative-curve delay against a shockwave solution on a triangular fundamental diagram.

Check	Mean |diff| (%)	Max |diff| (%)	Cases
Simulation against vertical-queue triangle	0.0008	0.0135	39

Simulation against spatial shockwave solution	13.41	24.44	39
Recovered wave properties			
Back-of-queue wave speed (km/h)	-11.73	-1.13	range
Speed within the queue (km/h)	4.11	20.53	range
Maximum queue extent (km)	0.041	7.600	range
Delay exponent on occupancy duration	$T^{2.000}$		3

Note. 39 cases spanning one and two running lanes, encroachments of 1.2, 1.8 and 2.4 m, demand at 0.55, 0.70 and 0.85 of unobstructed capacity, and occupancies of 2, 5 and 10 minutes. The first row compares the simulated cumulative-curve delay with the analytic vertical-queue triangle. The two share a point-queue storage assumption, so agreement to 0.0008 percent checks the arithmetic rather than the representation. The second row is the independent check: it compares the simulation with the excess vehicle-hours held inside the congested region of the time-space plane, computed from the two shock trajectories and making no vertical-queue assumption. The two differ by 13.4 percent on average and up to 24.4 percent, with the simulation always the larger, because a vertical queue stacks vehicle at a point while the spatial solution lets them occupy road upstream. Every delay reported elsewhere in this paper should therefore be read as an upper bound carrying roughly that margin. Fundamental diagram: triangular, free-flow speed 48 km h^{-1}, jam density 130 veh km^{-1} per lane, capacity 1545 pcu h^{-1} per lane. The T^2 exponent is fitted across the three occupancy durations at one lane, 1.8 m encroachment and demand ratio 0.70, and it is the signature of the bottleneck account: delay from an isolated obstruction grows with the square of its duration, which is why the spread of dwell times enters at all. The largest modeled queue extent, 7.6 km, exceeds the length of every study corridor, which run from 3.61 to 5.22 km; the model has no upstream junction and no spillback, so it lets a queue occupy road that does not exist. This is a failed benchmark and Section 7.4 states what it costs.

7 Discussion

7.1 How far the results can be pressed

The queueing structure is secure. It is standard machinery, it validates against an independent shockwave solution to within a known and one-directional margin, and the quadratic scaling of delay with blockage duration is a property of the mathematics rather than of our parameters.

The direction of the configuration result is secure, though its magnitude is less so. It follows from how much of the running lane a standing vehicle occupies, and that is geometry rather than behavior. But it is a continuous quantity rather than a binary one: the apron case still puts a third of the intrusion in the lane, still takes the whole lane during both maneuvers, and still takes the whole lane whenever the apron is full. What survives is that a park with usable off-carriageway space and enough of it behaves very differently from one without. Where a given park falls on that scale is a measurement nobody has made, and the overhang values that define the two cases here are analyst-set. The mean-versus-variance ratio is the part most exposed. Its inputs are dwell distributions with no published comparator, and while the ordering is stable across our sweeps, the ratio of 2.6 to 1 rest on six parks in one region and could move a long way on a wider sample. One interruption sits outside the model. Passengers boarding and alighting at a park cross the carriageway and accept gaps in the through stream (Pawar and Patil 2015), which the model does not represent and which acts in the same direction as the loading itself.

7.2 Implications for practice

Three kinds of body would have to act on this: the state ministries of works in Osun and Oyo, which hold the corridors the six parks sit on; the local government authorities that license motor parks and could require an apron as a condition of licensing; and the Federal Road Safety Corps where these corridors carry federal traffic.

First, the intervention that works is a lay-by, not a rule. Moving loading off the running lane takes modeled delay to near zero regardless of how long or how erratically vehicles dwell, and no regulation of dwell behavior comes close to that. The sweeps also show that not every park cost something: holding the rest of the reference case fixed, delay is zero below about 1.5 m of encroachment and below about 850 vehicles per hour of peak-direction demand. The corridors worth spending on are the busy ones with a narrow running lane.

Second, size the apron for the demand rather than for the vehicle count. Our sweeps show delay returning when the apron has too few berths or the arterial is busy enough that arrivals overlap, because vehicles then spill back into the lane and the configuration advantage is lost. An undersized apron is closer to no apron than to an adequate one.

Third, do not spend enforcement effort on dwell duration where loading is already off the lane. It buys nothing measurable, and the same effort spent on keeping the apron clear would preserve the benefit that matters. These results do not support treating motor parks as a nuisance to be removed. The parks exist because the transport system runs on them (Cervero and Golub 2007; Behrens et al. 2016), and the finding here is about where loading happens rather than whether it should.

7.3 Transferability

Nothing in the bottleneck construction is particular to a motor park. It needs only that vehicles stop where through traffic runs and that the stopping is intermittent, so a curbside minibus stop, a market frontage where goods vehicles unload and a taxi rank without a bay all sit inside the same account. What does not travel is any number in this paper. The elasticities are properties of the dwell distributions they were fitted to, those distributions have no published comparator, and the gap between the two configurations is set by an overhang value an analyst chose. Read the model somewhere else and the ordering of mean against variance is the claim worth carrying; read the magnitudes across and they are borrowed rather than transferred.

What travels beyond the setting is the modeling choice? Representing an intermittent facility as continuous friction does not merely lose precision, it removes the variable the intervention acts on, because a friction adjustment carries no term for where the loading happens. A capacity manual adjustment cannot distinguish a park with an apron from a park without one, and that distinction is the entire result of this study.

7.4 Limitations

All delays reported here are upper bounds. The check against the spatial shockwave solution puts them 13.4 percent high on average and 24.4 percent high at worst, always in the same direction, because the cumulative-curve construction stacks the queue at a point rather than along the road.

The queue is longer than the road it stands on. The largest modeled queue extent in the validation cases is 7.6 km, against study corridors of 3.61 to 5.22 km, because the model carries no upstream junction and no spillback. On a real corridor the queue would reach that junction first, so part of the delay reported here would be borne by the roads feeding it rather than removed. Resolving that needs a network loading model.

8 Conclusion

We modeled informal motor park loading as an intermittent bottleneck rather than as continuous side friction, across six arterial corridors totaling 25.74 km, and asked which property of the dwell distribution drives arterial delay.

The dwell mean dominates the dwell variance, with elasticities of 1.56 and 0.30 under in-lane loading, so the hypothesis that variability would dominate is rejected. Delay from an isolated blockage grows as the square of its duration, and the measured elasticity across the mean sweep is 1.56, so halving the average dwell cuts delay by about two thirds rather than by half.

Both are dominated by configuration. Under an off-carriageway apron the two elasticities fall to −0.02 and 0.00 and delay stays near 0.1 seconds per through vehicle across both sweeps, against 3.9 to 66.1 seconds when the same distributions are served in the running lane, so the dwell distribution stops mattering once the standing vehicle is mostly clear of the lane. The overhang parameter that produces that contrast is analyst-set, and the advantage is a matter of degree rather than of kind. Variance matters only where loading touches the lane, which is the same result from the other side, and at two berths or at 1,250 vehicles per hour, where the apron cannot hold its arrivals, the variance effect returns.

The recommendation is therefore physical rather than regulatory. A lay-by removes delay that no rule about dwell behavior can, but only if it is sized for its arrivals: at two berths in our sweeps the variance effect returns almost in full, and an apron that cannot hold its arrivals behaves much closer to no apron than to a working one.

Two further measurements would sharpen this most. The first is a count of how many parks in the region have usable off-carriageway space and how much, since that is the parameter the whole result turns on and no survey has yet made it. The second is a longer observation program at three or four parks, a week of continuous recording rather than three sessions, since dwell has no published comparator and the replication spread reported in Table 5 shows that a single observation day per site carries little information about that site.

References

Afolabi, Oluwaseyi Joseph, and Kolawole Taofeek Gbadamosi. 2017. "Impact of Commercial Motorcycle Operation on Urban Mobility in Ogun State, Nigeria." *Logistics & Sustainable Transport* 8 (1): 62–71. https://doi.org/10.1515/jlst-2017-0006.

Ahmed, Y. A. 2015. "Menace of Illegal Motor Parks in Nigerian Urban Environment: Example from Ilorin City." *Journal of Geography and Regional Planning* 8 (2): 37–46. https://doi.org/10.5897/jgrp2014.0466.

Ajayi, Adeyinka Peter. 2017. "Assessment of the Influence of Lagos Bus Rapid Transit Scheme (BRT-Lite) on Road Traffic Crashes (RTC) on Lagos Mainland-Island Corridor." *The Open Transportation Journal* 11 (1): 102–9. https://doi.org/10.2174/1874447801711010102.

Behrens, Roger, Dorothy McCormick, and David Mfinanga. 2016. *Paratransit in African Cities: Operations, Regulation and Reform*. Routledge. https://doi.org/10.4324/9781315849515.

Biswas, Sabyasachi, Souvik Chakraborty, Indrajit Ghosh, and Satish Chandra. 2018. "Saturation Flow Model for Signalized Intersection Under Mixed Traffic Condition." *Transportation Research Record: Journal of the Transportation Research Board* 2672 (15): 55–65. https://doi.org/10.1177/0361198118777407.

Cervero, Robert, and Aaron Golub. 2007. "Informal Transport: A Global Perspective." *Transport Policy* 14 (6): 445–57. https://doi.org/10.1016/j.tranpol.2007.04.011.

Chand, Subhash, Neelam Jain Gupta, and S. Velmurugan. 2017. "Development of Saturation Flow Model at Signalized Intersection for Heterogeneous Traffic." *Transportation Research Procedia* 25: 1662–71. https://doi.org/10.1016/j.trpro.2017.05.216.

Chung, Koohong, Jittichai Rudjanakanoknad, and Michael J. Cassidy. 2007. "Relation Between Traffic Density and Capacity Drop at Three Freeway Bottlenecks." *Transportation Research Part B: Methodological* 41 (1): 82–95. https://doi.org/10.1016/j.trb.2006.02.011.

Dowling, Richard, Alexander Skabardonis, and Vassili Alexiadis. 2004. *Traffic Analysis Toolbox Volume III: Guidelines for Applying Traffic Microsimulation Modeling Software*. FHWA-HRT-04-040. Federal Highway Administration.

Dueker, Kenneth, Thomas Kimpel, James Strathman, and Steve Callas. 2004. "Determinants of Bus Dwell Time." *Journal of Public Transportation* 7 (1): 21–40. https://doi.org/10.5038/2375-0901.7.1.2.

Fernández, Rodrigo. 2010. "Modelling Public Transport Stops by Microscopic Simulation." *Transportation Research Part C: Emerging Technologies* 18 (6): 856–68. https://doi.org/10.1016/j.trc.2010.02.002.

Gross, Donald, John F. Shortle, James M. Thompson, and Carl M. Harris. 2008. *Fundamentals of Queueing Theory*. 4th ed. John Wiley and Sons. https://doi.org/10.1002/9781118625651.

Gu, Weihua, Michael J. Cassidy, and Yuwei Li. 2015. "Models of Bus Queueing at Curbside Stops." *Transportation Science* 49 (2): 204–12. https://doi.org/10.1287/trsc.2014.0537.

Gu, Weihua, Yuwei Li, Michael J. Cassidy, and Julia B. Griswold. 2011. "On the Capacity of Isolated, Curbside Bus Stops." *Transportation Research Part B: Methodological* 45 (4): 714–23. https://doi.org/10.1016/j.trb.2011.01.001.

Gulivindala, Pallavi, and Arpan Mehar. 2018. "Analysis of Side Friction on Urban Arterials." *Transport and Telecommunication Journal* 19 (1): 21–30. https://doi.org/10.2478/ttj-2018-0003.

Han, Youngjun, and Soyoung Ahn. 2018. “Stochastic Modeling of Breakdown at Freeway Merge Bottleneck and Traffic Control Method Using Connected Automated Vehicle.” *Transportation Research Part B: Methodological* 107: 146–66. https://doi.org/10.1016/j.trb.2017.11.007.

Hollander, Yaron, and Ronghui Liu. 2008. “The Principles of Calibrating Traffic Microsimulation Models.” *Transportation* 35 (3): 347–62. https://doi.org/10.1007/s11116-007-9156-2.

Ibitayo, Olurominiyi O. 2012. “Towards Effective Urban Transportation System in Lagos, Nigeria: Commuters’ Opinions and Experiences.” *Transport Policy* 24: 141–47. https://doi.org/10.1016/j.tranpol.2012.07.009.

Kittelson and Associates, Parsons Brinckerhoff, KFH Group, Texas A&M Transportation Institute, and Arup. 2013. *Transit Capacity and Quality of Service Manual*. 3rd ed. TCRP Report 165. Transportation Research Board. https://doi.org/10.17226/24766.

Klopp, Jacqueline M., and Clemence Cavoli. 2019. “Mapping Minibuses in Maputo and Nairobi: Engaging Paratransit in Transportation Planning in African Cities.” *Transport Reviews* 39 (5): 657–76. https://doi.org/10.1080/01441647.2019.1598513.

Koshy, Reebu Zachariah, and V. Thamizh Arasan. 2005. “Influence of Bus Stops on Flow Characteristics of Mixed Traffic.” *Journal of Transportation Engineering* 131 (8): 640–43. https://doi.org/10.1061/(asce)0733-947x(2005)131:8(640).

Krivda, Vladislav. 2013. “Use of Video-Apparatus During Monitoring of Conflict Situations in Road Traffic on Roundabout in the Czech Republic.” *Promet - Traffic&Transportation* 25 (3): 295–303. https://doi.org/10.7307/ptt.v25i3.296.

Kumar, Ajay, and Fanny Barrett. 2008. *Stuck in Traffic: Urban Transport in Africa*. World Bank, Africa Infrastructure Country Diagnostic.

Laureshyn, Aliaksei, Åse Svensson, and Christer Hydén. 2010. “Evaluation of Traffic Safety, Based on Micro-Level Behavioural Data: Theoretical Framework and First Implementation.” *Accident Analysis & Prevention* 42 (6): 1637–46. https://doi.org/10.1016/j.aap.2010.03.021.

Lighthill, Michael James, and Gerald Beresford Whitham. 1955. “On Kinematic Waves II. A Theory of Traffic Flow on Long Crowded Roads.” *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 229 (1178): 317–45. https://doi.org/10.1098/rspa.1955.0089.

Lopez, Clélia, Chuan-Lin Zhao, Stéphane Magniol, Nicolas Chiabaut, and Ludovic Leclercq. 2019. “Microscopic Simulation of Cruising for Parking of Trucks as a Measure to Manage Freight Loading Zone.” *Sustainability* 11 (5): 1276. https://doi.org/10.3390/su11051276.

Mallikarjuna, C., A. Phanindra, and K. Ramachandra Rao. 2009. “Traffic Data Collection Under Mixed Traffic Conditions Using Video Image Processing.” *Journal of Transportation Engineering* 135 (4): 174–82. https://doi.org/10.1061/(asce)0733-947x(2009)135:4(174).

Meng, Qiang, and Xiaobo Qu. 2013. “Bus Dwell Time Estimation at Bus Bays: A Probabilistic Approach.” *Transportation Research Part C: Emerging Technologies* 36: 61–71. https://doi.org/10.1016/j.trc.2013.08.007.

Milkovits, Martin N. 2008. “Modeling the Factors Affecting Bus Stop Dwell Time.” *Transportation Research Record: Journal of the Transportation Research Board* 2072 (1): 125–30. https://doi.org/10.3141/2072-13.

Mobereola, Dayo. 2009. *Lagos Bus Rapid Transit: Africa’s First BRT Scheme*. SSATP

Discussion Paper No. 9. Sub-Saharan Africa Transport Policy Program, World Bank.

Ndibatya, Innocent, and M. J. Booysen. 2020. "Minibus Taxis in Kampala's Paratransit System: Operations, Economics and Efficiency." *Journal of Transport Geography* 88: 102853. https://doi.org/10.1016/j.jtrangeo.2020.102853.

Newell, G. F. 1993. "A Simplified Theory of Kinematic Waves in Highway Traffic, Part II: Queueing at Freeway Bottlenecks." *Transportation Research Part B: Methodological* 27 (4): 289–303. https://doi.org/10.1016/0191-2615(93)90039-d.

Newell, Gordon Frank. 1982. *Applications of Queueing Theory*. 2nd ed. Chapman and Hall.

Nguyen, Huynh Duc. 2016. "Saturation Flow Rate Analysis at Signalized Intersections for Mixed Traffic Conditions in Motorcycle Dependent Cities." *Transportation Research Procedia* 15: 694–708. https://doi.org/10.1016/j.trpro.2016.06.058.

Ogunrinola, I. Oluranti. 2011. "Informal Self-Employment and Poverty Alleviation: Empirical Evidence from Motorcycle Taxi Riders in Nigeria." *International Journal of Economics and Finance* 3 (2): 176–85. https://doi.org/10.5539/ijef.v3n2p176.

Olapoju, Olabisi Michael. 2016. "Culture of Distracted Driving Among Intra-City Commercial Bus Drivers in Ile-Ife, South-Western Nigeria." *Transportation Research Part F: Traffic Psychology and Behaviour* 42: 425–32. https://doi.org/10.1016/j.trf.2016.07.006.

OpenStreetMap contributors. 2020. *OpenStreetMap: Arterial Geometry for Osogbo, Ibadan and Ilesa*. Retrieved from the Overpass API, https://overpass-api.de.

Osoba, Samson Babatunde. 2012. "Appraisal of Parking Problems and Traffic Management Measures in Central Business District in Lagos, Nigeria." *Journal of Sustainable Development* 5 (8). https://doi.org/10.5539/jsd.v5n8p105.

Park, Byungkyu (Brian), and Hongtu (Maggie) Qi. 2005. "Development and Evaluation of a Procedure for the Calibration of Simulation Models." *Transportation Research Record: Journal of the Transportation Research Board* 1934 (1): 208–17. https://doi.org/10.1177/0361198105193400122.

Patkar, Manish, and Ashish Dhamaniya. 2020. "Developing Capacity Reduction Factors for Curbside Bus Stops Under Heterogeneous Traffic Conditions." *Arabian Journal for Science and Engineering* 45 (5): 3921–35. https://doi.org/10.1007/s13369-019-04309-4.

Pawar, Digvijay S., and Gopal R. Patil. 2015. "Pedestrian Temporal and Spatial Gap Acceptance at Mid-Block Street Crossing in Developing World." *Journal of Safety Research* 52: 39–46. https://doi.org/10.1016/j.jsr.2014.12.006.

Richards, Paul I. 1956. "Shock Waves on the Highway." *Operations Research* 4 (1): 42–51. https://doi.org/10.1287/opre.4.1.42.

Robinson, Steve. 2013. "Measuring Bus Stop Dwell Time and Time Lost Serving Stop with London iBus Automatic Vehicle Location Data." *Transportation Research Record: Journal of the Transportation Research Board* 2352 (1): 68–75. https://doi.org/10.3141/2352-08.

Roca-Riu, Mireia, Jin Cao, Igor Dakic, and Monica Menendez. 2017. "Designing Dynamic Delivery Parking Spots in Urban Areas to Reduce Traffic Disruptions." *Journal of Advanced Transportation* 2017: 1–15. https://doi.org/10.1155/2017/6296720.

Salau, Taofiki. 2015. "Public Transportation in Metropolitan Lagos, Nigeria: Analysis of Public Transport Users' Socioeconomic Characteristics." *Urban, Planning and Transport Research* 3 (1): 132–39. https://doi.org/10.1080/21650020.2015.1124247.

Sugiarto, Sugiarto, and Thirayoot Limanoond. 2013. "Impact of on-Street Parking on Urban Arterial Performance: A Quantitative Study on Travel Speed and Capacity

Deterioration." *Aceh International Journal of Science and Technology* 2 (2): 63–69. https://doi.org/10.13170/aijst.0202.04.

Tirachini, Alejandro. 2013. "Bus Dwell Time: The Effect of Different Fare Collection Systems, Bus Floor Level and Age of Passengers." *Transportmetrica A: Transport Science* 9 (1): 28–49. https://doi.org/10.1080/18128602.2010.520277.

Toledo, Tomer, and Haris N. Koutsopoulos. 2004. "Statistical Validation of Traffic Simulation Models." *Transportation Research Record: Journal of the Transportation Research Board* 1876 (1): 142–50. https://doi.org/10.3141/1876-15.

Transportation Research Board. 2010. *Highway Capacity Manual 2010*. Transportation Research Board of the National Academies.

Yuan, Kai, Victor L. Knoop, and Serge P. Hoogendoorn. 2015. "Capacity Drop." *Transportation Research Record: Journal of the Transportation Research Board* 2491 (1): 72–80. https://doi.org/10.3141/2491-08.