Public and Private Transit: Evidence from Lagos

Daniel Björkegren, Alice Duhaut, Geetika Nagpal, Nick Tsivanidis

August 2026

Private minibuses dominate transport in developing-country cities, serving 62% of motorized trips in Lagos. We collect panel data measuring how minibuses respond to the rollout of a public bus network. When public buses enter a route, minibuses depart less frequently, profits fall, and drivers switch to connected routes, lowering prices. A randomized experiment measures how commuters trade off prices and wait times. The private response makes treated-route commuters wait longer but lowers prices on connected routes. Overall, commuters benefit, and over a quarter of commuter welfare gains arise from the private response. Driver welfare losses equal 66% of these gains.

Björkegren: Columbia University, dan@bjorkegren.com. Duhaut: World Bank, aduhaut@worldbank.org. Nagpal: World Bank, gnagpal@worldbank.org. Tsivanidis: UC Berkeley, ntsivanidis@berkeley.edu. We thank many colleagues and seminar participants for helpful comments and discussions, and many research assistants for excellent work. This work would not have been possible without the tireless work of field managers Nura Sambo Umar and Segun Adebamiji. This paper is a product of the staff of the International Bank for Reconstruction and Development/the World Bank. The findings, interpretations, and conclusions expressed do not necessarily reflect the views of the World Bank, the Executive Directors of the World Bank or the governments they represent. The World Bank does not guarantee the accuracy of the data included in this work. This work was supported by the World Bank's Research Support Budget, the World Bank Knowledge for Change Programme, the European Union's Directorate-General for International Partnerships (DG INTPA) and the UK International Development from the UK government under the World Bank's Umbrella Facility for Impact Evaluation, the Clausen Center at UC Berkeley, IGC, J-PAL King Climate Action Initiative, the Brown University Seed Fund, the Weiss Fund, and STEG. The experimental portion of this project was preregistered in the AEA RCT registry (AEARCTR-0010283).

1. Introduction

Developing country cities are growing rapidly, and their governments are making large investments in public transit systems.[1] But these cities are already served by private minibus networks. Public systems are unlikely to displace private systems completely: for example, Dar es Salaam's modern BRT accounts for only 1.3% of trips, while minibuses still carry around 60%. Many evaluations focus on the new public system, measuring the time saved (Small and Verhoef 2007) or the accessibility created (Tsivanidis 2026; Zárate 2024). But such evaluations miss the response of the private incumbents. If public entry changes minibus fares, service frequency, or where drivers operate, then building public transit affects commuters who never ride it. It also may affect the drivers it competes with.

We study these interactions in Lagos, sub-Saharan Africa's largest city with a population of 22 million. Before the reform we study, private minibuses (known as danfo) carried 62% of the motorized trips, on 759 routes. Starting in 2019, the government rolled out a new public bus system in waves, resulting in 64 public bus routes served by 820 buses.[2] We analyze this staggered rollout using newly collected data on the minibus market. This allows us to measure responses both on routes that the government entered, and on other routes linked to them through common terminals.

We find that public entry reallocates private transit supply. On routes the government enters, minibus departures fall by 22%, increasing passenger waiting times, and fares fall by a suggestive 5–10%. Drivers shift away from these routes towards connected routes. Routes sharing a terminal with newly entered routes suggestively experience more waiting minibuses and fares that are around 8% lower, despite receiving no new public service themselves. Because there are many such routes, these spillovers are quantitatively important.

We clarify these forces using an equilibrium model. Because minibuses depart only when full, when they serve fewer passengers, departures are less frequent and passengers spend more time waiting. Drivers respond to changes in earnings by reallocating across routes. A driver's association sets fares, and cuts them when capacity is idle. The welfare impacts of introducing public transit are therefore ambiguous: lower minibus fares benefit commuters, while reduced departure frequency makes them wait longer. Driver reallocation spreads these effects beyond routes receiving public service. The model describes sufficient statistics for the impacts on commuter and driver welfare. These depend on the adoption share of the new service; how the new service changes prices, wait times, and minibus driver trips; how commuters value those changes; how commuters substitute between modes; and how drivers substitute across routes.

To measure these effects, we had to construct data on a sector that is largely invisible to the

[1] 60% of World Bank spending within cities is on transport projects (The World Bank (2025a,b)).

[2] Public buses are larger and newer than the existing minibuses, but run mostly on the same roads and have the same travel times in vehicle. They differ mainly in fares, wait times, comfort, and safety, as shown in Appendix Table S1.

government. Public ridership is recorded whenever a commuter boards a public bus, but for the minibus network the government had no current map, let alone data on how its fares or service changed over time. We mapped Lagos's minibus network by having enumerators identify and ride 759 routes, covering nearly 30,000 km. We also stationed enumerators at terminals to measure fares, driver queues, and departures on 278 routes over 13 rounds and 15 months, and followed 854 minibus drivers over five survey rounds. Our resulting panel data contains *treated* routes directly entered by the government, *connected* routes sharing a terminal with a newly entered route, and other routes in the minibus network.

Several empirical patterns confirm that driver reallocation, not demand, drives network spillovers. When a route receives public transit, drivers complete fewer trips, yet minibus queues shorten; if drivers stayed, queues would lengthen. On connected routes, departures are unchanged. Since buses depart only when full and there is almost always a queue of waiting buses, the departure rate reflects passenger arrivals: neither lower fares nor a new public transit connection generates additional minibus demand on these connected routes. In-vehicle travel times are unaffected: the new buses share the same roads, and Google Maps data show no change in road speeds. We estimate these effects from the staggered rollout using two regression designs. A within-terminal design compares treated routes with other routes at the same terminal. It absorbs possible differential trends at the larger terminals the government chose to enter, but is biased if treatment spills onto the comparison routes. A second design compares routes that become treated or connected with the rest of the minibus network.[3] Comparing the two, we find the bias is small on any individual route, but spillovers are economically important in aggregate because connected routes outnumber treated ones.

A natural concern is that the government may have entered routes whose outcomes were already changing. We exploit routes that were included in government rollout plans but never opened, including eight routes whose October 2020 launch was unexpectedly canceled when a terminal was burned during a protest. These placebo routes show no effects around their planned opening dates. Results are generally robust to wild bootstrap inference, which accounts for the small number of treated routes, though some estimates lose precision.

Valuing these changes requires knowing how commuters trade off lower fares against longer waits. We measure how much commuters dislike waiting in a field experiment designed to overcome two selection problems. First, we obtain broad representation by recruiting commuters at home since recruiting at bus stops would miss those in a hurry; we then offer daily randomized payments to wait at their stop, using a platform accessible from basic mobile phones. Second, as in many digital measurements of waiting, participation itself is endogenous. In our experiment, commuters are less likely to participate when rushed, so we separately randomize a check-in payment that shifts participation. Our selection-corrected estimate suggests commuters dislike

[3] Both specifications measure the effects relative to other routes in the network; they represent the total causal effect under the assumption that treatment affects only routes sharing an endpoint. A robustness test finds that results are unaffected by overlap along the route itself.

waiting at a value of ₦18.94 per minute ($1.42 per hour), 3.0 times the average wage. Ignoring endogenous participation understates this value by 56%; a conventional stated-preference exercise overstates it by 145%. We also estimate medium-run price sensitivity from abrupt changes in public bus fares.

We combine these estimates to measure the effects of introducing public transit on the welfare of commuters and minibus drivers, accounting for the response of the private market. On routes the government enters, commuters gain $0.25 per person per day directly from the new buses ($1.43m per month in aggregate). The private response reduces this by $0.02 ($0.09m per month), as commuters dislike the longer waits about as much as they value the lower fares. On connected routes, where the government added no service, commuters gain $0.03 per day ($0.59m per month) solely from reduced minibus fares. Across the treated and connected routes, commuters gain $1.93 million per month from the introduction of public transit, or $0.07 per commuter per day. The private response accounts for 26% of this commuter gain.[4] Drivers lose an average of $4.40 per day, around 37% of baseline surplus. These losses are the same on treated and connected routes, because driver mobility equalizes expected profits. Aggregate driver losses are $1.3 million per month, about 66% of the commuter gains.

Two implications follow. First, evaluating only the new public service misses substantial commuter benefits generated through the incumbent market: in Lagos, many of these gains arise on routes the government never enters and take the form of lower prices rather than improved access. Second, the reform redistributes surplus from incumbent drivers toward commuters. The world's fastest-growing cities are planning transit systems in markets already served by transit. Our results suggest these designs should carefully consider the interplay of public and private transit.

Related Literature. Our main contribution is to a body of previous work measuring the impacts of transit reforms in cities. Existing work focuses primarily on centralized public transit (Gibbons and Machin 2005; Glaeser, Kahn, and Rappaport 2008; Tsivanidis 2026; Kreindler et al. 2023; Almagro et al. 2025), or roads (Baum-Snow 2007). However, much of the world's population relies on decentralized private transit. There are various case studies of these systems (Cervero and Golub 2007), but few well-identified studies. An emerging literature has begun to investigate various aspects of decentralized minibus systems. The closest to ours is Conwell (2026), which conducts extensive data collection on minibus operations in Cape Town and develops an equilibrium model of queueing and waiting in the minibus industry, in which bus supply insures commuters against demand shocks, and finds that decentralized provision comes close to the planner's optimum. Mbonu and Eaglin (2024) analyses the impacts of fragmentation between drivers' associations in Johannesburg. Kelley, Lane, and Schönholzer (2021) introduces a monitoring technology to minibus owners in Kenya and finds that it improves contracting

[4]Our baseline results come from a multinomial logit model; results are qualitatively similar under a nested logit model.

with drivers. We contribute to this literature by measuring the response of a minibus network to the rollout of a new public network using a quasi-experiment, developing a sufficient statistics approach to measure its welfare implications using a queuing model of transit markets, and a randomized experiment to measure how commuters value price and wait time for transport options.[5]

Another literature considers the impacts of public good provision in developing country contexts where the good or service is already provided by private firms.[6] Recent work focuses on insurance (e.g. Mobarak and Rosenzweig (2013)) and education, where public sector investments can crowd-in or crowd-out private schools in different contexts (Dinerstein, Neilson, and Otero 2020; Andrabi et al. 2023; Dinerstein and Smith 2021). In contrast to these settings, transit is networked, representing multiple markets interlinked through driver and commuter choices. The interaction between private and public providers can be different and more nuanced in such settings; for example, we find crowd out on treated routes, and crowd in on connected routes. In general, there are many spillovers between public and private decisions in urban settings. The discrete spatial structure in transit allows us to measure such spillovers.

Another contribution is that our paper is one of few that causally estimates commuters' value of time in a developing country setting, and it demonstrates how to use selection correction to solve a common problem with measurement. Recent work in this literature has used experiments to identify how travelers trade off time and money. Kreindler (2024) conducts a field experiment that measures value of time and scheduling for drivers in Bangalore. There is a large literature studying travelers' value of time in wealthier settings, which commonly differentiates between waiting and in-vehicle travel time (Small and Verhoef 2007). Because the transit modes share the same roads, they differ only in waiting time, so it is the only time valuation we need to measure. Recent estimates of the value of wait time come from smartphone ride-sharing platforms (e.g., Goldszmidt et al. 2020; Buchholz et al. 2025), which face two selection problems in this setting. The first is selection into the platform: the apps, and the smartphones they require, are not widely used in developing countries (Milusheva, Bjorkegren, and Viotti 2021). The second is selection into when to use it: riders hail when they are in a hurry, biasing estimates of the value of time. Our design addresses both by allowing travelers with basic phones to participate and correcting for endogenous participation. Engineering details for the accompanying software are provided in a companion paper (Björkegren et al. 2026).

[5] In studying decentralized transport, we also connect with a recent strand of work studying this in contexts such as shipping and taxicab markets (Buchholz 2022; Brancaccio, Kalouptsidi, and Papageorgiou 2020; Fréchette, Lizzeri, and Salz 2019). Also related is Roberts et al. (2024), which finds that requests for motorcycle ride hailing increases in Jakarta around a recently opened mass transit line. An earlier, mostly theoretical literature considers how minibuses might interact with public transit in wealthy cities (Walters 1982; Mohring 1983; Bly and Oldfield 1986).

[6] There is a long tradition of similar work in developed countries, such as Brown and Finkelstein (2008) in health insurance.

2. Public Transit in Lagos

Prior to the period we study, the public transit system consisted of a single Bus Rapid Transit (BRT) route and a handful of regular routes that phased in and out of existence. As of 2009, only 5% of trips in Lagos were by government public transit (Lagos Travel Survey, described below).

As part of a modernization plan, the Lagos state government developed a Bus Reform Initiative to introduce a new public bus network, starting in 2019. The network served 64 routes, mapped in Figure 1A. The routes were served by 820 new buses, most of which have a capacity of 70 passengers; a smaller number have a capacity of 30. These buses offer amenities such as air conditioning. Users pay for service by tapping a prepaid electronic 'Cowry' card on entry, which can be topped up at stations (with cash or card) or online.

The Lagos Metropolitan Area Transport Authority (LAMATA) contracted out operations of the routes to four operators. In addition to establishing bus routes, new terminal infrastructure was built at major interchanges. Our study focuses on the large buses connecting these terminals on high demand axes. Most routes, and the routes we study, are standard buses without dedicated lanes (shown in solid lines in Figure 1A). Public buses travel alongside existing traffic, and as a result have similar travel times. The plan also included a handful of bus rapid transit routes with dedicated lanes (shown in dotted lines in Figure 1A). Many passengers ride buses from the origin to the final terminal: 57% of boardings are at the origin of the route.

Of the 64 routes that have opened, 35 launched during our sample period (late 2020 to the end of 2021). Route opening plans changed frequently due to the COVID-19 pandemic, operational challenges, and shocks such as the burning of a terminal during a protest in 2020.[7]

2.1. Data

We rely on two data sources to measure the public network.

Electronic Ticketing. We observe every tap in to the system with the associated fare, timestamp, and origin (though not destination), from August 2020 through February 2024. From this data we derive passenger volumes and estimate price sensitivity.

Opening Plans. To track the changes in public rollout plans, we collected a sequence of 22 route opening plans from LAMATA at regular dates between March 2020 and December 2021.

[7]The bus expansion we study was followed by two other expansions. One was a series of 'First and Last Mile' routes connecting terminals to nearby neighborhoods, operated by smaller 7-seater vehicles (some via narrow and muddy roads). A second was two rail lines connecting Lagos Island to the west and north, in September 2023 and 2024 respectively. The rail lines and most of the smaller routes opened after our period of study.

3. Private Transit in Lagos

As in many other fast growing cities, transportation in Lagos is dominated by private minibuses, which are locally called danfo (similar to matatus, or dala dalas). In 2009, 62% of motorized trips were via private minibus, as reported by the Lagos Travel Survey.

The minibus network is organized around major terminals (motorparks), a subset of which are near public terminals. Although public buses run on existing routes, they serve different stops—with public and private stops for the same origin or destination averaging 180m apart.[8]

Most drivers begin a journey by queueing at a terminal, to depart along a particular route. Once the vehicles in the front of the queue have departed and a driver's vehicle is next in line, it may load passengers. Minibuses almost always wait until full to depart. Photos of queues are shown in Appendix Figure S1. The most common size of minibus holds 14 passengers, although capacity can range from 7 to 22 passengers.[9]

The system tends to operate from terminal to terminal. Since minibuses tend to leave origin terminals full, passengers may not be able to board at intermediate stops.[10]

The government permits two bus drivers' associations to regulate the industry, setting fares and collecting fees, managing shared infrastructure, and enforcing rules (CPCS 2024).[11] However, these groups have evolved into complex entities characterized by revenue extraction, political patronage, and street-level authority. At bus stops outside terminals there is less oversight, and drivers have more leeway on setting prices in response to factors such as demand or weather.

3.1. Data

There are few existing statistics on this incumbent minibus network. Our main task is to collect data to better understand Lagos' private transit network—and how it responds to the expansion in public transit. Because of logistical reasons, not all data could be gathered simultaneously; a timeline of data collection is provided in Appendix Figure S2.

Bus Stop Observations. We hired enumerators to count bus departures and passengers throughout the network as the public system was rolled out, from November 2020 until December 2021. Our team collected 13 rounds of data, spaced approximately every month. At

[8]This is computed for 36 pairs of public and private stops, for which we were able to locate both stops on Google Earth.

[9]For simplicity we refer to all private shared transit vehicles as minibuses, including larger size buses (locally called molue) and smaller size vehicles (locally called korope).

[10]Data collected during our network mapping suggest 92% of passengers along minibus routes board at the origin terminal.

[11]The National Union of Road Transport Workers (NURTW, with Lagos chapter recently renamed to Lagos State Parks and Garages Management) controls most terminals in Lagos; some are also controlled by the Road Transport Employers' Association of Nigeria (RTEAN, with Lagos chapter recently renamed to Lagos State Park and Garage Administration).

Figure 1. Transit in Lagos

A. New Public Transit Network

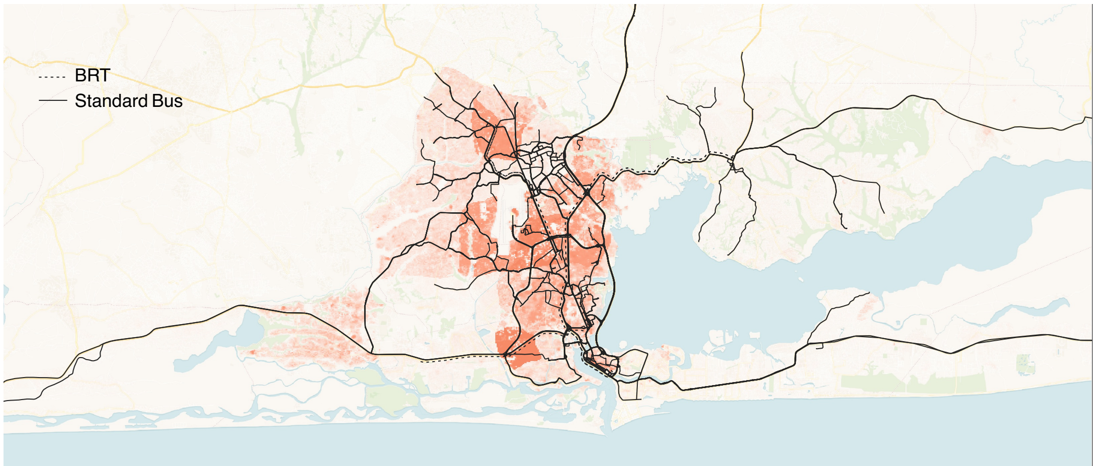

B. Private Transit Network

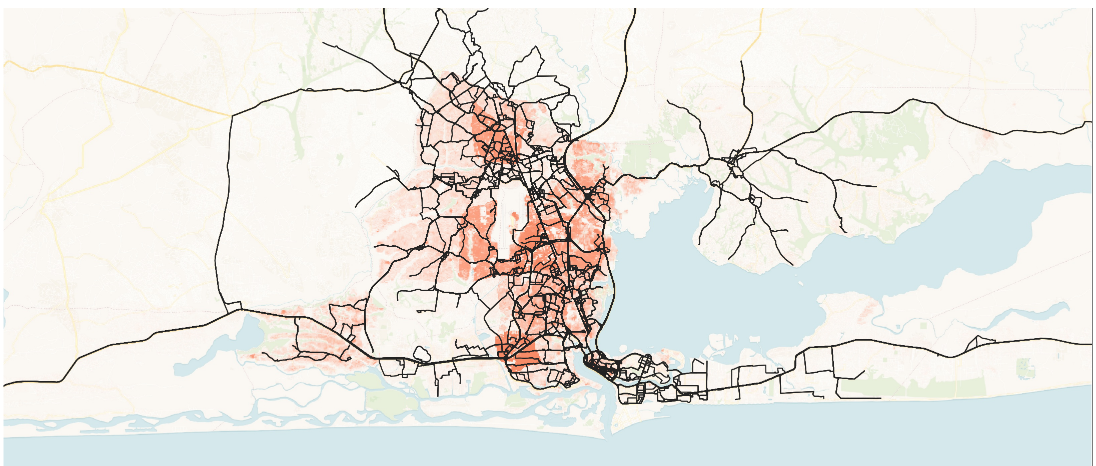

C. Sampled Private Transit Routes

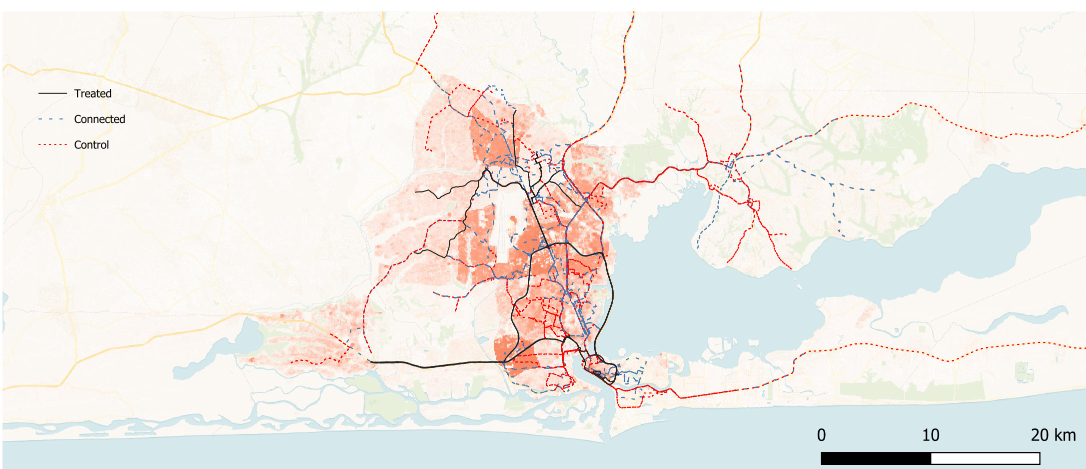

Note: The first panel shows the new public transit routes. Bus rapid transit (BRT) routes are shown with a dotted line. The second panel shows the existing private transit routes, as discovered in our network mapping exercise. The final panel shows the private transit routes we monitored for our study, including routes treated by public transit, connected routes that share an endpoint with a public transit terminal, and control routes that do not share an endpoint with public transit terminals. Note that routes that are not connected at endpoints may partially overlap on roads, but it tends to be difficult to transfer at intermediate stops. This final panel plots the 83% of sampled routes that we were able to match to the GIS database of routes. Base layer color shows population density.

terminals, enumerators collected data on departures, fares, and driver queues in half-hourly windows during morning peak, midday offpeak, and afternoon peak times.[12]

Our sampling frame used the government's rollout plan as of October 2020, and focused on routes it planned to open in the short term (by January 2021) as well as those it wanted to be part of its eventual network. We first classified private routes into three categories: *treated* routes that share both endpoints with a public bus route that was planned to open within our observation period, *connected* routes that share one endpoint, and *control* routes that share neither endpoint. A stylized illustration of these definitions is provided in Appendix Figure S3. We collected data on all planned treated routes; of the remainder, roughly two-thirds were targeted as connected and one-third as control. Around 85% of our sample came from this plan; the remaining 15% were routes departing from the same terminals that our enumerators were able to observe from a single position.

The government's plans experienced many changes during our sampling period. In total we collected data on 278 routes, of which 13 ended up as treated (representing 37% of all 35 routes receiving public transit during this period).[13] However, some routes planned for opening during our period did not see public service start: 58 of our sampled routes appeared in any deployment plans for 2020-2021, and 41 appeared in the March 2020 deployment plan closest to opening. [14,15] We later use changes in opening plans in placebo tests in Section 5.4. In addition to the terminal observations, we also collect similar data at 79 bus stops (where our enumerators could not identify the exact route, since it is not always displayed on vehicles).

Figure 1C maps sampled routes according to these categories.[16] Figure 2 shows the timing of public route openings on our sampled routes, and the dates of our bus stop observations.

Driver Surveys. We conducted a panel survey of 854 minibus drivers, starting with a baseline survey and following up in four subsequent rounds, covering November 2020 through February 2022. Drivers were sampled while waiting in queues at terminals, and outside terminals, on treated, connected, and control routes. The baseline survey was administered during recruitment, and drivers were asked for permission to participate in follow-up surveys by phone. This sampling method is less likely to capture drivers who often skip terminals (drive sole). We correct for this by asking the drivers who we survey what fraction of the time they skip terminals,

[12] Corresponding to 7-9 a.m., 1-2 p.m., and 4-6 p.m., respectively. Queue lengths and fare were recorded as of the start of the observation window; departures were measured cumulatively over the window.

[13] The remainder were not initially planned to open within our sample period and were not sampled. We consider routes as opened if they served passengers for at least 3 months. While we only collect data on a subset of treated routes, we use the full set of opened public routes to measure private routes' public transit connections.

[14] We had planned to begin data collection in February 2020, which was delayed due to the COVID-19 pandemic. Analysis of Google mobility data (Figure S4) showed that travel patterns reverted to January 2020 levels by late July 2020, prompting the collection of baseline data in October 2020.

[15] Appendix tables S2 and S3 compare routes that are treated to those that were planned to open in the short term as well as those part of the eventual network plan.

[16] For the 83% of sampled routes that we could match to our GIS data.

Figure 2. Sampled Routes Served by Public Transit and Bus Stop Observations

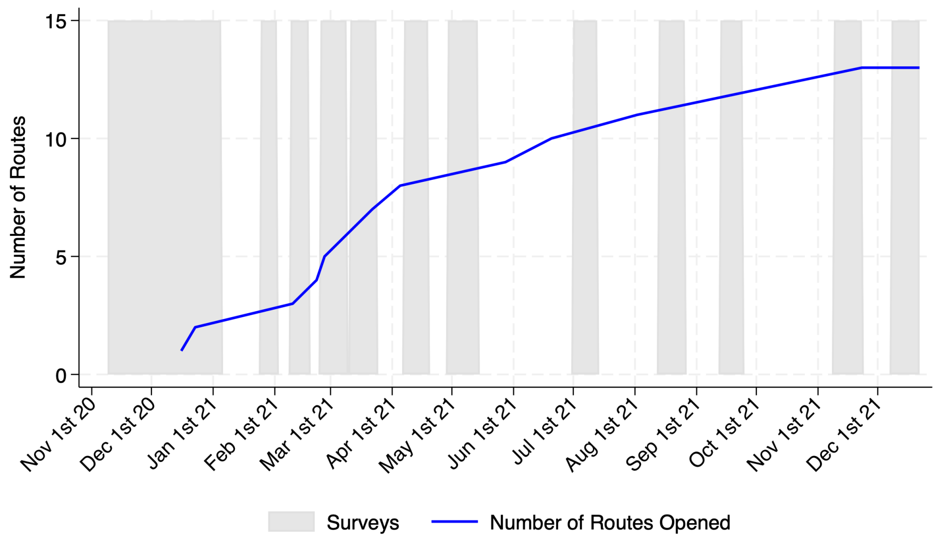

Note: Plots the number of routes served by public transit in our monitored sample. The timing of our bus stop observations are shaded as vertical boxes.

and use that to reweight the sample, as described in Appendix S3.1.[17]

Wait Time Experiment. We additionally ran an experiment on sensitivity to wait time for minibus commuters, which included its own data collection. It included a baseline survey at the time of onboarding, daily interaction with an enumerator recorded via an app at bus stops, and an endline survey, in the summer of 2023. This data is described in more detail in Section 6.2.

Other Sources. We use the Lagos Travel Survey, a representative 2009 survey of 12,274 individuals commissioned by LAMATA. To measure congestion impacts, a dataset of live travel times was collected from Google Maps multiple times per day throughout the study period. We also conducted a small survey of commuters near major public transit terminals in 2024; it provides choice data for the nested logit model of Online Appendix S4.4.2. Finally, the 2024 Nigerian Time Use Survey gives the share of residents who make any trip on a given day, which we use for a robustness test allowing the public rollout to affect the decision to take a trip (Online Appendix S4.4.1). To produce a map of the private network in Figure 1B, we commissioned a network mapping in the fall of 2022, and conducted one round of observation surveys to measure each route's characteristics.[18]

[17]The reweighted sample includes 849 drivers, as 5 drivers (recruited near terminals) had not worked in the past seven days and could not provide the information needed to construct these weights.

[18]This aimed to identify and georeference every usual private transit route in the Lagos Metropolitan area. It was conducted by a specialized firm, WhereIsMyTransport, by sending enumerators to identify and ride minibus routes while using a GPS tracker. Because we were not able to schedule the mapping before the other observations, it is used only for this map and some descriptive statistics used in the welfare quantification.

3.2. Features of Private Transit in Lagos

Here we use our data to document features of private transit in Lagos. We draw from baseline driver statistics presented in Appendix Table S4, and additional statistics generated from our baseline observation data.

Drivers are Relatively Highly Paid but Have a Challenging Job. On the last day worked, the average net income was ₦5,140, net of fuel, payments to conductors, fees paid at terminals, and any side payments. Minibus drivers earn more on average than commuters, and are more likely to be male and have less than a secondary education, as shown in Appendix Table S5. However, this pay likely compensates for the challenge of the job. Driving minibuses entails conflict: 48% of drivers report having disputes with passengers weekly, and 37% report having disputes with other drivers weekly. Minibuses frequently break down, with 92% of drivers reporting a repair in the last month.

Incentives to Adjust. Drivers complete 9.2 trips per day on average. Most drivers are residual claimants on fare payments they collect. Fuel is a substantial expense: ₦3,766 on average in the last day worked. 31% of drivers hired a conductor to collect fares and help passengers board, paid an average of ₦2,531. Also in the last month, 83% of drivers report paying a fine, commonly for being caught driving without proof of paying the terminal fee, or stopping in an inappropriate location.

Margins of Adjustment. Drivers can change routes: 85% of drivers report that they choose the routes they operate. However, switching terminals is costly: drivers must pay the association a one-time ₦15,494 on average to register to pick up passengers at a terminal. After paying that, drivers pay an average of ₦691 per day in additional fees to the association. 87% of drivers drive a single route, resulting in an average number of routes per driver close to one. As we discuss in Section 5.2.2, route switching is fairly frequent: 59% of drivers change routes over the year we monitor, with 86% of these changes coming within the same terminal.

It may be challenging for drivers to switch industries. The average driver surveyed has driven minibuses in Lagos for 12.7 years. Some have sunk investments: 51% of drivers own their buses; the remainder either rent or pay for their vehicle on installments (rent to own).

Excess Supply. In 91% of our terminal route observations, there was a minibus waiting for passengers to load, and on average, 4.75 buses were waiting in queue. And in the baseline driver survey, drivers report that 42% of working time is spent waiting in queues. That buses are queuing for passengers is suggestive of excess entry. It also suggests that minibus departure frequencies will mostly be determined by passenger arrival rates, and will tend to not be affected by marginal changes in supply of minibuses, since almost all buses (96%) wait until they are full before departing. Some drivers skip terminals and the queue to start driving along a route with an empty bus (locally this is called driving 'sole'); our survey suggests this constitutes a minority (21%) of trips.

Limited Expectations About Rollout. As of 2020, only 57% of drivers had heard about the public transit rollout. Among those, 82% of those who predicted it would not open on their route were correct. However, only 22% of those who predicted it would open on their route were correct, and none correctly predicted the month of entry. 36% expected public transit would reduce their passengers. See Appendix Table S6.

3.3. Comparison of Public and Private Service

We next compare public and private transit service. Table 1 compares characteristics of private transit to public transit, before and after entry of public transit. We restrict the sample to routes on which public transit entered. Statistics are reported for journeys beginning at a terminal.

Private Transit is More Frequent. Private minibuses had frequent departures (shown in the first pair of rows): before public entry, 3.8-4.5 minibuses passed per half hour on average, and this declined slightly to 3.7-3.8 after entry. In contrast, we estimate that after entry only 0.2-0.6 public buses passed per half hour.[19] Overall, the public buses constituted 6-14% of all departures. Because the two services serve different stops, commuters must decide which type of bus to wait for–it would be difficult to monitor both and hop on the first departing bus.

Minibuses Fill Before Departing Terminals. 96% of minibuses leave only when full (most buses have 14 seats; shown in the second pair of rows). In contrast, the public buses were not always full to capacity, especially during the midday offpeak times. Despite the lower capacity of each minibus, minibuses serve high passenger volumes because departures are frequent (third pair of rows in Table 1).

Public and Private Fares are Similar. The final pair of rows of Table 1 compares fares for traversing from the origin terminal to the route's endpoint. Fares are similar on both systems during peak periods, while public fares are slightly lower during off peak (₦212 vs ₦224). Fares across the entire system rose over time in line with inflation in the country; we report the average fare across all measured private transit routes in the final row which shows this overall trend. On private buses, the fare is lower if you get off at an earlier stop, but most public routes have a fare that does not depend on distance. As a result, minibus transit tends to be relatively more

[19] Commuters seldom queue to board the minibus system; however, during the morning peak there may be queues to board public buses.

Table 1. Private vs. Public Service on the Same Route

	Pre-Rollout		Post-Rollout	
	Private transit	Public transit	Private transit	Public transit
Departures (every 30 mins)				
Peak	4.47	-	3.78	0.62
	(2.46)	-	(2.70)	(0.33)
Off peak	3.79	-	3.66	0.23
	(1.93)	-	(2.35)	(0.19)
Passengers at Departure				
Peak	13.63	-	12.94	25.66
	(4.44)	-	(5.61)	(6.02)
Off peak	13.88	-	13.60	10.40
	(4.74)	-	(6.03)	(5.33)
Passenger Flows (per 30 mins)				
Peak	55.55	-	40.21	15.88
	(25.90)	-	(16.93)	(8.89)
Off peak	47.77	-	41.09	3.28
	(16.94)	-	(13.06)	(2.38)
Fare (₦)				
Peak	202.13	-	223.50	223.45
	(74.92)	-	(82.37)	(115.86)
Off peak	199.47	-	223.96	211.95
	(60.07)	-	(72.86)	(116.75)
Average fare across all routes (₦)	284.70	-	300.79	

Notes: Table reports means across routes for a given time of day for routes that are eventually served by public transit. We separately report means before and after public transit commences on the route. Standard deviations reported in parentheses. Computed for the 9 (of 13) treated private transit routes in our sample that can be matched to public transit e-ticketing data. Peak hours are defined as 7-9 a.m. and 4-6 p.m., and off-peak 1-2 p.m. Passengers counted upon departure from the origin terminal; additional passengers may board at later stops. For private transit we compute the route average over all our terminal observations before public transit began on that route, and after. For public transit, we compute the route average from e-ticketing data, using only the time periods and months that coincide with our private transit route observations. We then compute the average across routes, weighting routes based on the average volume of private transit passengers prior to rollout. For public transit, a bus departure is defined using the last swipe at a route endpoint, which is not followed by another swipe at the same endpoint within 30 minutes. Passenger counts are aggregated at each departure, calculated as the total swipes for the same bus, route, and endpoint up to and including the last swipe. Passenger flows per 30 min are defined as the passenger counts aggregated at each departure, multiplied by the number of departures within a 30-minute window.

attractive for shorter trips. Transfers are fairly common, but there is no discount for transferring: on both systems, one must pay the independent fare on each leg.

Total Transit Ridership is Relatively Unchanged. Total transit ridership is 1% higher at peak and 7.1% lower off-peak after entry (Table 1), and a route-level difference-in-differences finds a modest 15% rise (Appendix Table S7). Our model allows the new option to raise combined transit ridership, discussed further in Section 4.5.

4. A Model of Hybrid Transit Markets

To guide our analysis, we develop a model where public and minibus transit co-exist. A government bus enters a route served by minibuses; commuters choose between the minibus, the new public service, and an outside option. The price and frequency of minibus service emerge from the route choices and queueing behavior of drivers and the decisions of a monopoly association that sets prices. We take the prices and schedule of public buses as given.

The model characterizes how public entry affects commuters and incumbent minibus drivers, and how those effects propagate across the network. In theory, total welfare could increase or decrease. Commuters gain a new option and lower minibus fares but may wait longer; drivers on treated routes lose revenue and some reallocate, spilling the effects onto other routes. A set of sufficient statistics aggregates these direct and indirect effects and organizes the empirical analysis that follows. Online Appendix S4.1 develops the full dynamic queueing model; here we summarize its key steady state relationships.

4.1. Commuters

Time is discrete and the horizon is infinite. There are I locations, $i \in \{1, \ldots, I\}$. A pair of locations ij is called a route.

Each period, μ_{ij} commuters arrive at i intending to travel to j and choose from a set of available modes $\mathcal{M}_{ij}$. Prior to government entry, $\mathcal{M}_{ij} = \{M, 0\}$ includes minibus (M) and an outside option (0), which captures travel by other modes such as walking or automobile. More broadly, the outside option may also include not taking the trip; we assess robustness to this interpretation. When the government enters, treated routes gain public buses (P) so that $\mathcal{M}'_{ij} = \{M, P, 0\}$.[20]

Commuter n who travels from i to j using mode m obtains utility

$$u_{ijmn} = u_{ijm} + \epsilon_{ijmn} \tag{1}$$

where u_{ijm} is the common utility provided by mode m on ij, and ϵ_{ijmn} is an idiosyncratic preference shock, type 1 extreme value with scale parameter $1/\theta$. The share choosing mode m is

$$s_{ijm} = \frac{\exp\left(\theta u_{ijm}\right)}{\sum_{m' \in \mathcal{M}_{ij}} \exp\left(\theta u_{ijm'}\right)}. \tag{2}$$

The arrival rate of passengers on ij for mode m is therefore $s_{ijm}\mu_{ij}$.

When they become available, public buses provide common utility u_{ijP}, which is assumed to be constant: the welfare gain from this new option will be revealed by its ridership alone.

[20] The primary other options are walking and cars with 9% and 12% of trips in the 2009 travel survey respectively. We collapse other modes into the outside option since the payoffs to these are unlikely to be affected by the reform: Section 5.2.1 shows no effect on road speeds.

To account for the change in consumer welfare arising from the response of the minibus sector to public entry, we model the minibus payoff u_{ijM} explicitly. The queueing model in Online Appendix S4.1.1 delivers common utility from choosing the minibus

$$u_{ijM} = \alpha_M - \gamma p_{ijM} - \eta t_{ijM}, \tag{3}$$

which includes a fixed amenity term α_M (such as comfort or safety), plus the contributions of price p_{ijM} and total travel time t_{ijM}, which may change in response to public entry. Travel time is the wait at the stop t^W_{ijM} plus in-vehicle travel time t^T_{ij}. We take in-vehicle time as exogenous, since public buses use the same roads as minibuses and entry leaves road speeds unaffected (Section 5.2.1). Passenger wait times are endogenous, determined jointly by how many drivers serve the route, when they depart, and how fast passengers arrive, as we characterize below.

The parameters γ and η represent price and time sensitivities conditional on mode choice; $\theta\gamma$ and $\theta\eta$ represent the sensitivities when commuters can substitute across modes. We normalize the utility provided by the outside option to $u_{ij0} = 0$.

4.2. Supply

Minibus Drivers. Each period, B_i drivers enter terminal i and choose a route ij by joining its queue. Passengers board the bus at the front of the queue; buses depart once full at $\bar{n}$ passengers. A departing bus travels to the destination and either returns to queue ij or exits.[21]

Driver φ working route ij earns realized daily profit

$$V^Q_{ij\varphi} = V^Q_{ij} \cdot \nu_{ij\varphi},$$

which depends on the common term

$$V^Q_{ij} = \underbrace{\frac{T}{t^Q_{ij} + t^T_{ij}}}_{N^{\text{Trips}}_{ij}} \cdot \underbrace{[p_{ijM}\bar{n} - c_{ij} - \tau_{ij}]}_{\pi_{ij}}, \tag{4}$$

and an idiosyncratic component $\nu_{ij\varphi}$, which reflects heterogeneity in workday length. Driver φ works a day of length $T_{ij\varphi} = T \cdot \nu_{ij\varphi}$ on route ij, where T is the average workday length, and so completes $\nu_{ij\varphi}N^{\text{Trips}}_{ij}$ trips (Online Appendix S4.1.2).

Route profitability is determined by the average number of daily trips N^{Trips}_{ij} and per-trip profit π_{ij}. The average number of trips is the average workday length T divided by total trip time: queueing at the terminal t^Q_{ij} plus travel time t^T_{ij}. Per-trip profit is fare revenue $p_{ijM}\bar{n}$ net of operating costs c_{ij} and the association per-trip fee τ_{ij}. Queue time t^Q_{ij}, which we characterize

[21] We abstract from the return leg of trips. Drivers choose a route only at entry, since 87% of drivers in our driver survey drive from an origin to a destination and back again.

below, is a key margin through which government entry impacts drivers: the new public option draws away minibus users, so passengers arrive more slowly and queues lengthen if buses still wait to fill before departing.

A driver entering terminal i picks the route that maximizes his profit, $\max_j\{\nu_{ij\varphi} \cdot V_{ij}^Q\}$. We assume $\nu_{ij\varphi}$ is distributed Fréchet with unit mean and shape parameter $\sigma > 1$, so the share of drivers choosing route ij increases in its profitability, with sensitivity set by the Fréchet shape σ:

$$\rho_{ij} = \frac{\left(V_{ij}^Q\right)^\sigma}{\sum_k \left(V_{ik}^Q\right)^\sigma} = \frac{\left(N_{ij}^{\text{Trips}} \cdot [p_{ijM}\bar{n} - c_{ij} - \tau_{ij}]\right)^\sigma}{\sum_k \left(N_{ik}^{\text{Trips}} \cdot [p_{ikM}\bar{n} - c_{ik} - \tau_{ik}]\right)^\sigma}. \tag{5}$$

Drivers face a fixed cost F_i to register at terminal i. Properties of the Fréchet distribution imply that average profits of each driver at terminal i are the same on their chosen routes, and given by

$$\Pi_i = \underbrace{\left[\sum_k \left(N_{ik}^{\text{Trips}} \cdot [p_{ikM}\bar{n} - c_{ik} - \tau_{ik}]\right)^\sigma\right]^{1/\sigma}}_{\Pi_i^V} - F_i \tag{6}$$

where Π_i^V are average variable profits.

Entry. At each terminal, the number of entrants B_i is pinned down by free entry in the initial steady state so that $\Pi_i^V = F_i$. Our empirical results suggest that drivers switch routes within terminals after public entry but do not substantially exit terminals or the industry. We therefore hold B_i fixed when quantifying the impacts of public entry. This captures a medium-run response.

Minibus Association. The minibus association at each terminal chooses the price p_{ijM} and driver fees τ_{ij}, F_i to maximize its revenue.

Public Buses. Public buses charge a price p_{ijP}, determined outside the model. In contrast to minibuses, public buses leave at exogenous times regardless of passenger arrivals, though in practice they do not follow a reliable schedule.

4.3. Minibus Equilibrium

In equilibrium the choices of commuters, drivers, and the association are mutually consistent, pinning down steady-state driver queues and commuter wait times. Because minibuses depart only once full, passenger wait times depend on how fast both buses and passengers arrive. Competition from the new public option lowers the minibus passenger arrival rate $s_{ijM}\mu_{ij}$; we trace how waits, queues, and prices adjust as minibus demand falls.[22]

Commuter Wait Times. A minibus commuter faces expected wait time

$$t^W_{ijM} = \beta_{ij} \underbrace{\frac{1}{s_{ijM}\mu_{ij}} \left(\frac{\bar{n}-1}{2}\right)}_{\text{time for bus to fill}} + (1-\beta_{ij}) \underbrace{\frac{1}{\lambda_{ij}}}_{\text{time for bus to arrive}}.$$

With probability β_{ij} a bus is waiting in the queue, and the commuter boards it. A waiting bus holds $(\bar{n}-1)/2$ passengers in expectation, so once she boards another $(\bar{n}-1)/2$ must arrive before it reaches its capacity $\bar{n}$ and departs. Since passengers arrive at rate $s_{ijM}\mu_{ij}$, this takes $\frac{1}{s_{ijM}\mu_{ij}}\left(\frac{\bar{n}-1}{2}\right)$ minutes. With probability $1-\beta_{ij}$ no bus is waiting, so she waits for the next bus which arrives at rate λ_{ij}. In steady state at least $\bar{n}$ passengers have queued by the time it arrives, so it fills and departs immediately. The probability a bus is waiting is $\beta_{ij} = \frac{\lambda_{ij}\bar{n}}{s_{ijM}\mu_{ij}}$: a bus is more likely to be waiting when buses arrive faster, or minibus demand is lower (since buses then fill slowly and remain in the queue longer).

When minibuses are in excess supply, a bus is almost always waiting in the queue and commuter wait times are governed by demand: a commuter waits only for enough other passengers to fill the bus. In this case, when demand for minibuses s_{ijM} declines, passengers arrive more slowly, so buses take longer to fill and commuters wait longer. In our baseline data, a minibus is waiting in the queue in 96% of treated route observations. So public entry is likely to raise wait times for commuters who continue to use minibuses.

Minibus Queue Times and Route Choices. The time drivers spend in bus queues in steady state is given by

$$t^Q_{ij} = \underbrace{\frac{\bar{n}}{s_{ijM}\mu_{ij}}}_{\text{time for bus to fill}} \cdot \underbrace{(N^Q_{ij}+1)}_{\text{queue length including driver } \varphi}. \tag{7}$$

A driver joining a queue of N^Q_{ij} buses waits for every bus ahead, then his own, to depart, with each taking $\frac{\bar{n}}{s_{ijM}\mu_{ij}}$ minutes to fill.

When minibus demand $s_{ijM}\mu_{ij}$ falls on a treated route ij, the immediate effect is to increase fill times. However, steady state queue lengths may also change if some drivers cease to serve the route after it becomes less profitable.[23] Which force dominates is an empirical question: if the fill-time effect wins, queue times rise on treated routes and drivers there make fewer trips.

[22] Precise comparative statics are not possible given the many interactions in the model; Online Appendix S4.2.8 gives the full system of equations that must be solved simultaneously. The objective of the model is to deliver a set of sufficient statistics that can be measured in the data to use for welfare analysis of the observed intervention.

[23] In steady state, $N^Q_{ij} = \left[\frac{\mu_{ij}s_{ijM}}{\lambda_{ij}\bar{n}} - 1\right]^{-1}$; see Online Appendix S4.1.4.

When the expected profits of serving route ij decline, drivers may substitute to a connected route ik at the same terminal, following equation (5). If demand on connected routes is unchanged (as we will show to be the case), this increases queue lengths and times, lowering daily trips per driver there too.

Minibus Fares. The association sets prices to maximize revenues, which it captures through registration fees. This leads to an optimal markup

$$p_{ijM} - c_{ij}/\bar{n} = \frac{1}{\frac{t_{ij}^Q}{t_{ij}^Q + t_{ij}^T} \frac{\partial \ln t_{ij}^Q}{\partial p_{ij}}}. \tag{8}$$

The association balances the two effects of a higher fare on revenue: raising the fare earns more on every passenger, but also lowers demand, which lengthens queue times and cuts daily trips per driver. Equating these two forces leaves the standard monopoly markup: the per-passenger margin $p_{ijM} - c_{ij}/\bar{n}$ equals the inverse semi-elasticity of daily trips to the fare.

When the public option enters, commuters on treated routes get an additional travel option, which makes their demand for minibuses more elastic ($\theta\gamma(1 - s_{ijM})$ rises). Since equation (7) implies that the markup depends on $\frac{\partial \ln t_{ij}^Q}{\partial p_{ij}} = \theta\gamma(1 - s_{ijM}) + \frac{N_{ij}^Q}{1+N_{ij}^Q} \frac{\partial \ln N_{ij}^Q}{\partial p_{ij}}$, this tends to reduce minibus fares on treated routes, unless a large number of drivers leave the route (reducing queue times and lengths). If some of those drivers switch to connected routes (through the route choices in equation (5)), then connected routes will have more drivers but not more passengers, lengthening bus queues. A longer queue is idle capacity, and the association responds by cutting fares on connected routes as well, in an effort to attract more passengers (see Online Appendix S4.2.1).

4.4. Changes in Surplus from New Government Transit

Consumer Surplus. Public entry changes consumer surplus on a route through two channels:

$$\begin{aligned} \Delta CS_{ij} = {} & \frac{1}{\gamma\theta} \ln \left(\frac{1}{1 - s'_{ijP}} \right) \\ & + \frac{1}{\gamma\theta} \ln \left(s_{ij0} + s_{ijM} \exp \left(-\theta\eta\Delta t_{ijM} - \theta\gamma\Delta p_{ijM} \right) \right). \end{aligned} \tag{9}$$

The first line is the direct gain from the new public option. We do not need to know precisely *what* commuters value about it: the post-entry share s'_{ijP} is sufficient to reveal the utility it provides commuters, scaled by the parameter θ that governs how easy it is to substitute between alternatives.

The second line captures the indirect impacts that come from the private sector response. If minibuses change their fares or wait times, surplus shifts for travelers who continue using them. Driver reallocation can also lead to spillovers to untreated but connected routes. On these routes, the first line is zero because there is no public service, but minibus fares and wait times can still change as driver supply increases. The total change in consumer surplus $\Delta CS = \sum_{ij} \mu_{ij} \Delta CS_{ij}$ aggregates this across routes.[24]

Producer Surplus. Public entry impacts driver profits on both treated and connected routes. On treated routes, public transit directly competes with minibuses, lowering profits by cutting daily trips and fares. On connected routes, driver reallocation steals business from incumbents: as drivers switch from treated routes to others at the same terminal, bus queue lengths and queue times rise, reducing daily trips. Fares are also likely to fall, since longer queueing on connected routes lowers the optimal markup (see equation (8)).

Both effects are captured by the change in average variable profits, which is equalized across routes. Letting $\hat{x} \equiv x'/x$ denote the relative change in a variable between pre- and post-government entry equilibria, this is given by

$$\hat{\Pi}_i^V = \left[\sum_j \rho_{ij} \left(\hat{N}_{ij}^{\text{Trips}} \right)^{\sigma} \left[\hat{p}_{ijM} \pi_{ij}^{rev} + 1 - \pi_{ij}^{rev} \right]^{\sigma} \right]^{1/\sigma}, \tag{10}$$

where $\pi_{ij}^{rev} = p_{ijM}\bar{n}/(p_{ijM}\bar{n} - c_{ij} - \tau_{ij})$ is the ratio of gross to net revenues.

The change in aggregate producer surplus is $\Delta\Pi = \sum_i B_i \Pi_i^V (\hat{\Pi}_i^V - 1)$. Under long-run free entry, any gain would flow to the association through higher fees. However, as discussed below, our data show no significant changes in driver exit (either across terminals or out of the industry) or association fees. We interpret this as a medium-run response, and with fixed driver entry and unchanged fees $\Delta\Pi$ accrues to drivers. We therefore refer to this change as driver surplus, noting that under different assumptions, it could be shared with the association.

Measuring the Sufficient Statistics. The sufficient statistics in equations (9) and (10) can be computed with four categories of empirical objects.

First, we need descriptive statistics from the initial equilibrium. For commuters, these are traveler arrival rates (μ_{ij}), pre-entry minibus market share (s_{ijM}, which also gives the outside option share $s_{ij0} = 1 - s_{ijM}$), and the post-entry public bus market share (s'_{ijP}). For drivers, these are pre-entry terminal choices and route shares (B_i and ρ_{ij}) and profit margins (π_{ij}^{rev}). We measure the commuter objects by combining the 2009 Lagos Travel Survey, a representative survey of all trips, with the ridership data described in Sections 2.1 and 3.1; the driver objects

[24]See Appendix Section S4.2.7 for a derivation. The breakdown is similar to love-of-variety models in trade (Feenstra 1994). While government entry on one route could in practice affect minibus demand on connected routes, we find no evidence of this in Section 5.3, so we omit demand-side spillovers from the model for the purposes of welfare calculation.

come from our driver surveys and terminal observations. See Online Appendix S4.3.1 for details.

Second, we estimate the reduced-form effects of public entry on commuter wait times (Δt_{ijM}), fares ($\hat{p}_{ijM}$), and driver trips ($\hat{N}_{ij}^{\text{Trips}}$), across both treated and connected routes, from the staggered government entry in the next section.

Third, we measure commuters' price and time sensitivities (γ and η), conditional on mode choice. We estimate these from our wait time field experiment described in Section 6.

Fourth, we require the substitution parameters θ and σ, which govern how easily commuters substitute across modes and drivers substitute across routes. We identify θ from the response of commuter trips to abrupt price changes on the public system, and σ from the response of driver route choices to government entry.

4.5. Discussion and Limitations

Time Horizon of Analysis. Our sufficient statistics approach evaluates only short- and medium-run effects, to align with the empirical results observed in our sample period of just over a year. Several adjustments that we do not model or observe in our data might emerge over a longer horizon: for example, drivers might exit the industry in large numbers, or the association may change fees.

Extensive Margin Effects. In the model, total transit ridership can change after public entry as commuters substitute between transit and the outside option. We consider two definitions of the outside option. The main definition considers the outside option as taking the trip via any different mode (including walking), which implies that changes in transit ridership arise from substitution with other modes. This is the definition for which we have the most reliable data, since our main source of travel data conditions on making a trip. We also consider a definition that includes in the outside option the choice of not traveling, which allows changes in ridership also to arise from changes in the decision to take a trip at all; see Online Appendix S4.4.1. In either case, the model rules out larger adjustments such as households relocating or changing jobs. Such responses are unlikely in our setting, since public and private transit share the same roadway and offer the same in-vehicle travel times. Total transit trips are higher in the post period by 1% at peak (Section 3.3), and a difference-in-differences finds total transit trips rise by 15% when a route gets treated (Table S7), consistent with switching rather than large-scale trip generation.

Substitution Patterns. The multinomial logit model generates intuitive theoretical expressions but imposes independence of irrelevant alternatives (IIA). This restricts commuters to substitute between minibuses and public transit at the same rate as between either transit mode and the outside option. As a robustness test, we ease this restriction by also estimating our welfare results using a nested logit model that groups the two transit modes in one nest and the outside option in another, using our 2024 commuter survey (Section 3.1). Under that model the welfare gains of introducing public transit are slightly lower; see Online Appendix S4.4.2.

5. How Private Transit Responds to Public Rollout

This section uses the panel data we collected to measure how incumbents responded to the opening of the new public system. We are primarily interested in impacts on minibus departure frequencies (which determine wait times), prices, and driver queues. We also assess impacts on drivers.

5.1. Descriptive Evidence

Figure 3 compares overall public transit ridership with that on our three categories of private routes: *treated* routes that share both endpoints with a public bus route that opened within our observation period, *connected* routes that share one endpoint, and *control* routes that share neither endpoint.

In our sample of routes, three patterns are evident. First, public transit ridership (blue dot-dashed line) increases over this period as the system is rolled out over more routes. Second, private transit ridership on treated routes (orange dashed line) sees a commensurate decline. Third, there appears to be no meaningful trend on connected or control private transit routes. This suggests that most of the overall ridership impact of public transit appears to be local to treated routes, and justifies a specification that identifies effects on treated routes by comparing them to other private transit routes. We will more finely investigate impacts on connected routes at the end of the section, and find some evidence of supply-side spillover effects. Finally, note that despite the large investment in public transit, it remains a smaller share of trips than minibuses on the routes we sampled.

The three categories of private routes were similar at baseline, as shown in Appendix Table S8. Treated routes had similar route distances and driver queue lengths to connected and control routes, but were slightly higher volume routes (they had more frequent minibus departures and more passengers). Overall, we fail to reject tests of joint equality of these attributes between treated and either connected or control routes. However, *terminals* that had at least one treated route tend to be better connected and have higher passenger volumes than terminals that have no public transit connections, as shown in Appendix Table S9. We reject a test of joint equality between the characteristics of these two groups of terminals (F-stat 5.26).

Because the two groups of terminals are different, we will begin with specifications that absorb differential trends at the terminal level, using terminal-by-survey-round fixed effects. However, because this approach compares routes within terminals, it will be biased if untreated routes within a terminal are affected, such as through driver substitution or other spillovers (a violation of Stable Unit Treatment Value Assumption or SUTVA). In Section 5.3, we remove

Figure 3. Trends in Trips

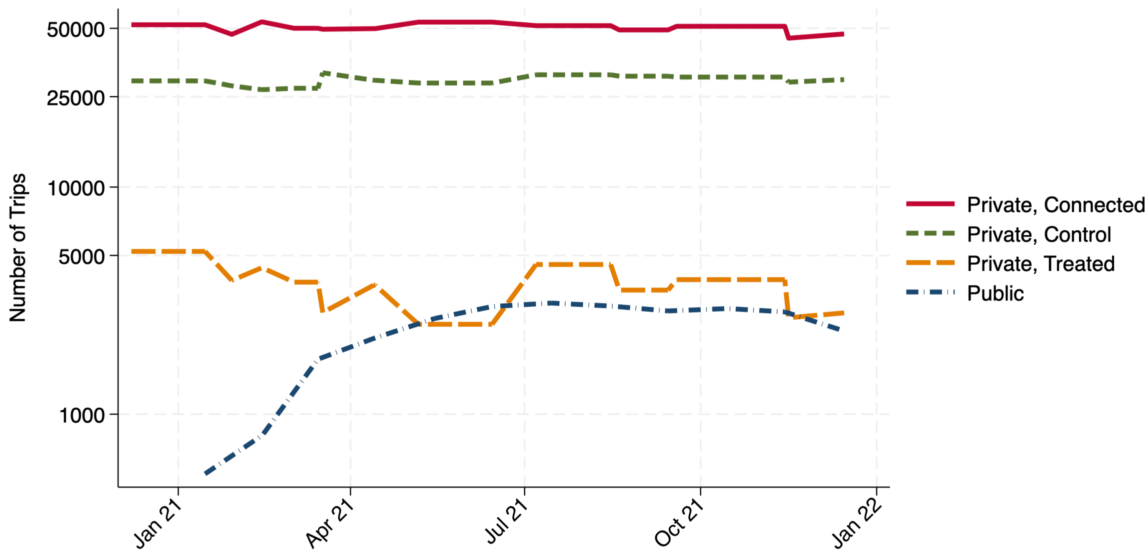

Note: Plots show the total monthly passenger counts in private and public transit. Public transit ridership comes from electronic ticketing data and covers 11 of the same 13 routes we observe in the private treated sample. Private transit ridership comes from our terminal observations. Not all routes were monitored in all survey rounds: we impute missing route observations with the average number of passengers in routes of the same group (private treated, private control, or private connected) in that round. For route-months with incomplete public transit ridership, we interpolate by averaging adjacent months when possible. For missing route-months at the end of the sample period, we carry forward the last observed ridership value for that route. Vertical axis shown in logarithmic scale.

these fixed effects and estimate specifications that account for spillovers of particular forms. We find small supply-side spillovers between routes at a terminal but show that the estimated effect on treated routes is similar whether we control for terminal trends or spillovers. Finally, Section 5.5 provides evidence that possible higher-order spillovers beyond those modeled do not appear to be at play in our setting.

5.2. Impact on Treated Routes

We begin by estimating how minibuses respond when public transit enters their routes, with terminal-by-survey-round fixed effects.

Base Regression Specification. Outcomes are at the level of route r, survey round t, and time of day τ (30 minute intervals within morning peak, afternoon peak, and afternoon off peak). Our main estimating equation is

$$Y_{rt\tau} = \beta\mathbb{I}\{\text{Open}_{rt}\} + \gamma_{m(r\tau)} + \delta_{i(r)t} + \boldsymbol{\kappa}_t{}' X_{rt\tau} + \epsilon_{rt\tau}, \tag{11}$$

which includes an indicator for whether route r is served by public buses as of survey round t as well as a variety of controls including fixed effects and observable controls $X_{rt\tau}$. The coefficient of interest is β which represents the impact of public transit entry.[25]

Commuting patterns in Lagos are highly asymmetric, with the same route varying significantly by time of day.[26] We define a 'transit market' $m(r\tau)$ as a route segmented into three periods: morning peak, afternoon peak, and afternoon off-peak. Transit market fixed effects $\gamma_{m(r\tau)}$ are included to control for time-invariant unobservables and to identify β from changes in outcomes driven by public entry. Control variables $X_{rt\tau}$, which include day of week, hour of departure fixed effects, and a set of fixed trip characteristics such as trip length, are interacted with survey round fixed effects to allow their effects to be time-varying. Specifications in this section include terminal-by-survey-round fixed effects $\delta_{i(r)t}$, with $i(r)$ representing the origin terminal of route r. These will absorb time-varying shocks at the terminal level.[27]

Our identification strategy compares how private transit shifts on a route when public transit enters, relative to other routes. In our focal specifications, this comparison is made relative to other routes sharing a terminal endpoint. This is because terminal-by-survey-round fixed effects absorb terminal wide shocks. As a result, this relies primarily on comparisons between changes in treated versus connected routes serving the same terminals. The disadvantage of this strategy is that estimates will be biased if there are spillovers between routes, which we will investigate later. Standard errors are clustered by route and terminal since treatment is at the route level but our sampling was at the terminal level. Since our sample of treated routes is small, we also report wild cluster bootstrap p-values (Cameron, Gelbach, and Miller 2008).

5.2.1. Departure Observations

We estimate the effect of public transit entry on outcomes observed in our departure observations at terminals, with results shown in Table 2.

Introducing public buses on a route decreases the departure rate of minibuses by 11-22%, as shown in Panel A. In the simplest specification in column (1), departures fall by 11% after public entry. Since on each route there are typically minibuses waiting for passengers to depart, these reductions reflect a decline in passengers. Because shared transport works by pooling passengers, less frequent departures mean higher waits for those who continue to use minibuses,

[25] When treatment effects are heterogeneous, staggered openings can bias two-way fixed effects estimates, as already-treated units serve as comparisons for later-treated units (de Chaisemartin and D'Haultfoeuille (2023), Sun and Abraham (2021)). Such comparisons represent a small fraction of our observations. In the within-terminal specification a treated route is compared with the other routes at its terminal; already-treated routes can serve as comparisons only when two treated routes at the same terminal open in different rounds, which occurs in only 15 of the 117 treated route-by-round observations. In our spillover specification, around half of route-by-round observations are the comparison group of 120 control routes, which are never treated and never connected, and already treated units serve as comparisons in 2.2% of observations. The close agreement between our two-way fixed effects and Sun and Abraham (2021) estimates, both in Section 5.4 and in the event studies of Figure S5, suggests these concerns are limited.

[26] Lagos' patterns are typical: morning travel flows from the outskirts to the center, and the reverse in the afternoon.

[27] In theory one might also wish to include fixed effects for the destination terminal of route r, but empirically that would absorb most of our variation. That is because we observe departures on sampled routes at each terminal: we capture many routes with the same origin, but few with the same destination.

as they will spend longer waiting in a bus as it fills to depart. This effect grows stronger as we add controls. Column (2) adds the terminal of observation-by-survey round fixed effects, allowing for aggregate shocks to locations, which sharpens up these effects. Column (3) adds controls for trip distance interacted with survey round fixed effects (to allow for differential trends by trip distance). Column (4) adds an indicator for whether the route was ever included in one of the government's opening plans for this phase of the system, again interacted with a time indicator. Column (5) adds controls for whether a route ends in one of Lagos' central business districts. Column (6) adds flexible controls for origin and destination characteristics through a 3rd order polynomial in their longitude and latitude. Finally, column (7) adds origin and destination region fixed effects interacted with time indicators. We define region by the local government areas (LGAs), which partition Lagos into 20 administrative units. Overall, the estimate of largest magnitude (22% decline in departures) comes from our most restrictive specification in column (7).

When public buses are introduced on a route, there is some evidence that private fares may decline slightly relative to other routes serving the same terminals, shown in Panel B. While point estimates suggest a reduction in fares between 2-6% across these specifications, these estimates are not statistically significant. Our wild bootstrap confidence intervals can rule out declines larger than 17%.

Combined, the slower arrival of passengers and possible reductions in fares make treated routes less desirable for drivers to serve. Correspondingly, we see declines in the number of minibuses waiting in queues for passengers of 23-29% depending on specification, as shown in Panel C. This suggests drivers on treated routes must be leaving those routes, which we will turn to in the next section.

We treat column (7) as the preferred specification going forward since it includes the richest set of controls.

Impacts on Congestion. We find no impact of public transit rollout on measures of traffic congestion derived from Google Maps (log road speed and travel time), reported in Appendix Table S11. This is not surprising given that the public buses we study share the roadways, and did not pack passengers much more densely than minibuses on average (Table 1).

Taking stock, we find that on treated routes, departures decrease, and there is suggestive evidence that fares decrease. Driver queues also shorten. This raises the question: where do affected drivers go? We assess this in the next two subsections.

5.2.2. Drivers

We next turn to our panel of driver surveys to understand how individual drivers are affected when public transit rolls out on their primary route.

Table 2. Effect of Public Transit on Private Transit

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: log(Departures)							
Open	-0.114	-0.182	-0.180	-0.189	-0.186	-0.219	-0.221
	(0.008)	(0.000)	(0.002)	(0.002)	(0.003)	(0.008)	(0.005)
	[0.144]	[0.040]	[0.076]	[0.062]	[0.058]	[0.078]	[0.060]
N	21,287	21,287	21,287	21,287	21,287	21,287	21,287
R^2	0.58	0.64	0.64	0.64	0.64	0.65	0.66
Mean Outcome (Levels)	3.45	3.45	3.45	3.45	3.45	3.45	3.45
Panel B: log(Fare)							
Open	-0.031	-0.051	-0.037	-0.055	-0.059	-0.060	-0.024
	(0.512)	(0.324)	(0.473)	(0.273)	(0.221)	(0.144)	(0.501)
	[0.563]	[0.368]	[0.531]	[0.351]	[0.299]	[0.229]	[0.585]
N	23,070	23,070	23,070	23,070	23,070	23,070	23,070
R^2	0.92	0.94	0.94	0.94	0.94	0.94	0.95
Mean Outcome (Levels)	209.51	209.51	209.51	209.51	209.51	209.51	209.51
Panel C: log(Queue)							
Open	-0.228	-0.282	-0.266	-0.270	-0.254	-0.275	-0.289
	(0.000)	(0.005)	(0.001)	(0.001)	(0.005)	(0.001)	(0.000)
	[0.047]	[0.132]	[0.074]	[0.070]	[0.084]	[0.049]	[0.033]
N	22,293	22,293	22,293	22,293	22,293	22,293	22,293
R^2	0.47	0.55	0.56	0.56	0.56	0.57	0.59
Mean Outcome (Levels)	4.75	4.75	4.75	4.75	4.75	4.75	4.75
Route × Period FE	X	X	X	X	X	X	X
Day of Week × Survey Round FE	X	X	X	X	X	X	X
Hour of Dep × Survey Round FE	X	X	X	X	X	X	X
Terminal × Survey Round FE		X	X	X	X	X	X
Trip Dist Controls × Survey Round FE			X	X	X	X	X
Dep. Plan × Survey Round FE				X	X	X	X
CBD Controls					X		
O & D Lat-Lon Poly × Survey Round FE						X	
O & D LGA × Survey Round FE							X

Notes: P-values clustered by route and terminal are reported in parentheses. Wild p-values computed using Cameron, Gelbach, and Miller (2008) are reported in square brackets. The rows at the bottom of the tables reflect control variables added to each specification. Route X Period FE are the 3 "transit market" fixed effects for each route. The remaining rows interact survey round fixed effects with the following variables. Rows 2 and 3 add fixed effects for each day of the week and hour of departure. Row 4 adds a fixed effect for each origin terminal. Row 5 adds dummies for each quartile of trip distance. Row 6 adds a dummy for whether the route ever appears in a LAMATA deployment plan, reflecting whether they planned to open public transit on that route at some point. Row 7 adds dummies for whether the route ends on Lagos Island or Victoria Island, containing the city's two main central business districts. Row 8 adds a third-order polynomial in origin and destination latitude and longitude. Row 9 adds fixed effects for the LGA of each origin and destination.

Drivers are a challenging group to track down for follow up surveys: they are highly mobile, and work long and erratic hours around the week, with unpredictable breaks. We followed a protocol to call multiple times and arrange to call back at convenient times. We recruited 854 drivers at baseline; 82% could be reached for at least one follow up but only 34% could be reached for all 4 follow ups.[28] Older drivers are more likely to participate in follow up rounds; we will control for age, though our results are not sensitive to this choice. Attrition is not significantly different between drivers whose routes at baseline were treated or connected. We present these results in Appendix Tables S14 and S15.

We return to the base regression specification, equation (11), and run it on outcomes from the driver survey. We compare drivers whose main route, as reported in the baseline survey, has the public service operating at the survey date ($\text{Open}_{r(d)t} = 1$ for driver d) with those whose main route does not. Results are shown in Table 3. Drivers whose baseline routes become served by public transit complete fewer trips per day (column (1)) and work more days (column (5)) compared to other drivers operating from the same terminal at baseline. They do not report significant differences in fares (column (2)) or the fee paid to the association for each trip (column (3)).[29] This reduction in the quantity of trips made aligns with the terminal observation results. In response, drivers are more likely to change routes. On average, 59% of drivers change routes over the year we monitor, and treatment results in an increase of 7.2 percentage points (column (6)). The vast majority of these changes are drivers selecting different routes from the same terminal (51% of all drivers, or 86% of all switchers): column (7) reports a sharper 11.6 percentage point increase in the likelihood a driver switches routes within the same terminal after the public service enters on their main route. In this specification, we do not find a difference in profits between drivers within treated terminals (column (4)).

These findings are consistent with the main terminal-level results but they suggest that changes in drivers' route choices may have affected other routes serving the same terminal. We explore this possibility next.

Table 3. Effect of Public Transit on Minibus Drivers

	N Trips	log(AvgFare)	log(Trip Fee)	log(Profit)	Days Worked	Change Route Any	Change Route Same Terminal
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Open	-1.550	-0.037	-0.039	0.016	1.533	0.072	0.116
	(0.005)	(0.298)	(0.501)	(0.761)	(0.000)	(0.136)	(0.006)
	[0.047]	[0.345]	[0.568]	[0.791]	[0.000]	[0.205]	[0.029]
N	2,177	2,015	2,030	2,085	2,648	1,411	1,411
R^2	0.76	0.87	0.78	0.59	0.58	0.84	0.80
Mean Outcome (Levels)	9.58	197.26	580.88	5004.84	4.41	0.59	0.51

Notes: P-values clustered by driver and terminal of recruitment reported in parentheses. Wild p-values computed using Cameron, Gelbach, and Miller (2008) are reported in square brackets. All columns include driver fixed effects, and survey round fixed effects interacted with terminal of recruitment, dummies for each quartile of driver age, the number of trips the driver reports in the previous workday at baseline, and a dummy for whether they own their vehicle or not at baseline. Number of trips are winsorized at 99th percentile due to large outliers. Regressions are weighted by the sampling weights discussed in Section 3.1.

[28] Appendix Table S12 reports the number of completed surveys in each round, and Appendix Table S13 reports the distribution of number of survey rounds completed.

[29] We only asked the per-trip payments to the association in all survey rounds; the terminal registration fee was asked only in the baseline survey since it is paid infrequently.

5.3. Impact on Connected Routes

We next test for the presence of spillover effects on private routes not directly exposed to public sector competition. We do so by dropping terminal-by-survey-round fixed effects and running regressions that compare treated, connected, and control routes. Connected routes are more likely to experience driver substitution, as the primary fee drivers pay the association is to serve a specific terminal, making it cheaper to switch to a route with the same origin or destination.[30] Conceptually, if a new public route opens on ij, the effect of treatment is identified by comparing changes on the treated route (ij) with those on untreated routes serving terminals with no public service (control routes, for example, lm). The spillover effect is identified by comparing how untreated routes change when their terminal receives new public service (connected routes, for example, ik) with those on untreated routes at terminals without public service (e.g., lm).

Our spillover specifications omit terminal-by-survey round fixed effects, and take the form

$$Y_{rt\tau} = \beta \mathbb{I}\{\text{Open}_{rt}\} + \alpha \text{ Connections to public transit}_{rt} + \gamma_{m(r\tau)} + \boldsymbol{\kappa}_t{}' X_{rt} + \epsilon_{rt\tau}.$$

where Connections to public transit$_{rt}$ describes route r's connections to public transit. In our main specifications this is the number of running public routes that share an endpoint with route r; we consider alternative specifications of public transit connections in robustness checks. We define these connection measures based on the max of the variable at the origin and destination, and set them to zero if route r itself ever receives public transit. This ensures the connection coefficients are estimated using variation in exposure on untreated routes rather than as treated routes are connected to more public bus routes. The impact of treating route r on treated route r is captured by β, and the impact on connected route r' is captured by α.

5.3.1. Departure Observations

For departure observation outcomes, results are shown in Table 4. Odd columns replicate the main specification (column (7) of Table 2), while even columns report the spillover specification which omits terminal-by-survey round fixed effects and add the connection measure.

The first two columns test for demand-side spillovers: when buses are waiting in queues (observed in 84% of spillover route observation windows at baseline), the departure rate of buses depends on the passenger arrival rate. The results show a precise zero effect of a route without new public service becoming connected to more public routes at either endpoint on its bus departure rate. This indicates that the introduction of new public service did not affect minibus passenger demand on untreated routes.

The last two columns test for supply side spillovers by examining the effects of new public

[30] Drivers could experience lower switching costs within terminals for other reasons, such as familiarity, and reports suggest that drivers are indeed "often tied to a single origin-destination (or sometimes a set of origin-destination), which are connected to a corresponding park" (CPCS 2024).

Table 4. Effect of Public Transit on Connected Private Transit Routes

	log(Departures)		log(Fare)		log(Queues)	
	(1)	(2)	(3)	(4)	(5)	(6)
Open	-0.221	-0.163	-0.024	-0.108	-0.289	-0.163
	(0.005)	(0.074)	(0.501)	(0.042)	(0.000)	(0.084)
	[0.060]	[0.134]	[0.585]	[0.117]	[0.033]	[0.135]
Number of open routes at terminal		0.002		-0.016		0.029
		(0.874)		(0.006)		(0.094)
		[0.887]		[0.023]		[0.141]
N	21,287	21,287	23,070	23,070	22,293	22,293
R^2	0.66	0.63	0.95	0.94	0.59	0.55
Terminal × Survey Round FE	X		X		X	
p-val: Same estimate on Open	0.59		0.05		0.30	

Notes: P-values clustered by route and terminal are reported in parentheses. Wild p-values computed using Cameron, Gelbach, and Miller (2008) are reported in square brackets. All columns include the same controls as the main specification (column (7) of Table 2), except for terminal-by-survey round fixed effects which are omitted in even columns. Number of open routes at a terminal is defined as the maximum of the number of public routes open at a minibus route's origin and destination, and is non-zero only for routes that never receive public service themselves. The average number of connected routes is 5.1. Final row shows the estimated p-value on a hypothesis test of equality between the coefficient on Open in the corresponding odd and even columns.

service on the length of driver queues on connected routes. As a route receives public transit, driver queues on that route shorten but queues to drive connected routes suggestively increase as shown in column (6). This is consistent with drivers substituting to serve other routes at the same terminal, affecting the supply of buses there and the number of daily trips drivers can make.

Connected routes experience an expansion of drivers with unchanged demand, lengthening bus queues. The model predicts that the association will respond to that idle capacity by cutting fares (Online Appendix S4.2.1). Column (4) supports this, with minibus fares declining on routes with new public transit connections.

5.3.2. Drivers

Although attrition in our driver survey is similar among those whose baseline routes become treated or connected to public transit, Appendix Table S15 finds evidence that the attrition of both could be at least 7-13% lower than those who served control routes.

Table 5 presents estimates of driver outcomes accounting for spillover effects. Despite the caveat that attrition may be differential, results are consistent with the broad story. We find suggestive evidence that drivers on connected routes experience similar effects as treated drivers, but to a smaller extent, consistent with second-order effects from the movement of drivers from treated to connected routes. While the specifications with terminal-by-survey round fixed effects found no change in profit between treated and connected routes, we find indications that profits decline relative to control in both types of exposed routes (column (4): by an insignificant 7% for treated and by a significant 2% per connected route for connected drivers). Other results for treated drivers remain consistent with the previous specification. Drivers on connected routes are more likely to switch routes within the same terminal (column (7)), receive lower average fares across their daily trips (column (2); estimate loses precision under wild-bootstrap inference), and work more days (column (5)). While there is no significant effect on the number of daily trips, multiplying the point estimate by the average of 5.1 open public routes at terminals with public service implies an effect for connected drivers approximately 10% as large as for treated drivers.

Driver Exit. Treated drivers' profits fall by approximately 7%, yet several pieces of evidence suggest drivers did not exit the industry over our sample period, either in aggregate or on exposed (treated or connected) routes. Exposed drivers are no more likely than control drivers to report not working in the past seven days (Appendix Figure S6). Their survey attrition is lower (Appendix Table S15). If drivers who exit the industry cease responding to our survey, that attrition result points in the opposite direction, though it could also be the case that drivers who exit become easier to reach. Minibus rental fees do not change differentially on their routes (Appendix Table S16). Instead, drivers adjust within the industry, switching routes at their terminal and working more days (Tables 3 and 5). At the industry level, exit would reduce demand for rented vehicles, yet rental fees show no downward trend over our sample (Appendix Figure S7). Exit is also costly: the average driver has 12.7 years of experience and half own their vehicles outright.

Table 5. Effect of Public Transit on Connected Minibus Drivers

	N Trips	log(AvgFare)	log(Trip Fee)	log(Profit)	Days Worked	Change Route Any	Change Route Same Terminal
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Open	-1.319	-0.064	-0.025	-0.072	1.854	0.091	0.169
	(0.008)	(0.082)	(0.642)	(0.129)	(0.000)	(0.014)	(0.001)
	[0.027]	[0.121]	[0.699]	[0.186]	[0.000]	[0.065]	[0.016]
Number of routes open at terminal	-0.028	-0.006	-0.002	-0.018	0.254	0.006	0.020
	(0.734)	(0.088)	(0.824)	(0.006)	(0.000)	(0.427)	(0.014)
	[0.798]	[0.168]	[0.832]	[0.028]	[0.000]	[0.479]	[0.046]
N	2,173	2,011	2,026	2,081	2,644	1,405	1,405
R^2	0.71	0.84	0.72	0.52	0.55	0.80	0.76
Mean Outcome (Levels)	9.58	197.26	580.88	5004.84	4.41	0.59	0.51

Notes: P-values clustered by driver and terminal of recruitment reported in parentheses. Wild p-values computed using Cameron, Gelbach, and Miller (2008) are reported in square brackets. All columns include driver fixed effects, and survey round fixed effects interacted with: dummies for each quartile of driver age, the number of trips the driver reports in the previous workday at baseline, and a dummy for whether they own their vehicle or not at baseline. Open and Number of open routes are defined analogously to the terminal observation specifications, but based on a driver's main route from the baseline survey. Regressions are weighted by the sampling weights discussed in Section 3.1.

5.4. Robustness

Placebo Tests. We assess whether these effects could be due to confounds by measuring impacts on 'placebo' routes that were planned but never opened, in Appendix S5.1. Following a fatal police shooting at Lekki Tollgate, protesters burned a public television station and the Oyingbo public bus terminal in October 2020. This terminal had been slated to receive 8 new routes at the end of October and late November, which were subsequently canceled. If our estimates were picking up differential trends on routes that the government planned to enter, then we might find that these canceled routes generate phantom effects. Our empirical strategy finds no phantom effects on these routes, nor on routes whose opening was scrapped during changes in opening plans across 22 variants of the rollout plan we digitized over our sample period. These results support the notion that our main treatment effects are due to entry itself.

Dynamics. In Appendix S5.2 we extend our base specifications to allow treatment effects to vary with time; there is some evidence that driver queues are affected after a delay.

Alternate Specifications. Our estimates are also robust to a number of alternate specifications.

Coefficients of the baseline specification are similar whether estimated with a two-way fixed effects estimator or a Sun and Abraham (2021) estimator that is more robust to treatment effect heterogeneity, shown in Appendix Tables S19–S22. Results are qualitatively similar when outcomes are reported in levels rather than logs, as shown in Appendix Table S10.

For the spillover specifications, Appendix Table S23 presents specifications with alternative measures of connections to public transit. Column (2) uses two dummies binning the number of open public transit connections (the baseline measure) into two groups: above and below 5 routes (the average number of open routes is 5.1), and column (3) uses a dummy for whether the route has any public connections. The findings of supply side and fare spillovers are robust to these alternative functional forms, with effects becoming larger on routes with more new public transit connections. For departures, we see some mixed effects but not convincing evidence of demand spillovers. Although there is a negative departure effect when a route receives any connection in column (3), column (2) suggests that the effect is stronger at terminals that are less affected, suggesting it may be picking up noise. Appendix Table S24 finds that driver results are generally similar when unweighted, and that spillover effects are larger when the continuous measure of transit connections is replaced with a dummy for whether a route has any connections.

5.5. Discussion

Reconciling the Baseline and Spillover Estimates. The previous analysis estimates two families of specifications. Since terminals with treated routes differ from those without, our first specification includes terminal-by-survey round fixed effects to flexibly control for unobservable terminal-level trends. The drawback is that treatment effects are identified through within-terminal comparisons, which may be subject to spillovers. Our second specification compares outcomes on routes at treated terminals (both treated and connected routes) with routes that do not interact with the new public system at either endpoint. This reduces concerns about spillovers but requires selecting a functional form, and sacrifices control over unobserved terminal-level trends. Thus a key question is whether our conclusions are sensitive to the approach.

Appendix Table S23 presents three alternative functional forms for spillovers in columns (1–3). For each, we report p-values for whether the coefficient on Open_{rt} differs from that in the specification with terminal-by-survey round fixed effects in column (4). In only two of the nine regressions do we reject equality at the 5% level—those where fare is the outcome. For queue length, coefficients are smaller in two of the three spillover specifications but not significantly so.

Overall, our results remain consistent regardless of whether the specification accounts for spillovers. This suggests that SUTVA violations are relatively minor in the terminal-by-survey-round fixed effects specifications, aligning with the small point estimates on the number of connected routes in Table 4. Small treatment effects on connected routes do not greatly influence our estimation strategy, but could still be important when aggregated across many routes.

Higher Order Spillovers. Our spillover specifications from the previous section capture causal effects relative to movements in the control group. If there are 'higher-order' spillovers beyond interactions at origin or destination terminals, then it could be that the control routes are also affected by entry (which would represent a violation of SUTVA). We test for one particular form of higher order spillover. While most trips in Lagos begin at route origins, some routes may interact with the public transit system at intermediate stops beyond their origin or destination.[31] In Appendix Section S5.3, we provide evidence that our estimates are not substantially affected by higher-order spillovers.

Testing Profit Equalization Across Routes. Our theory predicts that after public entry, drivers will reallocate between routes until expected profits are equalized. Appendix Section S4.2.9 tests this formally. The combination of departure, queue length, and price impacts implies a reduction of profits of approximately 0.376 log points on treated routes and 0.334 on connected routes, which are statistically indistinguishable (p=0.78). This is consistent with no significant differences in driver profits on treated versus connected routes (Table 3, column (4)).

[31] 92% of minibus passengers board at the origin (from our network mapping microdata). In the electronic ticketing data, 57% of public boardings occur at a route origin; however, on the average route the share is 88%, as the busiest routes are long, with more mid-route boarding.

6. How Commuters Value Changes in Transit Attributes

Next we estimate how commuters value changes in prices and wait times. We use quasiexperimental price variation in public buses to estimate how trips respond to prices, and an RCT to estimate how commuters value waiting.

6.1. Price Sensitivity

We first estimate the medium run price sensitivity (the combined parameters $\theta\gamma$): fares move ridership because commuters dislike paying (γ) and switch modes in response (θ). We use an event study approach on a sequence of abrupt price changes in the public transit system, in Appendix S6. We focus on 14-day windows around three events in 2023-24 in which fares were increased by 50%, increased by 33%, and decreased by 25% without altering the supply of public buses. To account for rounding differences, we instrument for route prices with event indicators. Estimates of the composite price sensitivity are similar in most events; our preferred estimate of $\theta\gamma = 0.0021$ results from pooling events. Using the average public fare of ₦220, this implies a price elasticity of approximately -0.37, which sits very close to values in the literature.[32]

6.2. Wait Time Experiment

We next measure separately the sensitivity to prices (γ) and waiting (η) by providing commuters offers to wait for random times before boarding a minibus. This experiment was preregistered in the AEA RCT registry (AEARCTR-0010283).[33]

Implementation. Our team developed a text message service which allows us to estimate commuters' value of time at minibus stops, and deployed it between June and August, 2023. We sampled 18 bus stops, and in the surrounding areas recruited 640 commuters at their homes on weekends.[34] We informed them that on weekday mornings when they commute over the next 3-5 weeks, they would find an enumerator waiting at their registered bus stop. The enumerator held a phone or tablet, which displayed a random code which changes every minute. The participant texted this code to our shortcode phone number (see Appendix Figure S8 for a photo of this interaction). For individual n at day and time t, the service sent an immediate airtime reward of $s_n^{checkin}$ for checking in, as well as a randomized offer to wait Δt_{nt} minutes for a payment of s_{nt} Naira. To accept the offer, a participant could wait at the bus stop, and after

[32] Davis (2021) estimates a price elasticity between -0.23 and -0.32 for Mexican metros, and Guzman, Gomez, and Moncada (2020) estimate -0.41 for Bogotá's BRT system. Almagro et al. (2025) estimate an average price elasticity of -0.45 at peak times for all transport modes (not just public transit).

[33] Deviations are detailed in Appendix S7.2. Additional engineering details for the software are provided in the companion paper (Björkegren et al. 2026).

[34] If we had alternately attempted to recruit at bus stops, it is likely that the only commuters who would be willing to go through our intake process would be only those with low value of time.

at least Δt_n minutes text back the new code that the enumerator's tablet then displayed. This code allowed us to verify they had waited the requisite amount of time. If the participant did not accept the offer, they could continue with their day. By observing the combinations of wait time and payments that are accepted and rejected, we are able to identify participants' value of time. Check in offers $s_n^{checkin}$ were randomized between individuals (₦200, 400, 700, or 1000) and then held fixed over the experiment. Random wait offers $(\Delta t_{nt}, s_{nt}) \sim \mathcal{O}$ were drawn for each individual-day. Individuals were told summary statistics about the distribution of wait offers at the beginning of the experiment.[35] During the first three days checking in, participants were given shorter (or no) waits and higher offers, to get them accustomed to checking in and to get a baseline on their typical arrival time at the bus stop.[36] Further implementation details are provided in Appendix S7.

A challenge to the measurement of time preference is that selection into the choice to wait may depend on how rushed a person is. In our case, on days when participants are rushed they may be less willing to check in with our experiment. The resulting choices we observe represent a selected sample, and would tend to yield estimates of the value of time that are biased downward. Selection biases are also likely to affect estimates from rideshare choices (Goldszmidt et al. 2020; Buchholz et al. 2025), since commuters may use rideshare more when rushed, which would lead to an opposite bias. A contribution of our paper is that we document and correct for this selection bias by exploiting randomized variation in $s_n^{checkin}$.

Theory. We generalize the utility function defined in equations (1) and (3) to allow for choice of both mode and waiting, in Appendix S7.3. In the experiment, a commuter has arrived at a bus stop having already chosen a travel mode, so we condition that general model on the mode choice of minibus, and consider the choice of waiting.

Because engaging with our experiment represents a hassle, we model the decision to participate. She decides whether to check in, and then whether to accept an offer to wait. Conditional on checking in, she will accept an offer to wait Δt_{nt} for payment s_{nt} if

$$\tilde{\alpha}_m - \gamma(p_m - s_{nt}) - \eta(t_m + \Delta t_{nt}) + \epsilon_{mnt+\Delta t_{nt}} > \tilde{\alpha}_m - \gamma p_m - \eta t_m + \epsilon_{mnt}, \tag{12}$$

where we omit route subscripts (ij). The idiosyncratic error $\epsilon_{mnt} = \epsilon_{mn} + \nu_{nt}$ includes two components: a preference for mode m (identical to that in equation (1)) and a preference for leaving at time t. The constant term $\tilde{\alpha}_m$ in this equation may differ from that in equation (3) since the choice of mode will be made assuming that commuters are already leaving at the

[35] They were informed, "You'll also get an opportunity to earn ₦20-1250 more by waiting between 3-25 minutes. The average offer will be ₦150 for 10 min."

[36] Participants were randomly assigned to either have one or two initial days with zero wait, and then high offers for the rest of this initial period.

optimal time, but other parameters will coincide.[37] Then the commuter will wait if

$$D^*_{nt} = \gamma s_{nt} - \eta \Delta t_{nt} + \nu^{wait}_{nt} > 0, \tag{13}$$

where we define $\nu^{wait}_{nt} = \nu_{nt+\Delta t_{nt}} - \nu_{nt}$.

Checking in incurs a time cost of $t^{checkin}$ but provides immediate payment $s^{checkin}_n$ plus any expected benefit of wait offer, captured by utility

$$C^*_{nt} = \gamma s^{checkin}_n - \eta t^{checkin} + \mathbb{E}_{\mathcal{O}}[D^*_{nt} | D^*_{nt} > 0] \Pr(D^*_{nt} > 0) + \nu^{checkin}_{nt}$$

where $\nu^{checkin}_{nt}$ captures idiosyncratic utility.

Estimation Procedure. We assume that the idiosyncratic errors are jointly normally distributed

$$\begin{bmatrix} \nu^{wait}_{nt} \\ \nu^{checkin}_{nt} \end{bmatrix} \sim N\left(\mathbf{0}, \begin{bmatrix} 1 & \rho\sigma^{checkin} \\ \rho\sigma^{checkin} & (\sigma^{checkin})^2 \end{bmatrix}\right),$$

fixing $\sigma^{wait} = 1$. In our data we observe the decision to check in $C_{nt} = \mathbb{I}\{C^*_{nt} > 0\}$ and, only if the person checks in ($C_{nt} = 1$), the decision to accept $D_{nt} = \mathbb{I}\{D^*_{nt} > 0\}$. Because $s^{checkin}_n$ is excluded from the waiting decision, we can use it as an instrument for the decision to check in. We calibrate the average hassle cost of checking in, $t^{checkin} = 1.31$ minutes, by having each enumerator time how long it took to walk to their spot, text the code, and wait for a response.[38] We estimate using maximum likelihood and cluster standard errors at the commuter level.

Results. We plot the experimental data in Figure 4. Figure 4A shows the distribution of offers and acceptances. We drew from a distribution in which shorter times were more likely and offer values tended to increase with time, but there is variation in both dimensions. Participants are much more likely to wait for shorter times and higher payments, though there is idiosyncratic variation, which depends on the individual and day. Figure 4B illustrates the selection issue. On average, respondents report planning to travel 91% of weekdays in the first week, and this does not vary by the (randomly assigned) check in offer. However, the proportion of days checked in is 55%, suggesting participants do not engage with the experiment every time they travel.[39] But participants randomly assigned higher check in offers are more likely to check in, confirming that we may use this as an instrument. Although respondents given higher check in offers check in more frequently, they are less likely to accept subsequent wait offers, which would be

[37] As shown in Appendix S7.3.

[38] Each enumerator timed the following sequence using a stopwatch: beginning at the bus stop, walk to the location where the enumerator typically stands, type in and send the random code, and then walk back to the bus stop. These estimates ranged between 28 and 144 seconds; we took the average.

[39] Participants check in on fewer days than they travel, but 90.3% checked in at least once.

Figure 4. Wait Experiment

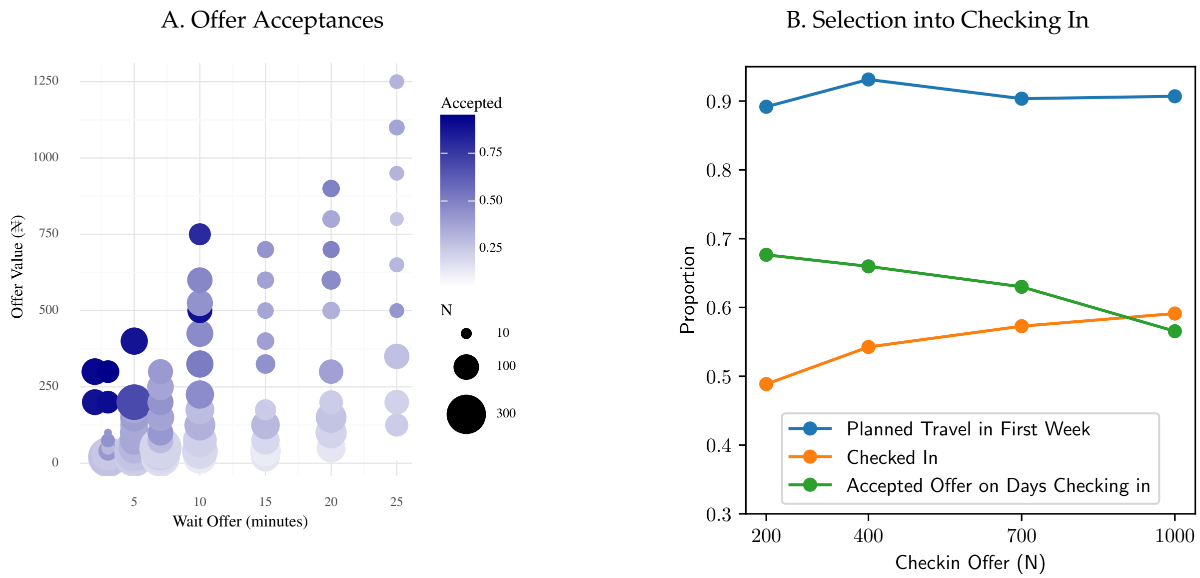

Note: Panel (A) plots offers (₦: y-axis) to wait (minutes: x-axis), with the proportion accepted shaded. The size of each point scales with the number of observations given that particular offer. Offers were randomly drawn from a weighted distribution that gave higher probability to lower times and offers. Panel (B) plots various outcomes by check in offer (₦: x-axis).

consistent with them being induced to check in on days they have particularly high value of time.

Estimated parameters using our selection-corrected maximum likelihood procedure are shown in Table 6. We obtain $\hat{\eta} = 0.0464$ and $\hat{\gamma} = 0.0025$; combined these suggest a disutility of waiting of ₦18.94 per minute. This corresponds to 3.0 times the average wage in our sample, or 2.6 times the civil service minimum wage.[40]

We estimate a positive ρ, suggesting a positive correlation of wait and checkin shocks, which is consistent with participants being less likely to both accept and check in on days they are rushed. We also find that the estimate of $\sigma^{checkin}$ is substantially larger than $\sigma^{wait} = 1$, suggesting the decision to commute and check in is more idiosyncratic than that to accept waits.

Comparison to Other Estimation Procedures. In Appendix Table S25 column (2), we show that a model that assumes perfect compliance yields an estimate of the value of time of ₦8.33 per minute: undervaluing the hassle of waiting by 56% relative to the selection-corrected estimate. This corresponds to 1.3 times the average wage in our sample or 1.1 times the average civil service wage closely following the study's conclusion.

Transportation studies commonly infer the value of time from hypothetical choices. In the

[40]Respondents reported monthly income in nine categories; we use each category's midpoint to compute an average monthly income of ₦61,289. While Lagos civil servants' minimum wage was ₦35,000 in 2023 (unchanged since 2019), it increased to ₦70,000 in 2024. We use the updated figure, as the 2023 wage likely lagged contemporary levels. Calculations assume 40 hours per week and 4 weeks per month.

Table 6. Commuting Utility Parameter Estimates

γ (utils/₦)	0.0025***
	(0.0002)
η (utils/min)	0.0464***
	(0.0031)
$\frac{\eta}{\gamma}$ (₦/min)	18.9357***
	(1.7123)
$\sigma^{checkin}$	14.3662***
	(3.901)
ρ	0.5163***
	(0.0588)
N	8640
Log Likelihood	-8720.18

Notes: Standard errors clustered at the user level. Estimation fixes $\sigma^{wait} = 1$. * p<0.1; ** p<0.05; *** p<0.01.

baseline survey we asked participants to make five hypothetical choices between options with varying costs and wait times. The value of time estimated from hypothetical choices is ₦46.46 per minute–overstating our main estimate by a factor of 2.5; see Appendix Table S25, column (3).

Many transportation planners assume a value of wait time in a range of 0.75-1.5 times the average wage. Our naïve estimate falls in that range, as do some estimates from stated preferences (Wardman, Neki, and Humphreys 2023; Small and Verhoef 2007), and (quasi)experiments (Goldszmidt et al. 2020; Buchholz et al. 2025). Our selection-corrected estimate of 3.0 times the average wage is closer to some more recent (quasi)experimental estimates: Castillo (2025) estimates roughly 4 times the average wage for ride-hail waits in Houston, and Kreindler (2024) 3.7 times the average wage for travel time in Bangalore. Waiting at minibus stops in Lagos may be unpleasant due to safety and the lack of shelter (Fan, Guthrie, and Levinson 2016).

Robustness. We assess several notions of robustness in Appendix S7.6, with alternate specifications in Appendix Table S26. When we allow η to differ by binary income category, we estimate that commuters with above median income value time at ₦20.62 per minute, relative to ₦18.17 for those with below median income (column (2)). We find slight diminishing disutility when we allow utility to include a squared term for waiting time (column (3)). Given that the heterogeneity and curvature we estimate are both slight, our main specification omits them. Estimates are similar under alternate sample definitions, shown in Appendix Table S27. We evaluate the possibility that participants arrive to the bus stop earlier to take advantage of waiting offers in Appendix S7.6.2. Estimates are also similar under a model that accounts for the possibility that, regardless of their hurry, participants accept a wait offer if a bus does not arrive; see Appendix S7.6.3.

7. Welfare Effects

Finally, we combine our estimates of treatment effects and parameters to assess how the rollout of public transit affected welfare. We compute welfare impacts on commuters following equation (9) and on minibus drivers following equation (10). Results are shown in Table 7. Clustered 95% confidence intervals are reported in parentheses and, where available, wild-bootstrap intervals in brackets. Appendix Section S4.3 provides details.

Commuters. At baseline, each commuter earns an average of $0.83 surplus per day from the preexisting system, as shown in the first row of Panel A. On treated routes, public transit increases overall surplus of commuters by $0.23 per day (28%), as shown in Column (1). The following three rows decompose this effect using equation (9). The benefit of introducing the new variety, without accounting for impacts on private transit, is $0.25 per day. But longer wait times cost travelers $0.04 per day, and lower fares would add $0.03 per day in benefits (taking our point estimates). These welfare adjustments are statistically significant when using clustered confidence intervals, with zero marginally included when using the wild bootstrap. Ultimately, commuters' disutility from increased wait times outweighs the benefits of reduced fares, so private market impacts harm treated commuters. The net loss from the private response is 7% of the total impact on treated routes.[41]

Commuters on connected routes also benefit from the price reductions we document in Table 4. Each commuter on a connected route gains $0.03 in surplus per day arising from price reductions (column (2)). Although small individually, there are around four times as many commuters on connected routes as on treated routes.

The third column aggregates these effects and reports the change in total commuter surplus in millions of dollars per month. Ignoring the private sector's response yields an estimated increase of $1.43 million per month. However, longer wait times for private transit reduce total surplus by $0.25 million per month, while lower prices on both treated and connected routes boost surplus by $0.75 million per month. Around four-fifths of the surplus from price reductions comes from connected routes, demonstrating how diffuse impacts add up across the network. On aggregate, the private sector's response raises aggregate commuter surplus from public entry by $0.50 million per month to $1.93 million per month—accounting for 26% of the total commuter welfare gains from public transit.

Drivers. While commuters benefit from public transit entry, minibus drivers face significant losses from increased competition. At baseline, each minibus driver earned an average surplus

[41] This is also the amount we would overestimate the benefits to commuters on treated routes if we ignored the private response. The 95% clustered and wild-bootstrap confidence intervals for this ratio are [-8%,31%] and [-16%,48%], respectively. We also note the quantification includes all routes that were opened, not just the sample we collected panel data on.

Table 7. Welfare Effects of Introducing Public Transit

		Individuals ($/individual/day)		**Aggregate** ($mn/month)
		Treated	Connected	
Panel A: Commuters	*Number:*	**240,353**	**877,066**	**1,117,419**
Baseline surplus from private transit		0.83 (0.63,1.23)	0.83 (0.63,1.23)	22.06 (16.66,32.61)
Effect of introducing public transit		+0.23 (0.16,0.36) [0.15,0.36]	+0.03 (0.01,0.05) [0.00,0.05]	+1.93 (1.33,2.71) [1.20,2.85]
Direct benefit of public transit		+0.25 (0.19,0.37)	0	+1.43 (1.08,2.12)
Additional impact from private Δdepartures		-0.04 (-0.08,-0.02) [-0.09,0.00]	0	-0.25 (-0.44,-0.09) [-0.52,0.02]
Additional impact from private Δprices		+0.03 (0.00,0.05) [-0.01,0.06]	+0.03 (0.01,0.05) [0.00,0.05]	+0.75 (0.32,1.15) [0.24,1.30]
Panel B: Minibus Drivers	*Number:*	**1,800**	**10,040**	**11,840**
Baseline surplus from private transit		11.87 (8.25,13.39)	11.87 (8.25,13.39)	3.44 (2.39,3.88)
Effect of introducing public transit		-4.40 (-5.10,-2.02) [-5.55,-1.24]	-4.40 (-5.10,-2.02) [-5.55,-1.24]	-1.27 (-1.48,-0.58) [-1.61,-0.36]
Accounting for private Δdepartures, ignoring route switching		-2.35 (-3.51,-0.84) [-4.11,0.15]	0	-0.10 (-0.15,-0.04) [-0.18,0.01]
Accounting for private Δprices and Δdepartures, ignoring route switching		-5.64 (-8.37,-2.27) [-9.17,-0.89]	0	-0.25 (-0.37,-0.10) [-0.40,-0.04]
Additional impact of allowing route switching		+1.24 (-0.67,4.28) [-1.69,5.72]	-4.40 (-5.10,-2.02) [-5.55,-1.24]	-1.03 (-1.17,-0.39) [-1.34,-0.21]
Panel C: Public Bus Drivers	*Number:*	**1,640**	**0**	**1,640**
Wages		-	-	+0.21
Panel D: Costs				
Operating Costs (Buses)		-	-	+2.15
Operating Costs (Terminals)		-	-	+0.11

Notes: 95% confidence intervals are constructed by drawing each of $\gamma, \eta, \theta, \sigma, \Delta t^W_{ijM}, \Delta p_{ijM}, \hat{N}^{\text{Trips}}_{ij}$ 1,000 times from a normal distribution with mean equal to its point estimate and standard deviation equal to its standard error. Square brackets report intervals constructed using wild-bootstrap standard errors for all empirical estimates for which wild-bootstrap inference is available. Driver surplus at baseline is measured using driver income net of all fees and expenses on the last day traveled in the baseline survey. Operating costs and driver wages provided by LAMATA; operating costs include fuel, maintenance and insurance for buses.

of \$11.87 per day, as shown in the first row of Panel B. Overall, drivers lose an average of \$4.40 in surplus per day (37%, second row). These losses are the same for drivers on treated and connected routes, because driver mobility equalizes expected profits.

The subsequent rows decompose this loss in three steps. If prices and route choices were fixed, the drop in demand (i.e., fewer daily trips) would lower drivers' profits on treated routes by \$2.35 per day, while drivers on connected routes would remain unaffected (since their demand is unchanged). When the price decline on treated routes is incorporated, the loss for these drivers more than doubles to \$5.64 per day while those on connected routes remain unaffected.[42] Finally, accounting for changes in routes benefits those on treated routes by \$1.24 per day, as they can shift to untreated alternatives, while it hurts drivers on connected routes by \$4.40 per day due to longer queues and lower fares. In aggregate, minibus drivers lose \$1.27 million per month, as shown in column (3) of Panel B.

Net Return. Combining commuters and minibus drivers, the net benefit is \$0.65 million. Public bus drivers earn wages of \$0.21 million per month (an upper bound for their surplus, shown in Panel C). If one includes those wages, the combined surplus of \$0.86 million falls short of the \$2.26 million in operating costs, mostly fuel expenses, reported in Panel D. These calculations exclude other factors, such as emissions, that may affect societal surplus. However, ignoring the private response would meaningfully change the net return: commuter benefits of \$1.43 million per month—with no driver losses—would come much closer to operating costs.[43]

Robustness. Online Appendix S4.4.1 evaluates results allowing for the entry of public transit to change not just the mode choice but whether trips are taken, by allowing the outside option to include not taking a trip, with the daily travel rate measured from the 2024 Nigerian Time Use Survey. The aggregate gain falls by under 1%, but this gain is distributed among more potential travelers, so the surplus per person falls, by 13% on treated routes and 12% on connected routes.

Online Appendix S4.4.2 estimates welfare effects using nested logit, which finds that transit modes are more substitutable with each other than with the outside option. Under nested logit, the welfare gains of introducing public transit fall by 3% using mode choices in our commuter survey. Only the direct benefit of the new mode changes, but the impact turns out to be small. The indirect effects of the minibus response are unchanged to first order: substitution patterns govern who switches modes, and switchers are indifferent between modes to begin with.

Appendix Table S28 presents key quantitative results using different parameter values across columns. In column (2), we use the lower value of $\sigma = 2.04$ estimated via PPML; driver results are similar. Columns (3) and (4) adjust the estimated value of θ downward and upward by 25%.

[42] We observe changes in supply but not demand on connected routes, and so assume the change in driver trips and fares on connected routes is generated by drivers switching from treated to connected routes.

[43] Mass transit systems in large cities such as London, New York, and Paris receive subsidies equal to 50-75% of operating costs, so the gap between societal benefits and costs is smaller here if we consider only commuters and about the same once drivers are accounted for.

This adjustment notably affects the estimate of the direct effect, as evidenced by the scaling of the first line in equation (9), though the results remain qualitatively similar overall. Finally, column (5) increases η, the sensitivity to waiting, by 25% relative to the estimate from our wait time experiment. In this case, commuter benefits fall by 4% because commuters are more averse to waiting.

8. Conclusion

Despite ambitious policy goals to replace decentralized private transit with centralized public transit, many developing country cities will have hybrid systems. This paper analyzes the interplay of public and private transit in Africa's largest city. We combine extensive data collected from observations, surveys, and ticketing, with quasiexperimental variation on the rollout of new public routes and an experiment on wait times with commuters at bus stops. Public entry partially displaced private minibuses, impacting fares, and causing drivers to reallocate across the system. We find that commuters have a high value of time. Altogether, we find that building a public transit system affects even commuters who continue to use the private network.

Quantifying these impacts may be important for better understanding resistance governments face from incumbents when introducing these types of reforms. Our results also suggest indirect impacts represent an important component of the total return to infrastructure investments. Our results suggest that distributional effects of government entry can be important in network markets dominated by private incumbents.

While our data is collected from a single city, our theory allows us to consider how introducing a competing transit sector may affect other settings. Generally minibus price reductions come at the expense of driver wages. In settings where drivers earn high rents, competition may lower prices without as much harm to service; but if rents are low then introducing competing transit could induce drivers to exit and reduce service. Our analysis illuminates one particular design decision: Lagos' public system serves different stops from the private system, which prevents passengers from pooling into the first passing bus, and results in larger wait times (Bly and Oldfield 1986). Shared stops may allow passengers to more easily select between competing providers. Alternately, if private operators depart on a schedule (rather than departing as a function of demand as in Lagos), that could reduce the need to wait, but may cause drivers more harm if buses run less full. Finally, the public bus system we study does not appear to have increased demand for connecting private transit, but other modes like rail that provide rapid trunk service might crowd in private connecting service.

Many further questions remain. When public transit is built, will private transit automatically rearrange into a socially useful structure (e.g., feeders), or does special care need to be taken? How should the design of public transit systems account for the service that can be provided by private networks? Are the regulations that are sufficient for a completely private system sufficient for one where private and public systems potentially connect and overlap? What combination of public and private transit best serves these growing cities?

References

Almagro, Milena, Felipe Barbieri, Juan Camilo Castillo, Nathaniel Hickok, and Tobias Salz. 2025. "Optimal Urban Transportation Policy: Evidence from Chicago." Working paper.

Andrabi, Tahir, Natalie Bau, Jishnu Das, Naureen Karachiwalla, and Asim Ijaz Khwaja. 2023. "Crowding in Private Quality: The Equilibrium Effects of Public Spending in Education."

Baum-Snow, Nathaniel. 2007. "Did Highways Cause Suburbanization?" *The Quarterly Journal of Economics* 122 (2): 775–805.

Björkegren, Daniel, Steven Reid Dickerson, Alice Duhaut, Geetika Nagpal, and Nick Tsivanidis. 2026. "Are You Willing to Wait? An SMS Platform to Measure the Value of Time in the Field." In *Proceedings of the 2026 ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies*, COMPASS '26: 96–100. New York, NY, USA: Association for Computing Machinery.

Bly, P. H., and R. H. Oldfield. 1986. "Competition between Minibuses and Regular Bus Services." *Journal of Transport Economics and Policy* 20 (1): 47–68.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2020. "Geography, transportation, and endogenous trade costs." *Econometrica* 88 (2): 657–691.

Brown, Jeffrey R., and Amy Finkelstein. 2008. "The Interaction of Public and Private Insurance: Medicaid and the Long-Term Care Insurance Market." *American Economic Review* 98 (3): 1083–1102.

Buchholz, Nicholas. 2022. "Spatial Equilibrium, Search Frictions, and Dynamic Efficiency in the Taxi Industry." *The Review of Economic Studies* 89 (2): 556–591.

Buchholz, Nicholas, Laura Doval, Jakub Kastl, Filip Matejka, and Tobias Salz. 2025. "Personalized Pricing and the Value of Time: Evidence From Auctioned Cab Rides." *Econometrica* 93 (3): 929–958.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *The Review of Economics and Statistics* 90 (3): 414–427.

Castillo, Juan Camilo. 2025. "Who Benefits from Surge Pricing?" *Econometrica* 93 (5): 1811–1854.

Cervero, Robert, and Aaron Golub. 2007. "Informal transport: A global perspective." *Transport Policy* 14 (6): 445–457.

de Chaisemartin, Clément, and Xavier D'Haultfoeuille. 2023. "Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey." *The Econometrics Journal* 26 (3): C1–C30.

Conwell, Lucas. 2026. "Privatized Provision of Public Transit." Job market paper, University College London. First version October 2022, this version March 2026.

CPCS. 2024. "Bus Industry Transition Program Diagnostic Report."

Davis, Lucas W. 2021. "Estimating the Price Elasticity of Demand for Subways: Evidence from Mexico." *Regional Science and Urban Economics* 87: 103651.

Dinerstein, Michael, Christopher Neilson, and Sebastian Otero. 2020. "The Equilibrium Effects of Public Provision in Education Markets: Evidence from a Public School Expansion Policy." *American Economic Review* 111 (10): 3376–3417.

Dinerstein, Michael, and Troy D. Smith. 2021. "Quantifying the Supply Response of Private Schools to Public Policies." *American Economic Review* 111 (10): 3376–3417.

Fan, Yingling, Andrew Guthrie, and David Levinson. 2016. "Waiting time perceptions at transit stops and stations: Effects of basic amenities, gender, and security." *Transportation Research Part A: Policy and Practice* 88: 251–264.

Feenstra, Robert. 1994. “New Product Varieties and the Measurement of International Prices.” *American Economic Review* 84 (1): 157–77.

Fréchette, Guillaume R., Alessandro Lizzeri, and Tobias Salz. 2019. “Frictions in a Competitive, Regulated Market: Evidence from Taxis.” *American Economic Review* 109 (8): 2954–2992.

Gibbons, Stephen, and Stephen Machin. 2005. “Valuing rail access using transport innovations.” *Journal of Urban Economics* 57 (1): 148–169.

Glaeser, Edward L., Matthew E. Kahn, and Jordan Rappaport. 2008. “Why do the poor live in cities? The role of public transportation.” *Journal of Urban Economics* 63 (1): 1–24.

Goldszmidt, Ariel, John A. List, Robert D. Metcalfe, Ian Muir, V. Kerry Smith, and Jenny Wang. 2020. “The Value of Time in the United States: Estimates from Nationwide Natural Field Experiments.” NBER Working Paper 28208, National Bureau of Economic Research.

Guzman, Luis A., Santiago Gomez, and Carlos A. Moncada. 2020. “Short Run Fare Elasticities for Bogotá’s BRT System: Ridership Responses to Fare Increases.” *Transportation* 47 (5): 2581–2599.

Kelley, Erin M., Gregory Lane, and David Schönholzer. 2021. “Monitoring in target contracts: Theory and experiment in Kenyan public transit.”

Kendall, David G. 1953. “Stochastic Processes Occurring in the Theory of Queues and Their Analysis by the Method of the Imbedded Markov Chain.” *The Annals of Mathematical Statistics* 24 (3): 338–354.

Kleinrock, Leonard. 1975. *Queueing Systems, Volume 1: Theory*. New York: Wiley-Interscience.

Kreindler, Gabriel. 2024. “Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications.” *Econometrica* 92 (4): 1233–1268.

Kreindler, Gabriel, Arya Gaduh, Tilman Graff, Rema Hanna, and Benjamin A. Olken. 2023. “Optimal Public Transportation Networks: Evidence from the World’s Largest Bus Rapid Transit System in Jakarta.”

Mbonu, Oluchi, and F. Christopher Eaglin. 2024. “Market Segmentation and Coordination Costs: Evidence from Johannesburg’s Minibus Networks.” *Working Paper*.

Milusheva, Sveta, Daniel Bjorkegren, and Leonardo Viotti. 2021. “Assessing Bias in Smartphone Mobility Estimates in Low Income Countries.” In *ACM SIGCAS Conference on Computing and Sustainable Societies*, COMPASS ’21: 364–378. New York, NY, USA: Association for Computing Machinery.

Mobarak, Ahmed Mushfiq, and Mark R. Rosenzweig. 2013. “Informal Risk Sharing, Index Insurance, and Risk Taking in Developing Countries.” *American Economic Review* 103 (3): 375–380.

Mohring, Herbert. 1983. “Minibuses in urban transportation.” *Journal of Urban Economics* 14 (3): 293–317.

Roberts, Mark, Maarten Bosker, Sailesh Tiwari, Putu Sanjiwacika Wibisana, Maria Monica Wihardja, and Ramda Yanurzha. 2024. “Are Ride-Hailing Services and Public Transport Complements or Substitutes? Evidence from the Opening of Jakarta’s Mrt System.”

Rosaia, Nicola. 2024. “Competing Platforms and Transport Equilibrium.” Working Paper.

Small, Kenneth A., and Erik T. Verhoef. 2007. *The Economics of Urban Transportation*. London and New York: Routledge.

Sun, Liyang, and Sarah Abraham. 2021. “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects.” *Journal of Econometrics* 225 (2): 175–199.

The World Bank. 2025a. “Transport.” https://www.worldbank.org/en/topic/transport/overview#2. Accessed February 11, 2025.

The World Bank. 2025b. “Urban Development.” https://www.worldbank.org/en/topic/urbandevelopment/overview#2. Accessed February 11, 2025.

Tsivanidis, Nick. 2026. “Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá’s TransMilenio.” *American Economic Review* 116 (2): 418–463.

Walters, A. A. 1982. “Externalities in urban buses.” *Journal of Urban Economics* 11 (1): 60–72.

Wardman, Mark, Kazuyuki Neki, and Richard Martin Humphreys. 2023. “Meta-Analysis of the Value of Travel Time Savings in Low- and Middle-Income Countries.” Mobility and Transport Connectivity

Series, World Bank.
Zárate, Román David. 2024. “Spatial Misallocation, Informality, and Transit Improvements: Evidence from Mexico City.” Working paper.

Online Appendix: Public and Private Transit

Contents

Appendix S1. Additional Figures Referenced in Paper

Figure S1. Minibus Queues

A. Terminal (Berger)

B. Roadside

Figure S2. Timeline of Data Collection

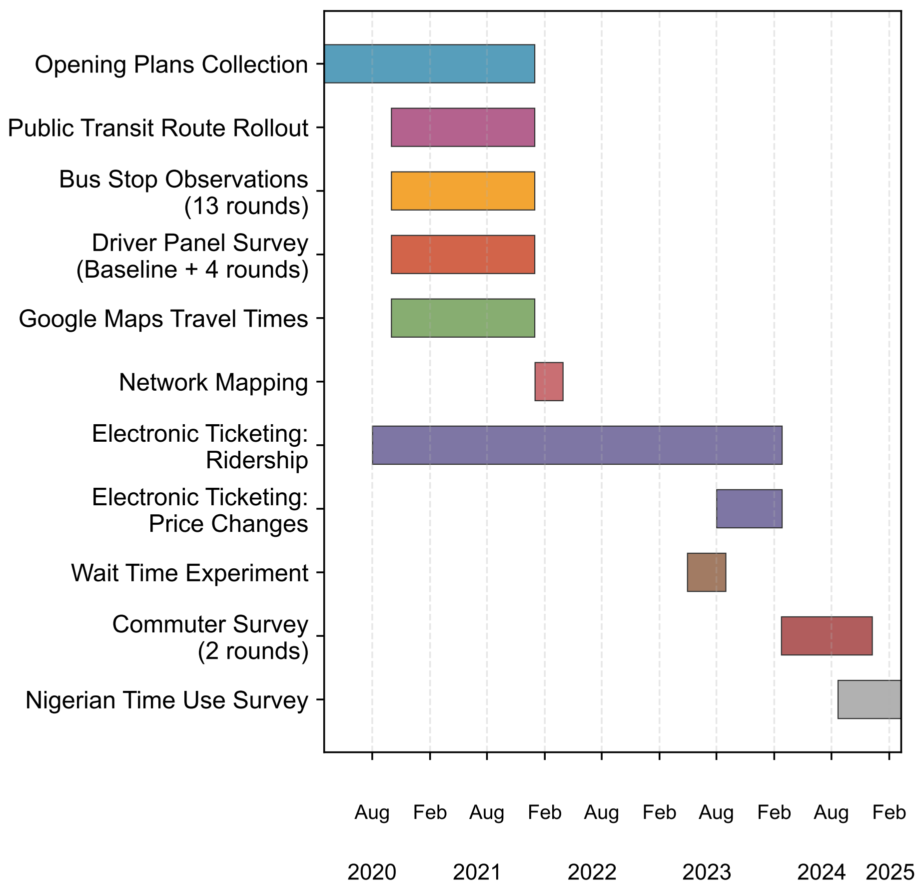

Figure S3. Depiction of Treatment Definitions

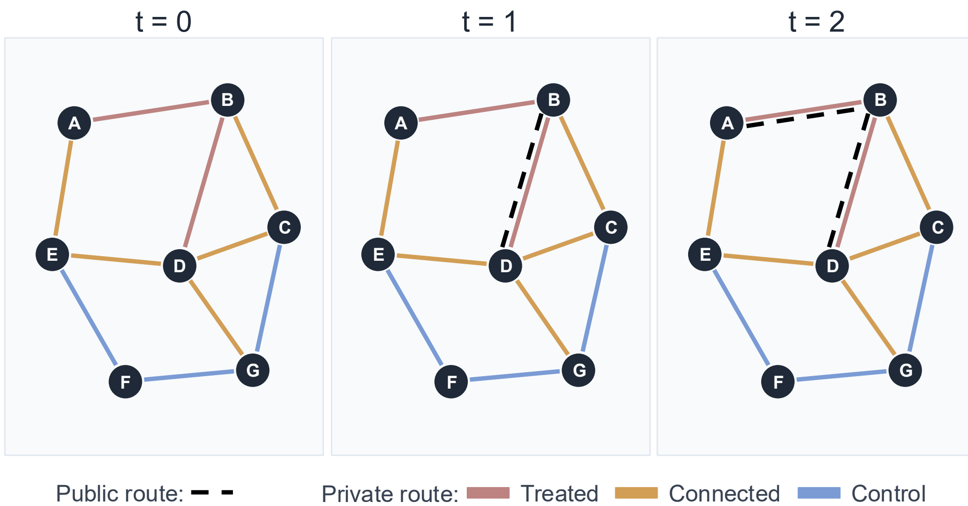

Note: Private routes are classified as treated if they are ever served by a public route during our time period, as connected if they are ever connected to a public route during our time period, or control if neither. While those definitions are static across time, we also define the variables $Open_{rt}$ to denote whether route r is served by public transit as of period t, and Connections to public transit$_{rt}$ to quantify whether r has connections to public routes as of period t and is not directly served.

Figure S4. Google Mobility Data: Visits to Transit Stations

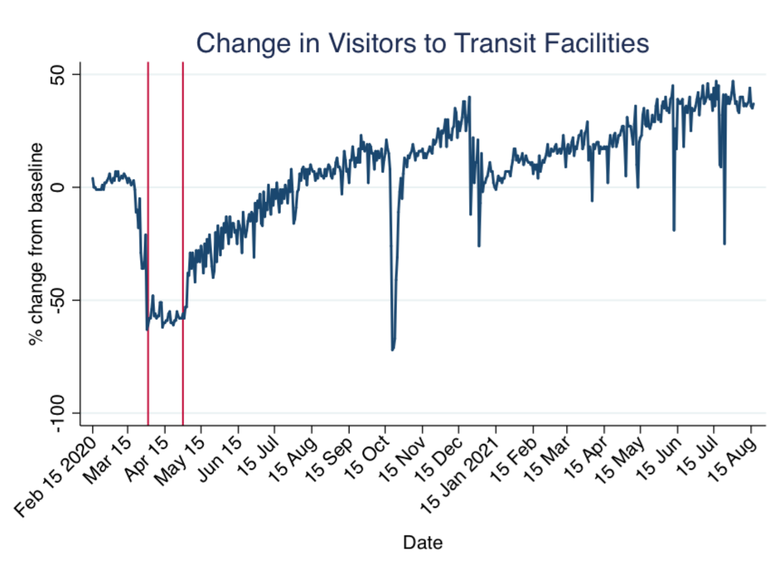

Figure S5. Event Study of Public Transit Entry

A. Minibus Departures

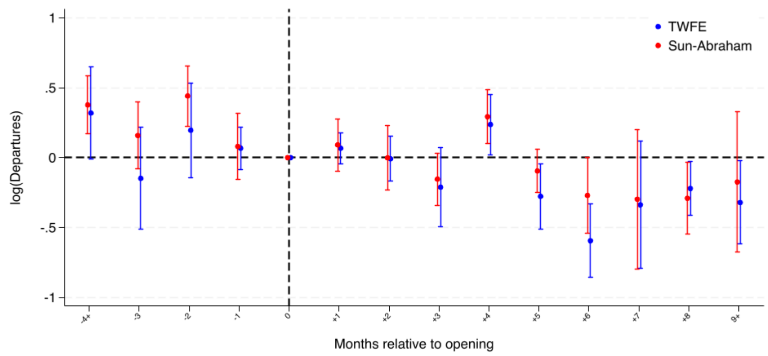

B. Minibus Fares

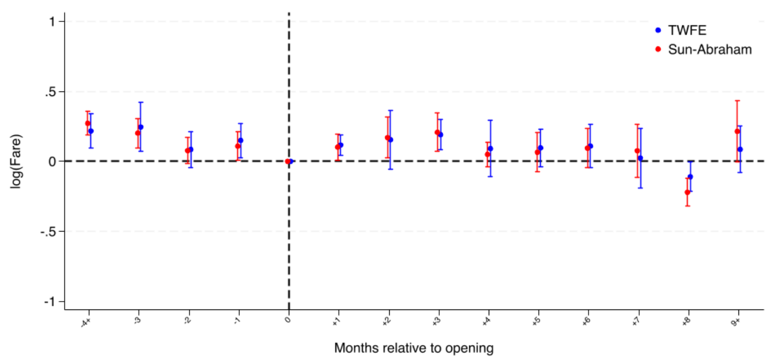

C. Minibus Driver Queues

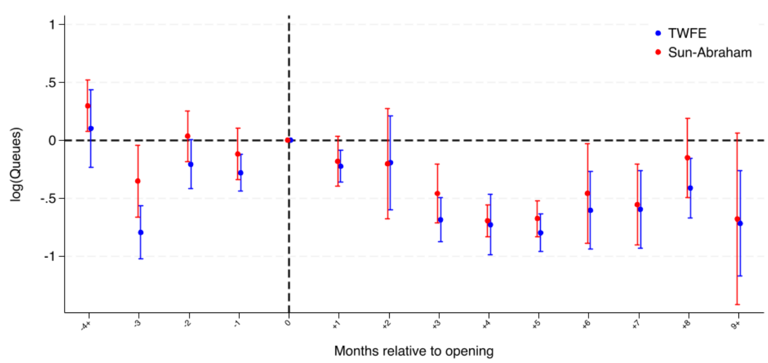

Note: Each panel reports coefficient estimates from a regression which replicates the main specification (column (7) of Table 2) but replaces the singular Open$_{rt}$ with a set of month-to-treatment dummies which reflect the month of observation relative to the month of opening. Blue points report coefficients and 95% confidence intervals from a TWFE regression, while red points represent those from the Sun and Abraham (2021) estimator.

Figure S6. Share of Drivers Not Working, by Survey Round

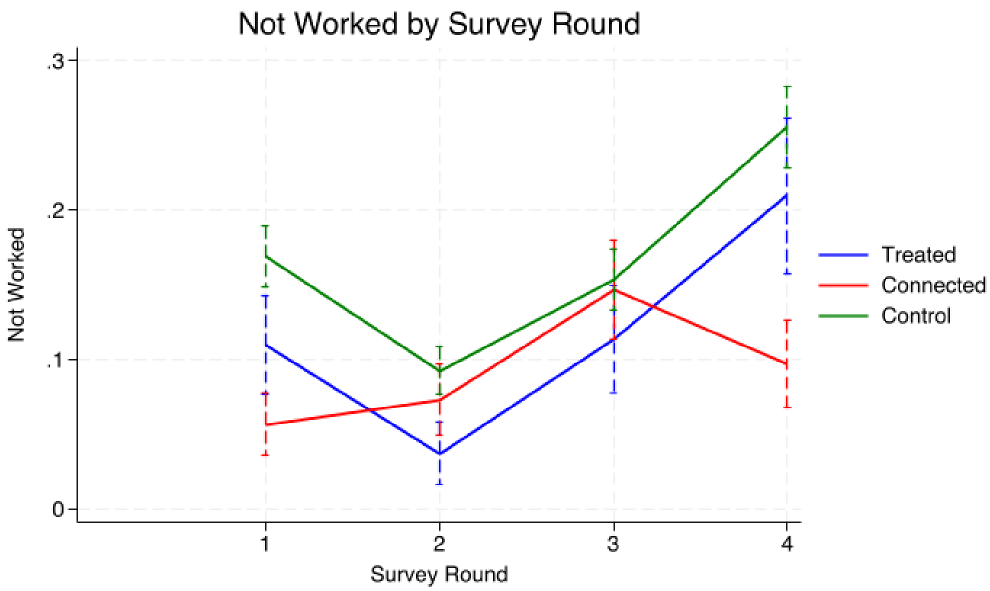

Note: The figure shows the share of surveyed drivers who report not having worked in the past seven days, by survey round, for drivers whose baseline main route becomes treated, is connected to a treated route, or is control.

Figure S7. Daily Rental Value for Minibuses

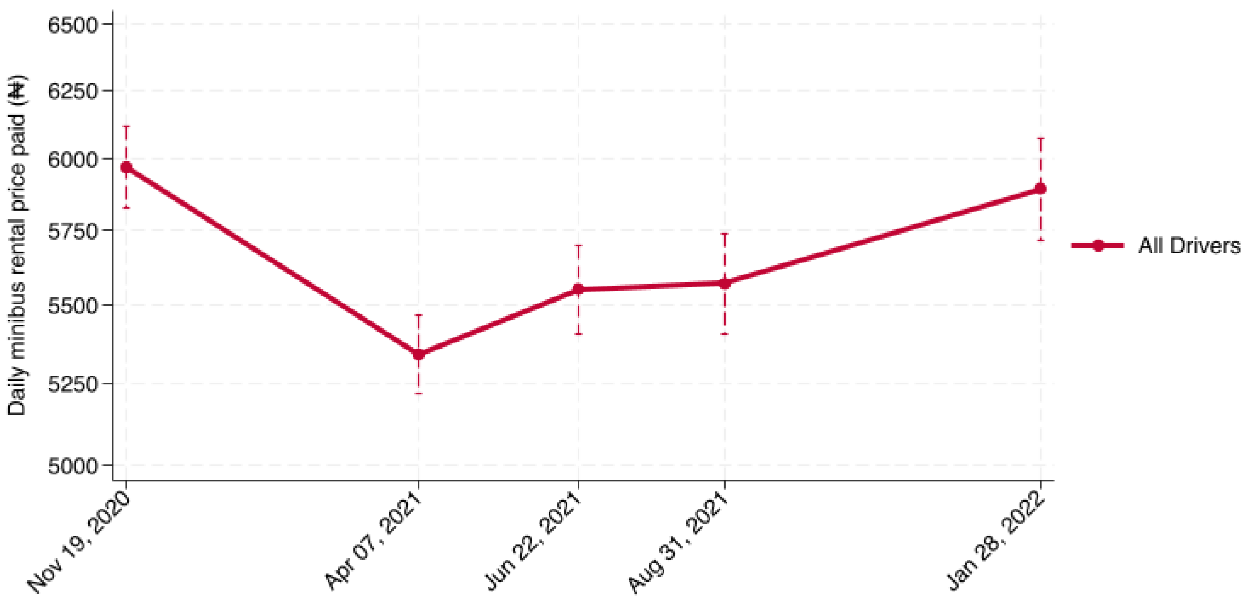

Note: The figure shows the average daily rental price paid by drivers to bus owners, as reported in the driver survey. The sample includes only drivers who rent rather than own their minibuses. Dates on the x-axis correspond to the median survey date for each round, from baseline through follow-up (FUP) 4.

Figure S8. Participant Checking in for Wait Experiment

Appendix S2. Additional Tables Referenced in Paper

Table S1. Public vs. Private Service Characteristics

	Public Transit	Private Transit	Difference Public - Private
Comfort	4.102	3.053	1.050***
	(0.057)	(0.082)	(0.100)
Safety	4.043	3.122	0.922***
	(0.046)	(0.060)	(0.075)
Reliability	3.851	3.031	0.820***
	(0.052)	(0.061)	(0.080)

Notes: Sample consists of commuters living within 2.5km of 10 major transit terminals, surveyed in 2024. Eligibility was restricted to regular users with at least two transit trips in the prior week. Commuters report characteristics of their longest trip taken the previous day, and characteristics of the alternate modes they could have used instead. The table reports means of characteristics ratings on a scale of 1–5. Ratings pool own-mode reports and hypothetical alternate-mode reports elicited for the other mode. Standard errors are in parentheses. * $p<0.1$; ** $p<0.05$; *** $p<0.01$.

Table S2. Baseline Average Characteristics of Private Transit Routes by Planned Rollout Timing

	Short-Term Network		Not Treated		
			Long-Term Network		Incidentally
	Treated	Not Treated, Dated	Planned, No Date	Later Phase	Observed
Departures (every 30 mins)	4.11	4.09	2.76	3.64	3.02
	(1.11)	(0.45)	(0.24)	(0.28)	(0.51)
Passengers at Departure	13.49	13.10	12.89	11.14	9.85
	(1.69)	(0.99)	(0.97)	(0.61)	(0.75)
Fare (₦)	246.15	207.59	292.79	201.62	215.08
	(30.62)	(14.35)	(25.42)	(12.27)	(30.24)
Distance (in km)	8.61	6.81	9.30	6.76	10.06
	(1.75)	(0.70)	(1.22)	(0.59)	(2.16)
Queue Length	4.90	5.31	4.69	4.71	4.57
	(0.75)	(0.51)	(0.47)	(0.32)	(0.54)
Min Dist to CBD (km)	11.98	10.10	8.11	9.16	11.07
	(1.88)	(0.82)	(0.87)	(0.71)	(1.43)
Max Dist to CBD (km)	18.65	14.76	14.86	13.84	18.34
	(0.86)	(0.83)	(1.36)	(0.81)	(2.10)
Min Pop Density (500m)	11060.90	13812.97	12952.79	10087.62	10376.46
	(1346.79)	(859.27)	(718.48)	(542.14)	(869.75)
Max Pop Density (500m)	18783.64	22723.85	21248.62	18197.85	19298.26
	(2194.37)	(1045.49)	(710.26)	(867.90)	(1460.22)
Number of Routes	13	58	52	119	36

Notes: Table reports means and standard errors of baseline characteristics for routes in our sample. Groups are mutually exclusive by construction. Treated routes opened during our sampling period. All Not Treated routes are further split by whether they had a firm rollout timeline (Short-Term Network) or not (Long-Term Network). "Not Treated, Dated" routes were assigned a planned opening date but had not opened as of the end of sampling. "Planned, No Date" routes appear on a government's initial deployment plan but were never assigned a firm date. "Later Phase" routes appear on future phases' deployment plans. "Incidentally Observed" routes were not part of our planned sample, and are routes that enumerators could observe and record while collecting data on sampled routes.

Table S3. Baseline Balance of Private Transit Routes by Rollout Status

	Treated vs.			
	Dated	Planned, No Date	Later Phase	Incidentally Observed
Departures (every 30 mins)	0.02	1.35*	0.47	1.09
	(1.09)	(0.72)	(0.92)	(1.08)
Passengers at Departure	0.39	0.60	2.35	3.64**
	(2.32)	(2.18)	(1.99)	(1.62)
Fare (₦)	38.57	-46.63	44.53	31.08
	(33.60)	(53.29)	(38.39)	(53.78)
Distance (in km)	1.80	-0.69	1.85	-1.45
	(1.68)	(2.60)	(1.88)	(3.76)
Queue Length	-0.42	0.21	0.19	0.32
	(1.13)	(1.02)	(1.01)	(1.00)
Min Dist to CBD (km)	1.89	3.87*	2.83	0.91
	(1.95)	(1.98)	(2.23)	(2.64)
Max Dist to CBD (km)	3.89**	3.79	4.81*	0.31
	(1.81)	(2.76)	(2.47)	(3.56)
Min Pop Density (500m)	-2752.07	-1891.89	973.28	684.44
	(1928.05)	(1587.87)	(1701.35)	(1659.41)
Max Pop Density (500m)	-3940.21	-2464.97	585.79	-514.62
	(2440.62)	(1785.61)	(2726.79)	(2767.36)
Number of Routes	58	52	119	36
F-stat	0.91	1.86	0.86	1.58

Notes: Each column reports the OLS coefficient on a group indicator from a regression of the row characteristic on that indicator, restricted to Treated routes and the named comparison group only, with the standard error in parentheses. Treated routes opened during our sampling period. "Dated" routes have a firm rollout timeline and an assigned opening date but had not opened as of the end of sampling. "Planned, No Date" routes appear on a government's initial deployment plan but were never assigned a firm date. "Later Phase" routes appear on future phases' deployment plans. "Incidentally Observed" routes were not part of our planned sample and were recorded by enumerators while collecting data on sampled routes. Groups are mutually exclusive by construction and are defined as in Table S2. The F-stat row reports the F-statistic from a test of joint equality of all characteristics between Treated routes and the named comparison group. * $p<0.1$; ** $p<0.05$; *** $p<0.01$.

Table S4. Minibus Drivers

	Weighted		Unweighted		N
	Mean	SD	Mean	SD	
Demographics					
Age (years)	43.40	8.80	43.69	8.82	849
Male	0.993	0.09	0.994	0.08	849
Has secondary education	0.611	0.49	0.603	0.49	849
Driving experience (years)	12.65	9.67	12.75	9.71	847
Operations					
Own bus	0.509	0.50	0.514	0.50	847
Choose own route	0.846	0.36	0.848	0.36	849
Number of routes plied	1.185	0.54	1.173	0.52	849
Proportion of trips skipping terminal start ('sole')	0.209	0.27	0.170	0.25	849
Number of days worked last week	5.622	0.93	5.631	0.89	849
Number of trips	9.201	5.82	9.092	5.76	849
Total hours between start and stop on last day worked	12.13	3.42	12.16	3.39	849
Proportion of work time spent driving	0.576	0.16	0.567	0.16	847
Proportion of work time spent queuing	0.424	0.16	0.433	0.16	847
Disputes					
Proportion of drivers in weekly passenger dispute	0.477	0.50	0.484	0.50	849
Proportion of drivers in weekly driver dispute	0.366	0.48	0.367	0.48	849
Income and Costs					
Last day worked					
Total revenue from all trips (₦)	17269.4	10953.98	17144.6	10818.89	811
Net income (₦)	5140.1	2913.87	5167.6	2951.67	818
Cost of fuel (₦)	3766.1	1581.31	3740.1	1606.35	846
Fees paid to the association (₦)	691.0	1525.38	705.9	1557.55	849
Employed conductor	0.306	0.46	0.295	0.46	843
If employed conductor, amount paid (₦)	2531.4	1024.76	2551.5	1028.85	237
If do not own vehicle, amount paid to the minibus owner (₦)	5972.8	2932.65	6021.2	2956.19	410
Monthly (last month)					
Had a repair	0.919	0.27	0.923	0.27	844
If had a repair, cost of last repair (₦)	11538.8	24666.40	11497.3	24022.27	769
Had a fine	0.830	0.38	0.827	0.38	480
If had a fine, cost of last fine (₦)	6505.3	14070.05	6463.2	14319.23	376
One-time fees (terminal level)					
Average terminal registration fee (₦)*			15494.13	7486.50	67

Notes: Survey of minibus drivers, sampled in queues and outside terminals. Statistics in columns (1)-(2) are weighted to account for differential ease of sampling those waiting in queues, as described in Section S3.1. Proportion of work time spent driving and queuing is based on responses to trip diary, which collects the first 8 trips on the last day worked. Number of trips and total revenue are winsorized at 99th percentile due to the presence of large outliers. *Average terminal registration fee is the current registration fee reported using a terminal survey, for terminals where drivers in the sample are registered, hence the smaller sample since this is recorded at the terminal-level.

Table S5. Minibus Driver and Commuter Characteristics

	Drivers	Commuters
Age (years)	43.4	37.6
	(0.3)	(0.4)
Male	0.993	0.515
	(0.003)	(0.022)
Has secondary education	0.611	0.943
	(0.017)	(0.010)
Monthly income (₦, adjusted to 2024 levels)	268,620	105,694
	(5,899)	(4,564)
N	849	540

Notes: Sample means with standard errors are reported in parentheses. Driver statistics are from the baseline survey (November – December 2020), weighted to account for differential ease of sampling those waiting in queues, as described in Section S3.1. Commuter statistics are from the commuter survey we ran in March and November 2024, described in Online Appendix S4.4.2. Monthly driver income is the net income on last day worked multiplied by the days driven in the previous week, scaled to a monthly figure assuming four weeks per month. Driver income has been inflation adjusted to 2024 levels using the ratio of World Bank Nigeria CPI (FP.CPI.TOTL) annual averages for 2024 and 2020. Monthly commuter income is calculated using the midpoint of an eight-category self-reported income variable.

Table S6. Minibus Drivers: Awareness of Public Transit

	Mean	SD	N
Expectations			
Have heard of public transit rollout (by 2020)	0.572	(0.50)	849
If heard and predicted public transit will open on surveyed route, correct*	0.222	(0.44)	9
If heard and predicted public transit will not open on surveyed route, correct*	0.816	(0.39)	234
If heard, expect public transit will not affect them	0.476	(0.50)	246
If heard, expect public transit to reduce their passengers	0.362	(0.48)	246

Notes: Survey of minibus drivers, sampled in queues and outside terminals. Statistics are not weighted to account for differential ease of sampling those waiting in queues. *Calculated only for those drivers who answered and whose main route had not yet seen the introduction of public transit. Of the 486 drivers who were aware of the public transit rollout, 34% had already seen public transit services introduced along their route by the time of the survey. The last two questions were asked to a random 50% of the respondents who had heard of the public transit rollout.

Table S7. Semi-Elasticity of Total Transit Ridership to Public Entry

	(1)
Open	0.149*
	(0.088)
N	6,041
R^2	0.65

Notes: An observation is a route by survey round by time-of-day window (morning peak, midday off-peak, afternoon peak). The outcome is the log of the total number of transit passengers per hour on the route, summing minibus and public bus boardings. Open is a dummy equal to one once public transit operates on the route, as in our main specification (Section 5.2). The specification includes route-by-window fixed effects and survey-round fixed effects interacted with terminal, trip-distance-quartile, and origin- and destination-LGA fixed effects. Standard errors two-way clustered by route and terminal are reported in parentheses. * p<0.1; ** p<0.05; *** p<0.01.

Table S8. Baseline Average Characteristics of Private Transit Routes

	Treated	Connected	Control	Difference T-Ctd	Difference T-Ctr
	Receives public transit	Will share endpoint with public transit	No public transit		
Departures per 30 mins	4.11	3.60	3.38	0.51	0.74
	(1.11)	(0.27)	(0.24)	(0.91)	(0.87)
Passengers at Departure	13.49	10.15	13.10	3.34	0.39
	(1.69)	(0.50)	(0.62)	(1.67)	(2.20)
Fare (₦)	246.15	203.75	239.03	42.41	7.13
	(30.62)	(14.78)	(11.49)	(46.46)	(39.12)
Distance (km)	8.61	7.50	7.91	1.11	0.70
	(1.75)	(0.86)	(0.54)	(2.71)	(1.85)
Driver Queue Length	4.90	5.30	4.56	-0.40	0.34
	(0.75)	(0.36)	(0.30)	(1.14)	(1.02)
Driver Queue Length > 0	0.88	0.88	0.84	-0.00	0.04
	(0.08)	(0.02)	(0.02)	(0.07)	(0.08)
Number of Routes	13	123	142		
F-Stat				1.26	0.41

Notes: Table reports means and standard errors of baseline characteristics for routes which received public service in our sample (Treated), those which share a node with a route which received public service (Connected) and those which did not (Control). Fares and passenger counts are reported for routes where buses are observed in our sample. Distance is computed on the straight line between start and end point of the route. Driver queue length is the number of buses waiting in the queue for a given route at the beginning of each time period.

Table S9. Baseline Average Characteristics of Terminals

	Treated	Control	Difference T-C
	Receives public transit	No public transit	
Total Departures	314.60	135.96	178.63
	(110.77)	(17.82)	(61.51)
Total Passengers across all Departures	3983.40	1292.16	2691.25
	(1474.30)	(178.43)	(731.62)
Number of Private Transit Routes Operational	10.86	8.05	2.80
	(2.51)	(0.84)	(2.21)
Number of Terminals	7	38	
F-stat			5.26***

Notes: Table reports means and standard errors of baseline characteristics of terminals which received any public service in our sample (Treated) and those which did not (Control). The F-stat row reports the F-statistic from a test of joint equality of all characteristics between Treated and Control terminals. * p<0.1; ** p<0.05; *** p<0.01.

Table S10. Effect of Public Transit on Private Transit (Linear Outcomes)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Departures							
Open	-0.383	-0.461	-0.469	-0.505	-0.467	-0.559	-0.730
	(0.027)	(0.000)	(0.000)	(0.000)	(0.001)	(0.007)	(0.003)
	[0.233]	[0.052]	[0.035]	[0.039]	[0.053]	[0.103]	[0.079]
N	24,601	24,601	24,601	24,601	24,601	24,601	24,601
R^2	0.57	0.62	0.63	0.63	0.63	0.63	0.65
Mean Outcome	3.45	3.45	3.45	3.45	3.45	3.45	3.45
Panel B: Fare							
Open	-4.007	-5.693	-2.321	-8.114	-8.422	-9.820	-0.973
	(0.716)	(0.673)	(0.865)	(0.523)	(0.500)	(0.321)	(0.907)
	[0.724]	[0.699]	[0.887]	[0.570]	[0.551]	[0.368]	[0.923]
N	23,070	23,070	23,070	23,070	23,070	23,070	23,070
R^2	0.93	0.94	0.94	0.95	0.95	0.95	0.95
Mean Outcome	209.51	209.51	209.51	209.51	209.51	209.51	209.51
Panel C: Queue							
Open	-0.812	-1.022	-0.927	-0.877	-0.849	-0.885	-0.978
	(0.030)	(0.056)	(0.044)	(0.065)	(0.078)	(0.033)	(0.047)
	[0.180]	[0.270]	[0.201]	[0.230]	[0.222]	[0.154]	[0.175]
N	24,591	24,591	24,591	24,591	24,591	24,591	24,591
R^2	0.53	0.59	0.60	0.60	0.61	0.62	0.63
Mean Outcome	4.75	4.75	4.75	4.75	4.75	4.75	4.75
Route × Period FE	X	X	X	X	X	X	X
Day of Week × Survey Round FE	X	X	X	X	X	X	X
Hour of Dep × Survey Round FE	X	X	X	X	X	X	X
Terminal × Survey Round FE		X	X	X	X	X	X
Trip Dist Controls × Survey Round FE			X	X	X	X	X
Dep. Plan × Survey Round FE				X	X	X	X
CBD Controls					X		
O & D Lat-Lon Poly × Survey Round FE						X	
O & D LGA × Survey Round FE							X

Notes: P-values clustered by route and terminal are reported in parentheses. Wild p-values computed using Cameron, Gelbach, and Miller (2008) are reported in square brackets.

Table S11. Effect of Public Transit on Traffic Congestion

	log(Speed)		log(Time)	
	(1)	(2)	(3)	(4)
Open	0.000		0.007	
	(0.023)		(0.026)	
Open $<$ 3 months		-0.002		0.004
		(0.021)		(0.024)
Open 3-6 months		-0.005		0.012
		(0.026)		(0.027)
Open $>$6 months		0.008		0.005
		(0.027)		(0.029)
N	7,362,322	7,362,322	7,362,322	7,362,322
R^2	0.63	0.63	0.86	0.86

Notes: Outcome is either log speed or log trip time. Data is from Google Maps, standard errors clustered at the route-level (288 routes). Controls include route fixed effects, and week-year, hour-month-year and trip distance quartile-month-year fixed effects.

Table S12. Minibus Driver Survey Observations

	Number of Respondents
Baseline	854
Follow up 1	564
Follow up 2	528
Follow up 3	514
Follow up 4	423

Notes: Number of respondents for each round of the minibus driver survey.

Table S13. Minibus Driver Survey Attrition

	Proportion of Sample Present
5 Surveys	0.342
4 or more Surveys	0.550
3 or more Surveys	0.663
2 or more Surveys	0.821
1 or more Surveys	1.000

Notes: Proportion of sampled minibus drivers who respond to at least given number of survey rounds, in minibus driver survey.

Table S14. Minibus Driver Survey Attrition Correlates

	By driver			By driver-survey-round		
	No follow ups			Completed		
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Age)	-0.143**		-0.149**	0.121**		0.119**
	(0.065)		(0.065)	(0.053)		(0.053)
Own Vehicle	-0.006		-0.001	0.030		0.029
	(0.028)		(0.028)	(0.023)		(0.023)
Log(Income)	0.000		-0.005	-0.026		-0.026
	(0.022)		(0.023)	(0.019)		(0.019)
Main Route Treated		-0.034	-0.043		-0.005	0.006
		(0.031)	(0.031)		(0.025)	(0.026)
N	808	847	806	4,040	4,235	4,030
R^2	0.01	0.00	0.01	0.01	0.00	0.00
Mean of Dependent Variable	.175	.179	.174	.68	.676	.68

Notes: In columns 1 to 3, outcome is a dummy for whether the driver does not appear in any follow up survey. Only data from the baseline is used. In columns 4 to 6, data is at the driver-survey level and the outcome is a dummy for whether a driver completed a survey. Right hand side variables are log of driver age, a dummy for whether they own their vehicle, log income on last day driving (all from the baseline), and a dummy for whether the driver's main route in the baseline is treated at any point in the data. Standard errors clustered by driver reported in parentheses. * $p<0.1$; ** $p<0.05$; *** $p<0.01$.

Table S15. Minibus Driver Survey Attrition Correlates By Treatment Group

	Completed by driver-survey-round					
	Treated vs Connected			Treated vs Control		
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Age)	0.117*		0.118*	0.036		0.017
	(0.068)		(0.069)	(0.091)		(0.091)
Own Vehicle	0.013		0.015	0.033		0.035
	(0.029)		(0.029)	(0.037)		(0.037)
Log(Income)	-0.066***		-0.067***	-0.016		-0.022
	(0.025)		(0.025)	(0.035)		(0.035)
Main Route Treated		-0.047	-0.045		0.067*	0.080**
		(0.031)	(0.032)		(0.038)	(0.039)
N	1,628	1,712	1,628	1,996	2,084	1,988
R^2	0.01	0.00	0.01	0.00	0.00	0.01
Mean of Dependent Variable	.778	.775	.778	.48	.475	.481

Notes: Data is at the driver-survey level and the outcome is a dummy for whether a driver completed a survey. Right hand side variables are log of driver age, a dummy for whether they own their vehicle, log income on last day driving (all from the baseline), and a dummy for whether the driver's main route in the baseline is treated at any point in the data. Standard errors clustered by driver reported in parentheses. * $p<0.1$; ** $p<0.05$; *** $p<0.01$.

Table S16. Minibus Rental Fees and Public Transit Entry

	Log (Daily minibus rental)
Open	0.052
	(0.087)
Number of routes open at terminal	-0.012
	(0.009)
N	1,026
R^2	0.52

Notes: Data is at the driver-survey level and the outcome is the log of the daily minibus rental fee reported by the driver. The specification includes driver fixed effects, and survey round fixed effects interacted with dummies for each quartile of driver age, the number of trips the driver reports in the previous workday at baseline, and a dummy for whether they own their vehicle at baseline. Open and Number of routes open at terminal are defined analogously to the terminal observation specifications, but based on a driver's main route from the baseline survey. Regressions are weighted by the sampling weights discussed in the paper. Standard errors clustered by driver and terminal of recruitment reported in parentheses.

Table S17. Placebo Tests

	log(Departures)			log(Fare)			log(Queues)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Open	-0.221	-0.221	-0.208	-0.024	-0.024	-0.029	-0.289	-0.289	-0.282
	(0.005)	(0.005)	(0.010)	(0.501)	(0.495)	(0.412)	(0.000)	(0.000)	(0.000)
	[0.060]	[0.061]	[0.070]	[0.585]	[0.580]	[0.500]	[0.033]	[0.033]	[0.035]
Open ENDSARS		-0.012			0.185			0.087	
		(0.938)			(0.085)			(0.603)	
		[0.955]			[0.440]			[0.782]	
Open Cancelled			0.088			-0.036			0.044
			(0.273)			(0.107)			(0.665)
			[0.485]			[0.137]			[0.765]
N	21,287	21,287	21,287	23,070	23,070	23,070	22,293	22,293	22,293
R^2	0.66	0.66	0.66	0.95	0.95	0.95	0.59	0.59	0.59

Notes: Specification is the same as column 7 of Table 2. P-values clustered by route and terminal are reported in parentheses. Wild p-values computed using Cameron, Gelbach, and Miller (2008) are reported in square brackets.

Table S18. Placebo Tests: Spillover Specification

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: log(Departures)							
Open	-0.221***	-0.221***	-0.172**	-0.161*	-0.208**	-0.166**	-0.158*
	(0.076)	(0.076)	(0.067)	(0.090)	(0.078)	(0.066)	(0.091)
Open ENDSARS		-0.012	0.109	0.112			
		(0.150)	(0.146)	(0.151)			
Number of open routes at terminal				0.002			0.002
				(0.013)			(0.013)
Open cancelled					0.088	0.056	0.055
					(0.079)	(0.072)	(0.069)
N	21,287	21,287	21,287	21,287	21,287	21,287	21,287
R^2	0.66	0.66	0.63	0.63	0.66	0.63	0.63
Panel B: log(Fare)							
Open	-0.024	-0.024	-0.030	-0.108**	-0.029	-0.034	-0.111**
	(0.035)	(0.035)	(0.037)	(0.052)	(0.036)	(0.038)	(0.052)
Open ENDSARS		0.185*	0.068	0.050			
		(0.105)	(0.075)	(0.073)			
Number of open routes at terminal				-0.016***			-0.016***
				(0.006)			(0.006)
Open cancelled					-0.036	-0.032	-0.025
					(0.022)	(0.020)	(0.020)
N	23,070	23,070	23,070	23,070	23,070	23,070	23,070
R^2	0.95	0.95	0.93	0.94	0.95	0.93	0.94
Panel C: log(Queues)							
Open	-0.289***	-0.289***	-0.296***	-0.160*	-0.282***	-0.296***	-0.164
	(0.072)	(0.073)	(0.054)	(0.094)	(0.075)	(0.058)	(0.098)
Open ENDSARS		0.087	0.148	0.178			
		(0.166)	(0.170)	(0.181)			
Number of open routes at terminal				0.029*			0.029*
				(0.017)			(0.017)
Open cancelled					0.044	0.010	-0.002
					(0.101)	(0.099)	(0.105)
N	22,293	22,293	22,293	22,293	22,293	22,293	22,293
R^2	0.59	0.59	0.55	0.55	0.59	0.55	0.55
Terminal X Survey Round FE	X	X			X		

Notes: Columns (1), (2) and (5) replicate each column from Table S17. Columns (3) and (6) run the same placebo specifications for the ENDSARS and Cancelled placebos, in the spillover specification without terminal-by-survey round fixed effects. Columns (4) and (7) then add to this the measure of public transit connections. Standard errors clustered by route and terminal are reported in parentheses. * p<0.1; ** p<0.05; *** p<0.01.

Table S19. Effect of Public Transit on Private Transit: Sun–Abraham Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: log(Departures)							
Open	-0.127*	-0.202***	-0.186***	-0.195***	-0.196***	-0.220***	-0.210***
	(0.072)	(0.055)	(0.067)	(0.064)	(0.069)	(0.079)	(0.078)
N	21,284	21,284	21,284	21,284	21,284	21,284	21,284
R^2	0.58	0.64	0.64	0.64	0.64	0.65	0.66
Mean Outcome	3.45	3.45	3.45	3.45	3.45	3.45	3.45
Panel B: log(Fare)							
Open	-0.008	-0.022**	-0.005	-0.022**	-0.026**	-0.027**	-0.001
	(0.020)	(0.011)	(0.011)	(0.011)	(0.012)	(0.014)	(0.026)
N	23,067	23,067	23,067	23,067	23,067	23,067	23,067
R^2	0.92	0.94	0.94	0.94	0.94	0.94	0.95
Mean Outcome	209.51	209.51	209.51	209.51	209.51	209.51	209.51
Panel C: log(Queues)							
Open	-0.246***	-0.280***	-0.249***	-0.250***	-0.233***	-0.249***	-0.230***
	(0.037)	(0.038)	(0.047)	(0.048)	(0.054)	(0.062)	(0.078)
N	22,290	22,290	22,290	22,290	22,290	22,290	22,290
R^2	0.47	0.55	0.56	0.56	0.56	0.57	0.59
Mean Outcome	4.75	4.75	4.75	4.75	4.75	4.75	4.75
Route $\times$ Period FE	X	X	X	X	X	X	X
Day of Week $\times$ Survey Round FE	X	X	X	X	X	X	X
Hour of Dep $\times$ Survey Round FE	X	X	X	X	X	X	X
Terminal $\times$ Survey Round FE		X	X	X	X	X	X
Trip Dist Controls $\times$ Survey Round FE			X	X	X	X	X
Dep. Plan $\times$ Survey Round FE				X	X	X	X
CBD Controls					X		
O & D Lat-Lon Poly $\times$ Survey Round FE						X	
O & D LGA $\times$ Survey Round FE							X

Notes: Estimates use the interaction-weighted Sun and Abraham (2021) estimator, with cohort defined by calendar month of route opening and never-treated routes as the control cohort. The rows at the bottom of the table reflect control variables added to each specification. Route $\times$ Period FE are the 3 "transit market" fixed effects for each route. The remaining rows interact survey round fixed effects with the following variables. Rows 2 and 3 add fixed effects for each day of the week and hour of departure. Row 4 adds a fixed effect for each origin terminal. Row 5 adds dummies for each quartile of trip distance. Row 6 adds a dummy for whether the route ever appears in a LAMATA deployment plan, reflecting whether they planned to open public transit on that route at some point. Row 7 adds dummies for whether the route ends on Lagos Island or Victoria Island, containing the city's two main central business districts. Row 8 adds a third-order polynomial in origin and destination latitude and longitude. Row 9 adds fixed effects for the LGA of each origin and destination. Standard errors clustered by route and terminal reported in parentheses. * $p<0.1$; ** $p<0.05$; *** $p<0.01$.

Table S20. Effect of Public Transit on Minibus Drivers: Sun–Abraham Robustness

	N Trips	log(AvgFare)	log(Trip Fee)	log(Profit)	Days Worked	Change Route	Change Route at Terminal
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Open	-1.544***	-0.046	-0.037	0.004	1.539***	0.070	0.114
	(0.501)	(0.031)	(0.069)	(0.051)	(0.129)	(0.086)	(0.084)
N	2,161	1,991	2,015	2,068	2,635	1,383	1,383
R^2	0.76	0.87	0.78	0.59	0.58	0.83	0.80
Mean Outcome (Levels)	9.58	197.26	580.88	5004.84	4.41	0.59	0.51
Driver FE	X	X	X	X	X	X	X
MP × Survey Round FE	X	X	X	X	X	X	X
Age Qrtl × Survey Round FE	X	X	X	X	X	X	X
Num Trips at Baseline × Survey Round FE	X	X	X	X	X	X	X
Own Vehicle × Survey Round FE	X	X	X	X	X	X	X

Notes: Estimates use the interaction-weighted Sun and Abraham (2021) estimator, with cohort defined by calendar month of route opening and never-treated routes as the control cohort. All columns include driver fixed effects, and survey round fixed effects interacted with terminal of recruitment, dummies for each quartile of driver age, the number of trips the driver reports in the previous workday at baseline, and a dummy for whether they own their vehicle or not at baseline. Number of trips are winsorized at 99th percentile due to large outliers. Regressions are weighted by the sampling weights discussed in Section 3.1. Standard errors clustered by driver and terminal of recruitment reported in parentheses. * p<0.1; ** p<0.05; *** p<0.01.

Table S21. Effect of Public Transit on Connected Private Transit Routes: Sun–Abraham Robustness

	log(Departures)		log(Fare)		log(Queues)	
	(1)	(2)	(3)	(4)	(5)	(6)
Open	-0.210***	-0.159*	-0.001	-0.079**	-0.230***	-0.138*
	(0.078)	(0.083)	(0.026)	(0.037)	(0.078)	(0.083)
Num routes open at terminal		0.002		-0.016***		0.029*
		(0.013)		(0.006)		(0.017)
R^2	0.66	0.63	0.95	0.94	0.59	0.55
N	21,284	21,284	23,067	23,067	22,290	22,290
Terminal × Survey Round FE	X		X		X	
p-val: Same estimate on Open	0.65		0.05		0.47	

Notes: Estimates use the interaction-weighted Sun and Abraham (2021) estimator, with cohort defined by calendar month of route opening and never-treated routes as the control cohort. All columns include the same controls as the main specification (column (7) of Table 2), except for terminal-by-survey round fixed effects which are omitted in even columns. Number of open routes at a terminal is defined as the maximum of the number of public routes open at a minibus route's origin and destination, and is non-zero only for routes that never receive public service themselves. Final row shows the estimated p-value on a hypothesis test of equality between the coefficient on Open in the corresponding odd and even columns. Standard errors clustered by route and terminal reported in parentheses. * p<0.1; ** p<0.05; *** p<0.01.

Table S22. Effect of Public Transit on Connected Minibus Drivers: Sun–Abraham Robustness

	N Trips	log(AvgFare)	log(Trip Fee)	log(Profit)	Days Worked	Change Route	Change Route at Terminal
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Open	-1.325***	-0.078**	-0.015	-0.086*	1.845***	0.103	0.183***
	(0.502)	(0.032)	(0.058)	(0.048)	(0.146)	(0.067)	(0.068)
Num Routes Open at Terminal	-0.028	-0.006*	-0.002	-0.018***	0.254***	0.006	0.020**
	(0.083)	(0.004)	(0.008)	(0.006)	(0.033)	(0.008)	(0.008)
N	2,173	2,011	2,026	2,081	2,644	1,405	1,405
R^2	0.71	0.84	0.72	0.52	0.55	0.80	0.76
Mean Outcome (Levels)	9.58	197.26	580.88	5004.84	4.41	0.59	0.51
Driver FE	X	X	X	X	X	X	X
Age Qrtl × Survey Round FE	X	X	X	X	X	X	X
Num Trips at Baseline × Survey Round FE	X	X	X	X	X	X	X
Own Vehicle × Survey Round FE	X	X	X	X	X	X	X

Notes: Estimates use the interaction-weighted Sun and Abraham (2021) estimator, with cohort defined by calendar month of route opening and never-treated routes as the control cohort. All columns include driver fixed effects, and survey round fixed effects interacted with dummies for each quartile of driver age, the number of trips the driver reports in the previous workday at baseline, and a dummy for whether they own their vehicle or not at baseline. Open and Number of open routes are defined analogously to the terminal observation specifications, but based on a driver's main route from the baseline survey. Regressions are weighted by the sampling weights discussed in Section 3.1. Standard errors clustered by driver and terminal of recruitment reported in parentheses. * p<0.1; ** p<0.05; *** p<0.01.

Table S23. Comparison of Connectedness and Fixed Effect Specifications

	(1)	(2)	(3)	(4)
Panel A: log(Departures)				
Open	-0.163	-0.197	-0.232	-0.221
	(0.074)	(0.020)	(0.001)	(0.005)
	[0.134]	[0.070]	[0.038]	[0.060]
Number of open routes at terminal	0.002			
	(0.874)			
	[0.887]			
1-4 routes at terminal open		-0.101		
		(0.015)		
		[0.045]		
5+ routes at terminal open		-0.048		
		(0.493)		
		[0.574]		
Any routes at terminal open			-0.099	
			(0.021)	
			[0.054]	
N	21,287	21,287	21,287	21,287
R^2	0.63	0.63	0.63	0.66
p-val: Same estimate on Open	0.40	0.67	0.76	
Panel B: log(Fare)				
Open	-0.108	-0.082	-0.041	-0.024
	(0.042)	(0.105)	(0.321)	(0.501)
	[0.117]	[0.201]	[0.405]	[0.585]
Number of open routes at terminal	-0.016			
	(0.006)			
	[0.023]			
1-4 routes at terminal open		-0.016		
		(0.573)		
		[0.632]		
5+ routes at terminal open		-0.079		
		(0.013)		
		[0.043]		
Any routes at terminal open			-0.019	
			(0.489)	
			[0.560]	
N	23,070	23,070	23,070	23,070
R^2	0.94	0.93	0.93	0.95
p-val: Same estimate on Open	0.01	0.02	0.48	
Panel C: log(Queues)				
Open	-0.163	-0.181	-0.291	-0.289
	(0.084)	(0.022)	(0.000)	(0.000)
	[0.135]	[0.058]	[0.003]	[0.033]
Number of open routes at terminal	0.029			
	(0.094)			
	[0.141]			
1-4 routes at terminal open		0.003		
		(0.960)		
		[0.966]		
5+ routes at terminal open		0.176		
		(0.077)		
		[0.109]		
Any routes at terminal open			0.011	
			(0.867)	
			[0.883]	
N	22,293	22,293	22,293	22,293
R^2	0.55	0.55	0.55	0.59
p-val: Same estimate on Open	0.21	0.16	0.97	
Terminal FE				X

Notes: Column 1 repeats the baseline spillover specification from even columns in Table 4. Column 2 replaces the public transit connection measure with two dummies that bin the number of open routes at a terminal variable from column 1 into two groups above and below 5 open routes. Column 3 does the same with a single dummy for whether a route has any public transit connection (i.e. whether Number of open routes at terminal is greater than zero). Column 4 repeats the main specification with terminal by survey round fixed effects (from column 7 of Table 2). The last row of each panel reports p-value from a hypothesis test of equality between the coefficient on Open in each spillover specification with the specification with terminal by survey round fixed effects in column (4). P-values clustered by route and terminal are reported in parentheses. Wild p-values computed using Cameron, Gelbach, and Miller (2008) are reported in square brackets.

Table S24. Driver Robustness

	N Trips	log(AvgFare)	log(Trip Fee)	log(Profit)	Days Worked	Change Route Any	Change Route Same Terminal
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Unweighted Spillover Specification							
Open	-1.052	-0.072	-0.024	-0.091	1.872	0.075	0.152
	(0.065)	(0.049)	(0.711)	(0.033)	(0.000)	(0.087)	(0.004)
	[0.085]	[0.074]	[0.745]	[0.064]	[0.000]	[0.156]	[0.031]
Number of routes open at terminal	-0.053	-0.004	0.000	-0.023	0.256	0.006	0.021
	(0.474)	(0.273)	(0.987)	(0.001)	(0.000)	(0.393)	(0.012)
	[0.544]	[0.347]	[0.990]	[0.009]	[0.000]	[0.459]	[0.052]
N	2,173	2,011	2,026	2,081	2,644	1,405	1,405
R^2	0.71	0.84	0.73	0.52	0.55	0.79	0.75
Panel B: Spillover Dummy Specification							
Open	-1.183	-0.053	0.013	-0.081	2.097	0.077	0.128
	(0.048)	(0.111)	(0.834)	(0.078)	(0.000)	(0.108)	(0.025)
	[0.069]	[0.138]	[0.854]	[0.118]	[0.000]	[0.169]	[0.066]
Any routes at recruitment terminal open	-0.452	0.016	0.061	-0.081	1.703	0.029	0.052
	(0.156)	(0.469)	(0.273)	(0.109)	(0.000)	(0.225)	(0.105)
	[0.177]	[0.475]	[0.294]	[0.128]	[0.000]	[0.255]	[0.125]
N	2,177	2,015	2,030	2,085	2,648	1,411	1,411
R^2	0.71	0.84	0.73	0.52	0.57	0.80	0.75
Mean Outcome (Levels)	9.58	197.26	580.88	5004.84	4.41	0.59	0.51

Notes: Panel A replicates the main spillover Table 5 without sampling weights. Panel B replicates it using a dummy for whether the route has any public transit connections as the connection measure. P-values clustered by driver and terminal of recruitment reported in parentheses. Wild p-values computed using Cameron, Gelbach, and Miller (2008) are reported in square brackets.

Table S25. Wait Experiment: Alternate Estimation Strategies

	(1)	(2)	(3)
	Main	Perfect Compliance	Stated Preference
γ (utils/₦)	0.0025***	0.0031***	0.0002**
	(0.0002)	(0.0002)	(0.0001)
η (utils/min)	0.0464***	0.0254***	0.0109***
	(0.0031)	(0.0035)	(0.0013)
$\frac{\eta}{\gamma}$ (₦/min)	18.9357***	8.3295***	46.4624**
	(1.7123)	(0.983)	(20.4845)
$\sigma^{checkin}$	14.3662***		
	(3.901)		
ρ	0.5163***		
	(0.0588)		
N	8640	4722	4398
Log Likelihood	-8720.18	-2843.73	-2641.69

Notes: Left column reports the selection corrected estimates. The middle column reports estimates from equation (S26) assuming compliance is perfect ($C_{nt} \equiv 1$ and $F^{\nu^{checkin}}(-\underline{C}) = F(-\underline{D}, -\underline{C}) = 0$). The right column estimates the perfect compliance specification using stated preferences from the baseline survey. Standard errors clustered at the user level. Estimation fixes $\sigma^{wait} = 1$. * p<0.1; ** p<0.05; *** p<0.01.

Table S26. Wait Experiment: Alternate Utility Specifications

	(1)	(2)	(3)
	Main	Income Heterogeneity	Curvature in Wait Time
γ (utils/₦)	0.0025***	0.0025***	0.0022***
	(0.0002)	(0.0002)	(0.0002)
η (utils/min)	0.0464***		0.0723***
	(0.0031)		(0.0086)
η^H (utils/min)		0.0507***	
		(0.0028)	
η^L (utils/min)		0.0447***	
		(0.0038)	
$\eta^{SquaredWait}$			-0.0012***
			(0.0004)
$\frac{\eta}{\gamma}$ (₦/min)	18.9357***		
	(1.7123)		
$\frac{\eta^H}{\gamma}$ (₦/min)		20.6187***	
		(1.6831)	
$\frac{\eta^L}{\gamma}$ (₦/min)		18.1737***	
		(1.9662)	
$\sigma^{checkin}$	14.3662***	14.35***	12.3846***
	(3.901)	(3.9145)	(3.5069)
ρ	0.5163***	0.5168***	0.68***
	(0.0588)	(0.0591)	(0.0705)
N	8640	8640	8640
Log Likelihood	-8720.18	-8718.71	-8715.92

Notes: η^H represents the coefficient for above median income; η^L for below. Standard errors clustered at the user level. Estimation fixes $\sigma^{wait} = 1$. * p<0.1; ** p<0.05; *** p<0.01.

Table S27. Wait Experiment Robustness

	(1) Main	(2) Omit Nontravel Days†	(3) Exclude First Checkins	(4) Include Early User-Days
γ (utils/₦)	0.0025***	0.0024***	0.0026***	0.0025***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
η (utils/min)	0.0464***	0.0465***	0.0426***	0.0389***
	(0.0031)	(0.0031)	(0.0036)	(0.0022)
$\frac{\eta}{\gamma}$ (₦/min)	18.9357***	19.4297***	16.4055***	15.7436***
	(1.7123)	(1.7821)	(1.6209)	(1.224)
$\sigma^{checkin}$	14.3662***	10.3003***	36.0153	46.3072
	(3.901)	(2.1381)	(26.3611)	(34.0152)
ρ	0.5163***	0.5551***	0.3839***	0.5234***
	(0.0588)	(0.0608)	(0.0644)	(0.0419)
N	8640	7886	7880	17291
Log Likelihood	-8720.18	-8024.04	-7955.94	-16700.48

Notes: The second column (†) omits days in the first week that participants told us they did not plan to travel (based on the baseline survey), and days in the last week that participants told us they did not travel (based on the endline). These are likely to represent a subset of the days that participants did not travel. The third column excludes the first checkins which had an optimistic offer. The fourth column includes even users who faced an early version of the design, as described in Section S7.2. Standard errors clustered at the user level. Estimation fixes $\sigma^{wait} = 1$. * p<0.1; ** p<0.05; *** p<0.01.

Table S28. Welfare Effects of Introducing Public Transit: Robustness

	(1) Baseline	(2) σ PPML	(3) 25% Lower θ	(4) 25% Higher θ	(5) 25% Higher η
Panel A: Commuters, Individual ($/day)					
Treated Routes					
Effect of introducing public transit	0.23	0.23	0.32	0.18	0.22
Direct benefit of public transit	0.25	0.25	0.33	0.20	0.25
Additional impact from private Δdepartures	-0.04	-0.04	-0.04	-0.04	-0.05
Additional impact from private Δprices	0.03	0.03	0.03	0.03	0.03
Connected Routes					
Direct benefit of public transit	0	0	0	0	0
Impact from private Δprices	0.03	0.03	0.03	0.03	0.03
Panel B: Commuters, Aggregate ($mn/month)					
Effect of introducing public transit	1.93	1.93	2.40	1.64	1.86
Direct benefit of public transit	1.43	1.43	1.91	1.14	1.43
Additional impact from private Δdepartures	-0.25	-0.25	-0.25	-0.25	-0.31
Additional impact from private Δprices	0.75	0.75	0.74	0.75	0.74
Panel C: Minibus Drivers, Individual ($/day)					
Effect of introducing public transit	-4.40	-4.48	-4.40	-4.40	-4.40
Panel D: Minibus Drivers, Aggregate ($mn/month)					
Effect of introducing public transit	-1.27	-1.30	-1.27	-1.27	-1.27

Notes: Robustness exercises of the results in Table 7, under alternative assumptions for parameters in the column headings.

Appendix S3. Data Collection

S3.1. Minibus Driver Survey Sampling

Minibus drivers are difficult to sample. We found that terminal queues were a natural place to recruit drivers into our survey, but drivers travel frequently and some skip terminal queues. We attempt to obtain representative statistics on the population of minibus drivers using two sampling streams and a reweighting correction.[44]

We recruit most of our drivers (632) from terminal queues. For these drivers we apply a sampling correction. We seek to estimate the average value of a characteristic x in a population. However, each individual i may be in different states: although we sample individuals in a queue, the likelihood that i is in the queue is W_i. We ask our sampled drivers what proportion of their trips begin at terminals (w_i; that is one minus the proportion skipping the terminal start).[45] This yields an estimate of W_i: w_i. We apply this sampling probability in an inverse probability weight for our results using the driver survey. Note that if individuals are in the queue with the same probability ($w_i \equiv w$), this would not reweight individuals in this sample relative to one another. Alternately, note that this correction cannot help us for individuals who are *never* observed in the queue: if $w_i = 0$, we would never observe the individual in this sample.

To attempt to capture some of these drivers who queue less, we specified that a portion of our sample was to be recruited from drivers waiting outside terminals. These drivers (217 in number) are more likely to skip starting at the terminal: 12% report always skipping the terminal start. We had initially planned to apply the above correction in reverse for this sample, which would have assumed that a driver outside the terminal was planning on skipping the queue; however, 50% of these drivers report never skipping the queue in the past workday. So some of these drivers may be simply taking an extended break. For that reason, we do not apply a sampling correction to these drivers; we include them in the sample each with a weight of 1 (equivalent to a 100% probability of being sampled).

Appendix Table S4 compares the weighted statistics to their unweighted counterparts.

Appendix S4. Theory

S4.1. Complete Model Details

Model Overview. Given the queuing structure we document in these markets, the supply side consists of an M/M/1 queuing model with bulk service (Kendall 1953; Kleinrock 1975). Having chosen a route to work on, drivers queue, load passengers, complete their trip, and return to the origin, taking prices and the number of entrants as given. Daily profits depend on both the number of trips they expect to make and the profit per trip. The number of trips they can make in a workday depends on the time each trip takes plus the time spent queuing. The queuing time is shorter when demand is higher (as buses fill and depart faster) and when fewer buses are in the queue (since buses reach the front more

[44] This reweighting procedure was inspired by a conversation with Rob Jensen.

[45] We compute this proportion based on the trip diary for the previous day (including up to 8 trips) plus the trip they are about to complete on the day of survey (where we count them as beginning at the queue if they are sampled in the queue).

quickly). Prices in the market are set by an association acting as a monopolist seeking to maximize its profits (akin to a rideshare platform).

Government entry introduces a new variety of travel mode that commuters value. As commuters adopt the new public mode, the arrival rate of passengers into the private market falls, creating a shock to its equilibrium.

S4.1.1. Commuters

Setup. Time is discrete and the horizon is infinite. Agents do not discount the future, but expect to exit the model after some period of time. There are I locations, $i \in \{1, \ldots, I\}$. Each period, commuters arrive at i intending to travel to j according to a Poisson process with parameter μ_{ij}. We call the pair ij a route.

Mode Choice. Commuters, indexed by n, select a mode $m \in \mathcal{M}$ to travel the route. Prior to government entry, $\mathcal{M} = \{M, 0\}$ includes minibus (M) and an outside option (0) capturing other modes such as walking or automobile. After government entry, public buses (P) are added to treated routes ($m \in \mathcal{M}' = \{M, P, 0\}$). The individual utility provided by mode m is given by $u_{ijmn} = u_{ijm} + \epsilon_{ijmn}$ where ϵ_{ijmn} is an additive type 1 extreme value preference distributed with scale parameter $1/\theta$.

The fraction of commuters who choose minibuses is therefore

$$s_{ijm} = \frac{\exp\left(\theta u_{ijm}\right)}{\sum_{m' \in \mathcal{M}} \exp\left(\theta u_{ijm'}\right)} \tag{S1}$$

Overall consumer surplus on route ij is

$$CS_{ij} = \frac{1}{\gamma\theta} \ln\left(\sum_{m' \in \mathcal{M}} \exp\left(\theta u_{ijm'}\right)\right). \tag{S2}$$

We are agnostic about the utility provided by public buses u_{ijP}, and normalize the mean utility of the outside option $u_{ij0} = 0$. To account for the potential changing utility provided by minibuses following government entry, we next model u_{ijM} explicitly.

Value Function for Minibus Travelers. Conditional on having chosen to travel via minibus, a commuter arrives at a queue for route ij. Passengers arrive at rate $\mu_{ij} s_{ijM}$, where s_{ijM} is the fraction of commuters who choose to travel via minibus. In what follows, we drop the explicit dependence of minibus market-specific objects on M where it does not introduce confusion.

With probability β_{ij} there is a bus waiting at the top of the queue. The commuter boards the bus, and pays a per period wait cost η as they wait to depart. Buses depart once they are full to their $\bar{n}$ passenger capacity. With probability $1 - \beta_{ij}$ there is no bus waiting on the queue, so the commuter waits until one arrives and boards the bus. In steady state, it will turn out there are sufficient passengers for this bus to immediately depart, so we assume this from the outset and verify the conjecture below.

Buses arrive to the queue at Poisson rate λ_{ij}. The departure rate $\mu_{ij} s_{ijM}/\bar{n}$ depends on how fast passengers arrive and bus capacity, since each waits to fill before leaving. Once a bus departs, it arrives at its destination at Poisson rate δ_{ij}^A. Traveling passengers continue to pay the time cost η, but upon arrival at the destination receive an amenity α_M from the trip, and a payoff $y - p_{ijM}$ consisting of their earnings for the day net of the fare, which they value with parameter γ.

We show in Appendix Section S4.2.2 that the expected value function for a commuter who chooses to travel via minibus is

$$
\begin{aligned}
& u_{ijM} = \alpha_M + \gamma y - \gamma p_{ijM} - \eta t_{ijM} \\
\text{where} \quad & t_{ijM} = t_{ijM}^T + \underbrace{(1-\beta_{ij}) t_{ijM}^W + \beta_{ij} t_{ijM}^F}_{\bar{t}_{ijM}^W} \\
& t_{ijM}^W = 1/\lambda_{ij} \qquad \text{(S3)} \\
& t_{ijM}^F = \frac{1}{\mu_{ij} s_{ijM}} \left(\frac{\bar{n}-1}{2} \right). \qquad \text{(S4)}
\end{aligned}
$$

Since earnings y accrue under every mode, γy is absorbed by the normalization $u_{ij0} = 0$, yielding equation (3). Total travel time t_{ijM} consists of both expected travel time $t_{ijM}^T = 1/\delta_{ij}^A$ and expected wait time $\bar{t}_{ijM}^W$. With probability β_{ij} there is a bus in the queue, and wait time depends on the time it takes for the bus to fill t_{ijM}^F. Since a waiting bus holds $(\bar{n}-1)/2$ passengers in expectation, this is the time it takes for another $(\bar{n}-1)/2$ to arrive. With probability $1-\beta_{ij}$ there is no bus in the queue, and wait time depends on the arrival rate of buses to the queue $t_{ijM}^W = 1/\lambda_{ij}$.

S4.1.2. Supply

We begin by characterizing minibus behavior and profits conditional on having chosen a route, and then consider driver route choice.

Value Functions for Minibus Drivers. Each route follows a queuing model as follows. At terminal i, drivers may join a queue for route ij, and wait for buses ahead of them to fill up and depart on a first-in-first-out basis. A driver joining a queue with N_{ij}^Q buses in it expects to wait

$$
t_{ij}^Q = \frac{\bar{n}}{\mu_{ij} s_{ijM}} (N_{ij}^Q + 1) \qquad \text{(S5)}
$$

minutes in the queue. This is the time it takes for these N_{ij}^Q buses, along with the driver arriving in the queue, to fill and depart. Once the driver reaches the top of the queue and fills up, the bus departs and the driver receives net income per trip of $p_{ijM}\bar{n} - c_{ij} - \tau_{ij}$. This depends on the revenues from passenger fares, a lump sum cost of making the trip c_{ij}, and the association's per-trip fee τ_{ij}.

Buses arrive probabilistically at the destination, and at the end of each trip, drivers exit the model with a probability δ_{ij}^E proportional to the total length of a trip (micro-founded in Appendix Section S4.2.6). If they do not exit, they return to the end of the queue at i immediately. For tractability we abstract from the return leg of trips, and take the decision to leave when full as exogenous.[46]

Drivers, indexed by φ, are heterogeneous in the length of their workdays on each route $T_{ij\varphi} = T\nu_{ij\varphi}$, where T is the average workday length and $\nu_{ij\varphi}$ is an idiosyncratic component. Appendix Section S4.2.3

[46] In the data we find that 96% of minibuses leave when full.

shows that the value function for driver φ who enters route ij is $V^Q_{ij\varphi} = V^Q_{ij}\nu_{ij\varphi}$ where the common component is given by

$$V^Q_{ij} = \underbrace{\frac{T}{t^Q_{ij} + t^T_{ij}}}_{N^{\text{Trips}}_{ij}} \times \underbrace{\left[p_{ijM}\bar{n} - c_{ij} - \tau_{ij}\right]}_{\pi_{ij}}.$$

Intuitively, the value from entering a route depends on the number of trips a driver makes per day on that route and the profit per trip. Business stealing operates through queue lengths: more entrants on a route will, all else equal, increase time spent in queues through equation (S5) and reduce the number of trips the driver can complete in a day.

Route Choice. As explained in Section 3, drivers face high costs by the association to register to operate from a terminal. We therefore treat as separate the decisions of which terminal i to enter and which route ij to enter conditional on having chosen a terminal.

Driver φ who has chosen to enter terminal i solves the route choice problem $\max_j\{V^Q_{ij\varphi}\}$. We assume the idiosyncratic shocks $\nu_{ij\varphi}$ are drawn from a Fréchet distribution with unit mean and shape parameter σ. The fraction of drivers at terminal i choosing route ij is then[47]

$$\rho_{ij} = \frac{\left(V^Q_{ij}\right)^\sigma}{\sum_k \left(V^Q_{ik}\right)^\sigma} = \frac{\left(N^{\text{Trips}}_{ij} \times \left[p_{ijM}\bar{n} - c_{ij} - \tau_{ij}\right]\right)^\sigma}{\sum_k \left(N^{\text{Trips}}_{ik} \times \left[p_{ikM}\bar{n} - c_{ik} - \tau_{ik}\right]\right)^\sigma}. \tag{S6}$$

Each period B_i drivers enter terminal i, so that $B_{ij} = B_i\rho_{ij}$ join the queue for route ij. We discuss how total entry B_i is determined below.

Due to properties of the Fréchet distribution, average profits for entrants at terminal i are the same regardless of the route they choose. These are related to the denominator of the choice probabilities through

$$\Pi_i = \underbrace{\left[\sum_k \left(N^{\text{Trips}}_{ik} \times \left[p_{ikM}\bar{n} - c_{ik} - \tau_{ik}\right]\right)^\sigma\right]^{1/\sigma}}_{\Pi^V_i} - F_i \tag{S7}$$

where Π^V_i are average variable profits and F_i are registration costs charged by the minibus association to operate at terminal i (expressed as per-workday equivalents).

S4.1.3. Minibus Association

We will allow our model to accommodate free or fixed entry of drivers. We solve the association's problem under free entry, and then state how the results change under fixed entry.

The minibus association sets prices, per-trip fees, and registration fees to maximize its revenues at each terminal. Under free entry of minibus drivers, revenues equal registration fees $B_i\Pi^V_i$ plus per-trip

[47] We restrict attention to equilibria in which $V^Q_{ij} > 0$ for every active route.

fee income $\sum_j \tau_{ij}\bar{N}_{ij}B_{ij}$, where $\bar{N}_{ij} \equiv \frac{T}{t^Q_{ij}+t^T_{ij}}\left(\frac{B_{ij}}{B_i}\right)^{-1/\sigma}$ is the expected number of daily trips among drivers who choose route ij, reflecting selection on workday length. We formulate the problem in terms of choosing optimal prices p_{ijM}, per-trip fees τ_{ij}, and the numbers of entrants B_i and B_{ij}, where the fees can then be inverted to solve for optimal registration fees using the free entry condition.[48]

The association maximizes this objective subject to several constraints. First, it cannot directly control drivers' choice of which route to enter after entry into a terminal. Instead, it understands that as it shifts overall entry into the terminal, fees, and prices on each route, it affects the number of buses on each route B_{ij} via drivers' "route choice" constraint[49]

$$\frac{B_{ij}}{B_i} = \left(\frac{T}{t^Q_{ij}+t^T_{ij}}\frac{p_{ijM}\bar{n} - c_{ij} - \tau_{ij}}{\Pi^V_i}\right)^{\sigma}. \tag{S8}$$

Second, the association understands that by changing prices and entry it affects queue times faced by drivers

$$t^Q_{ij} = \frac{\bar{n}}{\mu_{ij}s_{ijM}(p_{ij},B_{ij})}(1+N^Q_{ij}(p_{ij},B_{ij})). \tag{S9}$$

Third, it faces an adding up constraint

$$B_i = \sum_j B_{ij}. \tag{S10}$$

The association's problem is therefore represented by the lagrangian

$$\begin{aligned}
\mathcal{L} = B_i\Pi^V_i &+ \sum_j \tau_{ij}\frac{T}{t^Q_{ij}+t^T_{ij}}\left(\frac{B_{ij}}{B_i}\right)^{-1/\sigma}B_{ij} \\
&+ \sum_j \kappa_{ij}\left[\frac{T}{t^Q_{ij}+t^T_{ij}}\left[p_{ijM}\bar{n} - c_{ij} - \tau_{ij}\right]\left(\frac{B_{ij}}{B_i}\right)^{-1/\sigma} - \Pi^V_i\right] \\
&+ \sum_j \zeta_{ij}\left[t^Q_{ij} - \frac{\bar{n}}{\mu_{ij}s_{ijM}(p_{ij},B_{ij})}(1+N^Q_{ij}(p_{ij},B_{ij}))\right] \\
&+ \mu_i\left[\sum_j B_{ij} - B_i\right]
\end{aligned}$$

where $\kappa_{ij}, \zeta_{ij}, \mu_i$ are the multipliers on the driver route choice constraint, the queue time constraint, and the adding up constraint respectively. The necessary first order conditions at an optimum are

$$\kappa_{ij}\frac{T}{t^Q_{ij}+t^T_{ij}}\bar{n}\left(\frac{B_{ij}}{B_i}\right)^{-1/\sigma} - \zeta_{ij}t^Q_{ij}\frac{\partial \ln t^Q_{ij}}{\partial p_{ij}} = 0 \tag{p_{ijM}}$$

[48] In a quantitative implementation we could follow the literature on ridesharing platforms (e.g. Rosaia (2024)) and assume the association maximizes a combination of profits and consumer surplus. This would capture that the government may pressurize the association towards benefiting commuters in exchange for the right to oversee transit provision and for the land given to it for terminals. However since we do not use the association side of the model in the sufficient statistics approach, we omit this for brevity.

[49] We note that choice probabilities s_{ijM} and queue lengths N^Q_{ij} depend on entry through its effects on the arrival rate of buses into the queue λ_{ij} and the probability of finding a bus on the queue β_{ij}, so denote this dependence implicitly.

$$\frac{T}{t_{ij}^Q + t_{ij}^T}\left(\frac{B_{ij}}{B_i}\right)^{-1/\sigma}[B_{ij} - \kappa_{ij}] = 0 \tag{τ_{ij}}$$

$$\Pi_i^V + \frac{1}{\sigma}\frac{1}{B_i}\left[\sum_j \tau_{ij}\bar{N}_{ij}B_{ij} + \Pi_i^V \sum_j \kappa_{ij}\right] - \mu_i = 0 \tag{B_i}$$

$$\left(1 - \frac{1}{\sigma}\right)\tau_{ij}\bar{N}_{ij} - \kappa_{ij}\frac{1}{\sigma}\frac{\Pi_i^V}{B_{ij}} - \zeta_{ij}\frac{\partial t_{ij}^Q}{\partial B_{ij}} + \mu_i = 0 \tag{B_{ij}}$$

$$B_i = \sum_j \kappa_{ij} \tag{Π_i^V}$$

$$\zeta_{ij} - [\tau_{ij}B_{ij} + \kappa_{ij}(p_{ijM}\bar{n} - c_{ij} - \tau_{ij})]\frac{\bar{N}_{ij}}{t_{ij}^Q + t_{ij}^T} = 0 \tag{t_{ij}^Q}$$

Optimal Prices. The (τ_{ij}) condition implies $\kappa_{ij} = B_{ij}$. Combining first order conditions (p_{ijM}) and (t_{ij}^Q) yields

$$\bar{n} - [p_{ijM}\bar{n} - c_{ij}]\frac{t_{ij}^Q}{t_{ij}^Q + t_{ij}^T}\frac{\partial \ln t_{ij}^Q}{\partial p_{ij}} = 0$$

When the association raises prices on route ij, it increases fare income by $\bar{n}$ per trip. But it also lowers the number of trips drivers make per day by increasing queue times as passenger demand falls. The optimal price equates these costs and benefits, leading to a markup

$$p_{ijM} - c_{ij}/\bar{n} = \frac{1}{\frac{t_{ij}^Q}{t_{ij}^Q + t_{ij}^T}\frac{\partial \ln t_{ij}^Q}{\partial p_{ij}}} \tag{S11}$$

Prices will therefore be lower whenever queue times are more sensitive to prices, since the association loses more revenues as a result of higher prices.

Optimal Entry. To characterize optimal fees, it is helpful to first characterize the allocation the association wants to implement. We first note that the route choice constraint can be rearranged to $\Pi_i^V = \bar{N}_{ij}(p_{ijM}\bar{n} - c_{ij} - \tau_{ij}) = V_{ij}^0 - \tau_{ij}\bar{N}_{ij}$, where $V_{ij}^0 \equiv \bar{N}_{ij}(p_{ijM}\bar{n} - c_{ij})$ is the gross surplus per driver on route ij (i.e. before fees). Multiplying by B_{ij}, summing over j, and imposing free entry $F_i = \Pi_i^V$ yields

$$B_i F_i + \sum_j \tau_{ij}\bar{N}_{ij}B_{ij} = \sum_j B_{ij}V_{ij}^0. \tag{S12}$$

In words, association revenue equals the entire gross surplus its drivers generate. The fees therefore do not determine how much the association captures; they determine which allocation is implemented. We can therefore consider the association choosing prices and the allocation of buses to maximize total gross surplus, and then find the two fee instruments that implement this allocation.

To compute optimal entry, we first define $\bar{V}_i^0 \equiv \sum_j \rho_{ij}V_{ij}^0$ as the average gross surplus across the terminal's routes, and $\varepsilon_{ij}^{N,B} \equiv \frac{\partial \ln N_{ij}^{\text{Trips}}}{\partial \ln B_{ij}} = -\frac{t_{ij}^{Q}}{t_{ij}^{Q}+t_{ij}^{T}} \frac{\partial \ln t_{ij}^{Q}}{\partial \ln B_{ij}} < 0$, the elasticity of daily trips with respect to entry on the route. Combining the (B_{ij}), (B_i), and (t_{ij}^{Q}) conditions with $\kappa_{ij} = B_{ij}$ and the route choice constraint, we begin by showing that, holding route shares constant, the total number of entrants satisfies

$$\sum_j \rho_{ij} V_{ij}^{0} \left(1 + \varepsilon_{ij}^{N,B}\right) = 0.$$

That is, the association admits entrants until the marginal entrant's gross value equals the business he steals from incumbents.

Taking B_i as given, entry on each route B_{ij} satisfies

$$V_{ij}^{0} \left(1 + \varepsilon_{ij}^{N,B} - \frac{1}{\sigma}\right) = -\frac{1}{\sigma} \bar{V}_i^{0},$$

i.e. the net marginal value of a bus on each route (gross value net of business stealing and lower productivity through worse selection) is equalized across routes. Combining the two yields the condition for optimal entry on each route

$$\varepsilon_{ij}^{N,B} = -1 + \frac{1}{\sigma}\left(1 - \frac{\bar{V}_i^{0}}{V_{ij}^{0}}\right). \tag{S13}$$

In the absence of driver heterogeneity ($\sigma \to \infty$), entry on each route is set at the point where the trips completed by the marginal bus equal the trips it steals from incumbents (i.e. $\varepsilon_{ij}^{N,B} = -1$) since this maximizes total route surplus (at the optimal fare). With heterogeneous workday lengths, expanding a high-value route ($V_{ij}^{0} > \bar{V}_i^{0}$) draws in relatively unproductive drivers, so the association chooses fewer entrants than it would with homogeneous drivers.

Optimal Fees. The association fees that implement this allocation induce drivers to internalize the business stealing effects they have on other drivers by affecting queue times. Combining the (B_{ij}) condition with $\kappa_{ij} = B_{ij}$ and the (t_{ij}^{Q}) condition, substituting the multiplier on the adding-up constraint from the (B_i) condition ($\mu_i = F_i + \frac{1}{\sigma}\bar{V}_i^{0}$ under free entry), and using the route choice constraint yields

$$\tau_{ij} = \frac{1}{\bar{N}_{ij}} \left[\underbrace{V_{ij}^{0} \left|\varepsilon_{ij}^{N,B}\right|}_{\text{business stealing}} + \underbrace{\frac{1}{\sigma}\left(V_{ij}^{0} - \bar{V}_i^{0}\right)}_{\text{relative selection}} - F_i \right] \tag{S14}$$

The trip fee charges the marginal bus for the externality it imposes on the route's incumbents through longer queue times and hence lower daily trips, plus a term that accounts for shifting driver productivities across routes. The registration fee captures the surplus that remains, $F_i = \Pi_i^{V}$. The $-F_i$ credit in (S14) prevents the association from charging for the same surplus through both instruments; together the two fees collect the route's entire gross surplus, $\tau_{ij}\bar{N}_{ij} + F_i = V_{ij}^{0}$.

Association Behavior Under Fixed Entry. When the overall number of entrants is fixed, the association takes B_i as given and chooses F_i, τ_{ij}, p_{ijM} to maximize $B_i F_i + \sum_j \tau_{ij} \bar{N}_{ij} B_{ij}$ subject to the same set of constraints as above, plus one additional constraint that $\Pi_i^V \geq F_i$, i.e. that drivers earn non-negative profits. Solving this problem yields exactly the same price setting condition (S11). The allocation of the fixed number of buses still equalizes $V_{ij}^0 \left(1 + \varepsilon_{ij}^{N,B} - \frac{1}{\sigma}\right)$ across routes. The association again captures driver surplus since the non-negative profit constraint binds.[50]

S4.1.4. Steady State Equilibrium

Queue Length. Letting B_{ij}^T denote the number of traveling buses on route ij in steady state, the arrival rate of buses into the queue is

$$\lambda_{ij} = \chi_{ij} \times \left[(1 - \delta_{ij}^E)\delta_{ij}^A B_{ij}^T + B_{ij}\right]. \tag{S15}$$

Each period, a fraction δ_{ij}^A of the traveling buses arrive at their destination, and $1 - \delta_{ij}^E$ of these re-join the queue, while B_{ij} entering buses also join. $\chi_{ij} \in (0, 1]$ is a parameter that reflects a possible delay in buses arriving in the queue.[51] The exit rate from the queue is $\mu_{ij} s_{ijM}/\bar{n}$, i.e. the arrival rate of $\bar{n}$ passengers.[52] We show in Appendix Section S4.2.5 that in steady state, the average number of buses in the queue is given by

$$N_{ij}^Q = \left[\frac{\mu_{ij} s_{ijM}}{\lambda_{ij} \bar{n}} - 1\right]^{-1}. \tag{S16}$$

This shows how greater entry from minibus drivers will tend to increase queue lengths (and queue times) through λ_{ij}. While greater demand through $\mu_{ij} s_{ijM}$ will tend to reduce queue lengths by increasing the departure rate from the queue, it will also have an effect on arrival rates λ_{ij} through the number of entrants.

Entry. Our assumption for long-run equilibrium in the market is that the number of entrants B_i is determined by free entry. Under free entry, B_i adjusts to affect queue times until

$$\Pi_i^V - F_i = 0 \; \forall i \tag{S17}$$

We assume this determines the initial equilibrium prior to public entry, and impose this in the equilibrium definition below.

[50] This holds when the association resets F_i and τ_{ij}. Our welfare calculations hold both at pre-entry levels, matching the absence of fee changes in our data, so changes in variable profits accrue to drivers. Fares still satisfy (S11), which maximizes joint driver and association surplus for any given B_i regardless of how it is divided.

[51] Equation (S16) requires $\lambda_{ij}\bar{n} < \mu_{ij} s_{ijM}$. In equilibrium entry keeps queues finite: the value of joining a route falls as its queue lengthens, so drivers enter and choose routes only while this holds, and we restrict attention to such equilibria.

[52] Although passengers arrive at Poisson rates, bus departures are not memoryless: a bus leaves when its $\bar{n}$th passenger boards, so fill times are Erlang. We approximate departures as Poisson for tractability, and the approximation fits the data better: at the observed $\beta_{ij} = 0.84$ it predicts a queue of 5.3 buses and the Erlang form 3.2, against 4.75 in the data. In practice passengers and buses arrive in bunches, so fill times are more irregular than Erlang, and the Poisson form captures this. In any case, no result uses the level of the queue.

The model can also accommodate a fixed number of entrants, and we impose this assumption to evaluate the changes in surplus from the policy given we do not observe driver exit (Table 3). We view this as a medium-run response with exit potentially following in the long-run (in which case we would impose the free entry condition in the post-entry equilibrium once more).

Traveling Buses. In steady state, the number of traveling buses leaving the state $B_{ij}^T \delta_{ij}^A$ must equal the inflows $\mu_{ij} s_{ijM}/\bar{n}$, so that[53]

$$B_{ij}^T = \frac{\mu_{ij} s_{ijM}}{\bar{n}\delta_{ij}^A}. \tag{S18}$$

Equilibrium Definition. *Given model parameters $\{\gamma, \eta, \alpha_m, \theta; \bar{n}, c_{ij}, \chi_{ij}, \sigma, T\}$ and data $\{\mu_{ij}, t_{ij}^T\}$, an equilibrium is a vector $\{s_{ijm}, CS_{ij}, \rho_{ij}, \Pi_i, t_{ij}^W, t_{ij}^F, t_{ij}^Q, p_{ij}, \tau_{ij}, F_i, N_{ij}^Q, \lambda_{ij}, B_i, B_{ij}, B_{ij}^T\}$ such that i) commuter surplus and mode choices satisfy (S1), (S2); ii) driver surplus and route choices satisfy (S6), (S7); iii) equilibrium waiting and queue times satisfy (S3), (S4), and (S5); iv) prices and fees satisfy association optimality* (S11) *and* (S14)*, with route entry $B_{ij} = B_i \rho_{ij}$, and v) equilibrium queue lengths, arrival rates of buses into the queue, the entry rate and number of traveling buses satisfy steady state conditions (S15), (S16),* (S17) *and (S18).*

S4.2. Additional Results and Derivations

S4.2.1. Comparative Statics for Fares

Since the optimal markup depends on $\frac{\partial \ln t_{ij}^Q}{\partial p_{ij}}$, we differentiate steady state queue times with respect to fares. Substituting the steady-state expressions for queue length (S16), bus arrivals (S15), and traveling buses (S18) into the queue-time equation (S5), and then differentiating with respect to the fare while holding B_{ij} fixed using $\frac{\partial \ln s_{ijM}}{\partial p_{ij}} = -\theta\gamma(1 - s_{ijM})$, yields

$$\frac{\partial \ln t_{ij}^Q}{\partial p_{ij}} = \underbrace{\theta\gamma(1 - s_{ijM})}_{\text{fill speed}} + \underbrace{\chi_{ij} B_{ij} t_{ij}^Q \, \theta\gamma(1 - s_{ijM})}_{\text{queue length}} = \left(1 + \chi_{ij} B_{ij} t_{ij}^Q\right) \theta\gamma(1 - s_{ijM}),$$

so that the optimal markup is $p_{ijM} - c_{ij}/\bar{n} = \left[\frac{t_{ij}^Q}{t_{ij}^Q + t_{ij}^T}\left(1 + \chi_{ij} B_{ij} t_{ij}^Q\right)\theta\gamma(1 - s_{ijM})\right]^{-1}$. Fares fall whenever public entry raises the price semi-elasticity of demand $\theta\gamma(1 - s_{ijM})$ or the number of buses serving the route B_{ij}.

On treated routes public transit competes for the same commuters. The direct effect is that minibus demand becomes more price elastic ($\theta\gamma(1 - s_{ijM})$ rises) as commuters have more travel options. An indirect effect comes through queue lengths, but its direction is ambiguous: lower passenger arrivals lengthen queues while drivers exiting the route shorten them. If the direct effect dominates, we expect minibus fares to fall.

[53]This last term is the probability a bus in the queue has $\bar{n} - 1$ passengers $1/\bar{n}$ times the probability an additional passenger shows up $\mu_{ij} s_{ijM}$

On connected routes, we find that demand for minibuses is unchanged while the supply of minibuses expands. The only impact on fares therefore comes through queue times, which rise (at existing fares) as drivers reallocate into these routes while the fill rate of buses at the top of the queue is unchanged. A longer queue is idle capacity: each queued bus completes a trip only once passengers arrive to fill it. Since the association captures the full gross surplus on trips, it tries to convert this idle capacity into trips by lowering fares. Seen another way, part of the cost of a high fare is a longer queue faced by each driver. More drivers on connected routes increase that cost, leading the association to lower fares.

S4.2.2. Derivation of Commuter Value Functions

We derive the value function on a single route, and so omit dependency of variables on route ij. We begin by solving for utility of a traveler who has chosen to travel via minibus. When this traveler arrives at the origin there is either a bus waiting or not. If there is a waiting bus, the passenger boards and waits for it to depart. If not, the passenger waits for one to arrive, enters it, and waits for it to depart.

Value Function for Waiting Passengers. With probability β there is a bus waiting. As shown in Appendix Section S4.2.4, the expected number of passengers waiting in the bus is $(\bar{n}-1)/2$. We want to solve for the value of joining a bus with this number of passengers already on board.

Let $U^B(n)$ be the value of arriving at a bus with $n < \bar{n}$ passengers in it. A commuter waiting at a bus stop incurs a per-period wait cost of $-\eta$. Since the bus leaves when it has $\bar{n}$ passengers, if there are $\bar{n}-2$ passengers on the bus she boards, the commuter waits until one additional passenger arrives. The value function is therefore

$$\begin{aligned} U^B(\bar{n}-2) &= -\eta + \mu s_M U^T + (1-\mu s_M) U^B(\bar{n}-2) \\ &= -\eta \frac{1}{\mu s_M} + U^T, \end{aligned}$$

where U^T is the value of being on a traveling bus, and $1/\mu s_M$ is the expected time for an additional passenger to join the waiting bus.

The value of arriving at a bus with $\bar{n}-3$ passengers on it is therefore

$$\begin{aligned} U^B(\bar{n}-3) &= -\eta \frac{1}{\mu s_M} + U^B(\bar{n}-2) \\ &= -\eta \frac{2}{\mu s_M} + U^T \end{aligned}$$

Continuing, the expected value of arriving when there is a waiting bus (which has $(\bar{n}-1)/2$ passengers on it) is

$$\begin{aligned} U^B &= -\eta t^F + U^T \\ \text{where } t^F &= \frac{1}{\mu s_M}\left(\frac{\bar{n}-1}{2}\right) \end{aligned}$$

is the expected time it takes for the bus to fill $(\bar{n}-1)/2$ passengers under passenger arrival rate μs_M.

With probability $1-\beta$ there is no bus waiting when they arrive. One arrives at Poisson rate λ. The value of arriving when there is no bus waiting is therefore

$$\begin{aligned} U^{NB} &= -\eta + \lambda \tilde{U}^B + (1-\lambda) U^{NB} \\ &= -\eta t^W + \tilde{U}^B \end{aligned}$$

where $t^W \equiv 1/\lambda$ is the expected time to wait for a bus to arrive, and $\tilde{U}^B$ is the value of the passenger being able to load the bus after waiting t^W periods, at which point there are $\mu_m t^W$ passengers waiting. We know that if $\mu_m t^W > \bar{n}$, then the bus departs immediately. We will show later this is satisfied in steady state, so $\tilde{U}^B = U^T$. Therefore

$$U^{NB} = -\eta t^W + U^T.$$

Value Function for Traveling Passengers. A traveling bus arrives according to a Poisson process with parameter δ^A. While traveling, passengers pay the time cost $-\eta$. If they arrive, they receive an amenity α_M from the trip, and a payoff $y-p$ consisting of their earnings for the day net of the price of the fare, which they value with parameter γ. Their value function is therefore

$$\begin{aligned} U^T &= -\eta + \delta^A(\alpha_M + \gamma(y-p)) + (1-\delta^A) U^T \\ &= -\eta t^T + \alpha_M + \gamma(y-p) \end{aligned}$$

where $t^T = 1/\delta^A$ is expected travel time on the route.

Expected Value Function. Using these results to compute the expected value from traveling by minibus $U_m = (1-\beta)U^{NB} + \beta U^B$, and adding back the dependency of variables on route, we find that

$$\begin{aligned} U_{ijM} &= \alpha_M + \gamma y - \gamma p_{ijM} - \eta t_{ijM} \\ \text{where } \; t_{ijM} &= t^T_{ijM} + \underbrace{(1-\beta_{ij}) t^W_{ijM} + \beta_{ij} t^F_{ijM}}_{\bar{t}^W_{ijM}} \end{aligned}$$

where $t^T_{ijM} = 1/\delta^A_{ij}$ is expected travel time, and $\bar{t}^W_{ijM}$ is the expected time the passenger waits at the bus queue until they depart. This consists of $t^W_{ijM} = 1/\lambda_{ij}$, the expected wait time when no bus is on queue when the passenger arrives, and $t^F_{ijM} = \frac{1}{\mu_{ij} s_{ijM}} \left(\frac{\bar{n}-1}{2}\right)$, the expected time for the bus to fill and depart if there is one on the queue when the passenger arrives. Note that the endogenous probability $\beta_{ij} = \lambda_{ij}\bar{n}/\mu_{ij} s_{ijM}$ is derived in Appendix Section S4.2.5.

S4.2.3. Derivation of Minibus Driver Value Functions

We derive the value function on a single route, and so omit dependency of variables on route ij. For ease of notation, we abstract from driver heterogeneity in the derivation and introduce it at the end of the section.

Value Functions, Traveling Buses. When traveling, a bus arrives at its destination with probability δ^A. If it arrives, it exits the model with probability δ^E. With probability $1-\delta^E$ it returns to the start of the queue at the route origin. We also allow for drivers to have a time cost c when traveling or queuing. The value function of a traveling minibus is therefore

$$\begin{aligned} V^T &= -c + \delta^A(1-\delta^E)V^Q + (1-\delta^A)V^T \\ &= -ct^T + (1-\delta^E)V^Q. \end{aligned}$$

Value Functions, Queuing Buses. Consider a bus joining a queue with N buses. Departures from the queue occur at rate $\mu s_M/\bar{n}$ since each bus waits to fill to capacity $\bar{n}$ before departing. Using the recursion

$$\begin{aligned} V^Q(N) &= -c + \frac{\mu s_M}{\bar{n}} V^Q(N-1) + (1 - \frac{\mu s_M}{\bar{n}})V^Q(N) \\ &= -\frac{c\bar{n}}{\mu s_M} + V^Q(N-1) \end{aligned}$$

allows us to solve for the value of joining a queue with N buses as a function of being at the front of the queue $V^Q(0)$ as

$$V^Q(N) = -c\frac{N\bar{n}}{\mu s_M} + V^Q(0) \tag{S19}$$

This is intuitive, as $\frac{N\bar{n}}{\mu s_M}$ is the expected time it takes for the queue of length N to clear.

Once at the top of the queue, the bus starts from zero waiting passengers. Let $V^Q(0,n)$ be the value of being at the top of the queue with n waiting passengers in the vehicle. Then $V^Q(0) = V^Q(0,0)$ which satisfies the recursion

$$\begin{aligned} V^Q(0,0) &= -c + \mu s_M V^Q(0,1) + (1-\mu s_M)V^Q(0,0) \\ &= -\frac{c}{\mu s_M} + V^Q(0,1) \\ V^Q(0,1) &= -\frac{c}{\mu s_M} + V^Q(0,2) \\ &\vdots \\ V^Q(0,\bar{n}-1) &= -\frac{c}{\mu s_M} + p\bar{n} - c_M + V^T \end{aligned}$$

The last line says that once the bus has one vacant seat left, it pays the cost to wait the expected amount of time $1/\mu s_M$ for one more passenger to arrive, gets the fares from passengers $p\bar{n}$, pays the lump sum monetary cost of travel c_M, and gains the value from becoming a traveling bus.

Solving this recursion yields $V^Q(0) = p\bar{n} - c_M - \frac{c\bar{n}}{\mu s_M} + V^T$. Substituting this, and the expression for V^T into (S19) allows us to express the value of joining a queue on ij as

$$V^Q_{ij} = \frac{1}{\delta^E_{ij}} \times \left[p_{ij}\bar{n} - c_{ij} - c(t^Q_{ij} + t^T_{ij})\right] \quad \text{where} \quad t^Q_{ij} = \frac{\bar{n}}{\mu_{ij}s_{ijM}}(N^Q_{ij}+1)$$

Note that queue time t_{ij}^Q is an endogenous object which depends on the length of the queue in steady state N_{ij}^Q, and we replace c_M with c_{ij} to match the notation in the text.

Note that $1/\delta_{ij}^E$ is the expected number of trips the driver expects to make in a day. Appendix Section S4.2.6 shows that if we assume a finite workday length without driver exit, we get exactly the expression above for the value of entry into a route, with $T/(t^Q + t^T)$ instead of $1/\delta_{ij}^E$ premultiplying the per-trip profit function. This is because in a workday of length T, since each trip takes $t^Q + t^T$, drivers make exactly $T/(t^Q + t^T)$ trips in a day. We therefore parameterize $\delta_{ij}^E = (t_{ij}^Q + t_{ij}^T)/T$ to deliver this same expression. The overall value of joining a queue is then

$$V_{ij}^Q = \underbrace{\frac{T}{t_{ij}^Q + t_{ij}^T}}_{N_{ij}^{\text{Trips}}} \times \underbrace{\left[p_{ij}\bar{n} - c_{ij} - c\left(t_{ij}^Q + t_{ij}^T\right)\right]}_{\pi_{ij}}$$

which depends on the number of trips drivers expect to make in a day, times the profits per trip.

Since the time cost of queuing is already captured through its effect on the number of trips, we assume $c = 0$ to deliver the final expression for the value function used in the text.

In the paper we assume drivers are heterogeneous in their workday length on each route. Driver φ has a workshift of length $T_{ij\varphi} = T\nu_{ij\varphi}$ on route ij. The number of trips he can complete in a day is therefore $N_{ij\varphi}^{\text{Trips}} = T_{ij\varphi}/(t_{ij}^Q + t_{ij}^T) = \nu_{ij\varphi} N_{ij}^{\text{Trips}}$, and the value of entering a route and joining a queue is $V_{ij\varphi}^Q = N_{ij\varphi}^{\text{Trips}} \pi_{ij}$.

S4.2.4. Distribution of Passengers Waiting on a Bus

The probability distribution p_n of the number of people n waiting evolves according to

$$p_{n,t+dt}^{pass} = p_{n,t}^{pass}(1 - \mu_{ij} s_{ijM} dt) + p_{n-1}^{pass} \mu_{ij} s_{ijM} dt.$$

In steady state this becomes $p_n^{pass} = p_{n-1}^{pass}$,which implies a uniform distribution, so

$$p_n^{pass} = \begin{cases} \frac{1}{\bar{n}} & \text{if } 0 \leq n \leq \bar{n} - 1 \\ 0 & \text{otherwise.} \end{cases}$$

Therefore the expected number of passengers waiting on a bus if there is one in the queue is $(\bar{n} - 1)/2$.

S4.2.5. Derivation of Steady State Queue Length

This section characterizes equilibrium queue lengths. We again suppress dependence on route ij. The arrival rate of buses into the queue is

$$\lambda = \chi \times \left[(1 - \delta^E)\delta^A B^T + B\right]$$

where B^T, B are the number of traveling and entering buses in steady state and $\chi \in (0, 1]$ is a parameter that reflects a potential delay in buses arriving into the queue. The exit rate from the queue is $\mu s_M / \bar{n}$.

The Kolmogorov Forward Equation describing the distribution β_n of the number of buses n is then

$$\beta_{0,t+dt} = \beta_{0,t}\left(1-\lambda dt\right) + \beta_{1,t}\frac{\mu s_M}{\bar{n}}dt$$
$$\beta_{n,t+dt} = \beta_{n+1,t}\frac{\mu s_M}{\bar{n}}dt + \beta_{n-1,t}\lambda dt + \beta_{n,t}\left(1-\lambda dt - \frac{\mu s_M}{\bar{n}}dt\right) \quad \forall\, n>0$$

Letting $dt \to 0$ we get

$$\beta_1 = \frac{\lambda \bar{n}}{\mu s_M}\beta_0$$
$$0 = \left(\beta_{n+1}-\beta_n\right)\frac{\mu s_M}{\bar{n}} + \left(\beta_{n-1}-\beta_n\right)\lambda$$

Solving this recursively, we find

$$\beta_n = \xi^n \beta_0 \quad \text{where } \xi \equiv \frac{\lambda \bar{n}}{\mu s_M}$$

Since this is a probability distribution, we find β_0 from the normalization

$$\beta_0 \sum_{n=0}^{\infty} \xi^n = 1 \quad \Rightarrow \quad \beta_0 = 1-\xi$$

Note β_0 also provides the probability of commuters finding zero buses on the queue needed to compute expected wait times for passengers, where we define in the paper the probability of finding a bus in the queue as $\beta_{ij} \equiv 1-\beta_{0,ij}$.

The average number of buses in the queue ij is then

$$N_{ij}^Q = \sum_{n=0}^{\infty} n(1-\xi_{ij})\xi_{ij}^n = \frac{\xi_{ij}}{1-\xi_{ij}} = \left[\frac{\mu_{ij}s_{ijM}}{\lambda_{ij}\bar{n}} - 1\right]^{-1}$$

where we have used $\sum_{n=0}^{\infty} n\xi^n = \xi\frac{\partial}{\partial \xi}\sum_{n=0}^{\infty}\xi^n$ in the second equality. Note these results require $\lambda_{ij}\bar{n} < \mu_{ij}s_{ijM}$. In equilibrium this holds on the routes drivers serve, since the value of joining a route falls as its queue lengthens (equation (S6)) and entry stops before the queue explodes.

S4.2.6. Alternative Value Function in Finite Time

This section solves driver value functions under a finite work day length. The aim is to microfound the exit probability (which governs the average work day length, and number of trips each driver expects to make, in the main model). For ease of notation, we abstract from driver heterogeneity in the derivation and introduce it at the end of the section.

A workday has a finite length T. We first characterize the value functions of traveling buses. In a small period of length dt, traveling buses pay a time cost cdt, and arrive with probability $\delta^A dt$ (in which case they immediately return to the start of the queue, otherwise they continue traveling). The value function is then

$$V^T(t) = -cdt + (1-\delta^A dt)V^T(t+dt) + \delta^A dt V^Q(t+dt).$$

Subtracting $V^T(t)$ from both sides, dividing by dt and letting $dt \to 0$ yields

$$V^T(t) = -t^T \left[c - \dot{V}^T(t)\right] + V^Q(t)$$

where $t^T = 1/\delta^A$ is expected travel time, and $\dot{V}^T(t) \equiv \partial V^T(t)/\partial t$. In words, the value of traveling is the time cost (both the length of the trip itself, plus the reduction in work time), plus the value of rejoining the queue after the trip is complete. We again consider a single route and drop the notational dependence of variables on ij.

In equilibrium, a queue has length N^Q. In the main text, the expected time to leave the queue is $t^Q(N^Q) = \frac{\bar{n}}{\mu s_M}(N^Q + 1)$. We therefore consider a simplified process where the bus departs the queue with probability $\delta^Q = \mu s_M/(\bar{n}(N^Q + 1))$ giving the same expected time in the queue. The value of joining the queue is therefore

$$V^Q(t) = -cdt + \delta^Q dt\left[p\bar{n} - c_M + V^T(t + dt)\right] + (1 - \delta^Q dt)V^Q(t + dt),$$

where c_M is again the lump sum monetary cost of the trip. Subtracting $V^Q(t)$ from both sides, dividing by dt and letting $dt \to 0$ yields

$$V^Q(t) = p\bar{n} - c_M - t^Q \left[c - \dot{V}^Q(t)\right] + V^T(t)$$

At the end of the period, the value of joining the queue is zero so $V^Q(T) = 0$. The system of equations is then

$$\begin{aligned}
V^T(t) &= -t^T \left[c - \dot{V}^T(t)\right] + V^Q(t) \\
V^Q(t) &= p\bar{n} - c_M - t^Q \left[c - \dot{V}^Q(t)\right] + V^T(t) \\
V^Q(T) &= 0
\end{aligned}$$

Guess and Verify. We guess the value functions take the form

$$\begin{aligned}
V^Q(t) &= \alpha^Q + \beta^Q(T - t) \\
V^T(t) &= \alpha^T + \beta^T(T - t)
\end{aligned}$$

where the dependence of V^Q on N, which is constant in equilibrium, is absorbed in the coefficients. Then $\dot{V}^T(t) = -\beta^T$ and $\dot{V}^Q(t) = -\beta^Q$. Plugging this into these two equations gives

$$\begin{aligned}
\alpha^T + \beta^T(T - t) &= -t^T \left[c + \beta^T\right] + \alpha^Q + \beta^Q(T - t) \\
\alpha^Q + \beta^Q(T - t) &= p\bar{n} - c_M - t^Q \left[c + \beta^Q\right] + \alpha^T + \beta^T(T - t)
\end{aligned}$$

Equating coefficients we get $\beta^T = \beta^Q = \beta$, and substituting this in and doing the same for the α shifters yields

$$\alpha^T = -t^T \left[c + \beta\right] + \alpha^Q$$

$$\alpha^Q = p\bar{n} - c_M - t^Q [c + \beta] + \alpha^T$$

Substituting the second into the first yields

$$0 = p\bar{n} - c_M - t^T [c + \beta] - t^Q [c + \beta]$$
$$\beta = \frac{p\bar{n} - c_M - ct^T - ct^Q}{t^T + t^Q}$$

The boundary condition implies $\alpha^Q = 0$, which delivers a final expression for

$$\alpha^T = -t^T [c + \beta]$$

These conditions collectively verify the guess. Putting these together delivers the functional form for the value of joining a queue at time $t = 0$ as

$$V^Q(0) = \underbrace{\frac{T}{t^T + t^Q}}_{N^{\text{Trips}}} \times \underbrace{\left[p\bar{n} - c_M - c(t^T + t^Q)\right]}_{\pi^{\text{PerTrip}}}$$

where $N^{Trips} \equiv \frac{T}{t^T + t^Q}$ is the number of trips the driver can make in a period of length T. As in Section S4.2.3, we simplify by assuming $c = 0$.

In the paper, drivers are heterogeneous in their workday length on each route, so that their workday length is $T_{ij\varphi} = T\nu_{ij\varphi}$.

S4.2.7. Derivation of Change in Consumer Surplus

We first compute the change in welfare on a single route, dropping the dependence on ij. Consumer surplus is given by

$$CS = \frac{1}{\gamma\theta} \ln \left(1 + \sum_{m \in \mathcal{M}} \exp\left(\alpha_m - \theta\eta t_m - \theta\gamma p_m\right)\right)$$

where we impose the normalization of the utility of the outside option $u_0 = 0$.

We want to know how this changes after mode P is added to the set of available modes which become $\mathcal{M}' = \mathcal{M} \cup \{P\}$.

The change in consumer surplus is given by

$$\begin{aligned}
\Delta CS &= \frac{1}{\gamma\theta} \ln \left(\frac{1 + \exp\left(\alpha_M - \theta\eta t'_M - \theta\gamma p'_M\right) + \exp\left(\theta u_P\right)}{1 + \exp\left(\alpha_M - \theta\eta t_M - \theta\gamma p_M\right)}\right) \\
&= \frac{1}{\gamma\theta} \ln \left(\frac{1 + \exp\left(\alpha_M - \theta\eta t'_M - \theta\gamma p'_M\right)}{1 + \exp\left(\alpha_M - \theta\eta t_M - \theta\gamma p_M\right)} + \frac{\exp\left(\theta u_P\right)}{1 + \exp\left(\alpha_M - \theta\eta t_M - \theta\gamma p_M\right)}\right) \\
&= \frac{1}{\gamma\theta} \ln \left(\frac{1 + \exp\left(\alpha_M - \theta\eta t'_M - \theta\gamma p'_M\right)}{1 + \exp\left(\alpha_M - \theta\eta t_M - \theta\gamma p_M\right)} + \frac{\exp\left(\theta u_P\right)}{\sum_{n \in \mathcal{M}'} \exp\left(\theta u'_n\right)} \frac{\sum_{n \in \mathcal{M}'} \exp\left(\theta u'_n\right)}{1 + \exp\left(\alpha_M - \theta\eta t_M - \theta\gamma p_M\right)}\right) \\
&= \frac{1}{\gamma\theta} \ln \left(\frac{1 + \exp\left(\alpha_M - \theta\eta t'_M - \theta\gamma p'_M\right)}{1 + \exp\left(\alpha_M - \theta\eta t_M - \theta\gamma p_M\right)} + s'_P \exp\left(\gamma\theta\left(CS' - CS\right)\right)\right)
\end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{\gamma\theta} \ln \left(s_0 + \frac{\exp(\alpha_M - \theta\eta t_M - \theta\gamma p_M)}{1 + \exp(\alpha_M - \theta\eta t_M - \theta\gamma p_M)} \frac{\exp(\alpha_M - \theta\eta t'_M - \theta\gamma p'_M)}{\exp(\alpha_M - \theta\eta t_M - \theta\gamma p_M)} + s'_P \exp(\gamma\theta\Delta CS) \right) \\
&= \frac{1}{\gamma\theta} \ln \left(s_0 + s_M \exp(-\theta\eta\Delta t_M - \theta\gamma\Delta p_M) + s'_P \exp(\gamma\theta\Delta CS) \right)
\end{aligned}
$$

$$
\begin{aligned}
\exp(\gamma\theta\Delta CS) &= s_0 + s_M \exp(-\theta\eta\Delta t_M - \theta\gamma\Delta p_M) + s'_P \exp(\gamma\theta\Delta CS) \\
\exp(\gamma\theta\Delta CS) &= \frac{1}{1 - s'_P} \left[s_0 + s_M \exp(-\theta\eta\Delta t_M - \theta\gamma\Delta p_M) \right] \\
\Rightarrow \Delta CS &= \frac{1}{\gamma\theta} \ln \left[\frac{1}{1 - s'_P} \times \left[s_0 + s_M \exp(-\theta\eta\Delta t_M - \theta\gamma\Delta p_M) \right] \right]
\end{aligned}
$$

Summing across all routes gives the change in aggregate surplus

$$
\Delta CS = \sum_{ij} \mu_{ij} \left[\frac{1}{\gamma\theta} \ln \left(\frac{1}{1 - s'_{ijP}} \right) + \frac{1}{\gamma\theta} \ln \left(s_{ij0} + s_{ijM} \exp(-\theta\eta\Delta t_{ijM} - \theta\gamma\Delta p_{ijM}) \right) \right].
$$

S4.2.8. Computing General Equilibrium Counterfactuals

We begin by showing the system of equations that characterize equilibrium. In the initial equilibrium we impose the free entry condition, but in the system of equations governing the change in equilibrium variables in response to public entry we will take the number of entrants and association fees F_i, τ_{ij} to be held fixed as explained in Section S4.1.4.

Equilibrium Definition. Given model parameters $\{\gamma, \eta, \alpha_m, \theta; \bar{n}, c_{ij}, \chi_{ij}, \sigma, T\}$ and data $\{\mu_{ij}, t_{ij}^T, F_i\}$, an equilibrium is a vector $\{s_{ijM}, t_{ij}^W, t_{ij}^F, \lambda_{ij}, t_{ij}^Q, p_{ij}, N_{ij}^Q, B_{ij}^T, B_{ij}, B_i, \bar{U}_{ij}\}$ such that

$$
\begin{aligned}
s_{ijM} &= \frac{\exp(\alpha_M - \theta\eta t_{ijM} - \theta\gamma p_{ijM})}{\sum_{n \in \mathcal{M}} \exp(\alpha_n - \theta\eta t_{ijn} - \theta\gamma p_{ijn})} \\
t_{ij}^W &= 1/\lambda_{ij} \\
t_{ij}^F &= \frac{1}{\mu_{ij} s_{ijM}} \left(\frac{\bar{n} - 1}{2} \right) \\
\lambda_{ij} &= \chi_{ij} \left[(1 - \delta_{ij}^E) \delta_{ij}^A B_{ij}^T + B_{ij} \right] \\
t_{ij}^Q &= \frac{\bar{n}}{\mu_{ij} s_{ijM}} (N_{ij}^Q + 1) \\
N_{ij}^Q &= \left[\frac{\mu_{ij} s_{ijM}}{\lambda_{ij} \bar{n}} - 1 \right]^{-1} \\
p_{ijM} - c_{ij}/\bar{n} &= \frac{1}{\frac{t_{ij}^Q}{t_{ij}^Q + t_{ij}^T} \left[\theta\gamma(1 - s_{ijM}) + \frac{N_{ij}^Q}{1 + N_{ij}^Q} \frac{\partial \ln N_{ij}^Q}{\partial p_{ij}} \right]} \\
B_{ij}^T &= \frac{\mu_{ij} s_{ijM} t_{ij}^T}{\bar{n}} \\
B_{ij} &= B_i \times \frac{\left(N_{ij}^{\text{Trips}} \times [p_{ijM} \bar{n} - c_{ij} - \tau_{ij}] \right)^\sigma}{\sum_k \left(N_{ik}^{\text{Trips}} \times [p_{ikM} \bar{n} - c_{ik} - \tau_{ik}] \right)^\sigma}
\end{aligned}
$$

$$0 = \left[\sum_k \left(N_{ik}^{\text{Trips}} \times [p_{ikM}\bar{n} - c_{ik} - \tau_{ik}]\right)^\sigma\right]^{1/\sigma} - F_i$$
$$\bar{U}_{ij} = \frac{1}{\theta}\ln\left(1 + \sum_{m \in \mathcal{M}} \exp\left(\alpha_m - \theta\eta t_{ijm} - \theta\gamma p_{ijm}\right)\right)$$

where $t_{ijm}, \delta_{ij}^E, \beta_{ij}, N_{ij}^{\text{Trips}}$ are auxiliary variables

$$\begin{aligned}
t_{ijM} &= t_{ijM}^T + (1-\beta_{ij})t_{ij}^W + \beta_{ij}t_{ij}^F \\
\delta_{ij}^E &= \frac{t_{ij}^Q + t_{ij}^T}{T} \\
\beta_{ij} &= \frac{\lambda_{ij}\bar{n}}{\mu_{ij}s_{ijM}} \\
N_{ij}^{\text{Trips}} &= \frac{T}{t_{ij}^Q + t_{ij}^T}
\end{aligned}$$

General Equilibrium Counterfactuals. We now show how the equilibrium system of equations can be used to compute the impacts of government entry. The shock to this system is the market share of the government entrant s'_{ij} in the post-entry period.

Letting $\hat{x} \equiv x'/x$ denote the relative change in a variable x between the pre- and post-entry period, the system of equilibrium equations (assuming the number of entrants at each terminal is fixed) can be rewritten as

$$\begin{aligned}
\hat{s}_{ijM} &= \frac{\exp\left(-\theta\eta\Delta t_{ijM} - \theta\gamma\Delta p_{ijM}\right)}{(1-s_{ijM}) + s_{ijM}\exp\left(-\theta\eta\Delta t_{ijM} - \theta\gamma\Delta p_{ijM}\right)} \times \left[1 - s'_{ijP}\right] \\
\hat{t}_{ij}^W &= \hat{\lambda}_{ij}^{-1} \\
\hat{t}_{ij}^F &= \frac{1}{\hat{s}_{ijM}} \\
\hat{\lambda}_{ij} &= \pi_{ij}^\lambda(\widehat{1-\delta_{ij}^E})\hat{B}_{ij}^T + (1-\pi_{ij}^\lambda)\hat{B}_{ij} \\
\hat{B}_{ij}^T &= \hat{s}_{ijM} \\
\hat{t}_{ij}^Q &= \frac{\widehat{N_{ij}^Q + 1}}{\hat{s}_{ijM}} \\
\widehat{N_{ij}^Q + 1} &= \frac{1-\beta_{ij}}{1-\beta_{ij}\hat{\beta}_{ij}} \\
\hat{B}_{ij} &= \frac{\left(\hat{N}_{ij}^{\text{Trips}}\right)^\sigma \left[\hat{p}_{ijM}\pi_{ij}^{rev} + 1 - \pi_{ij}^{rev}\right]^\sigma}{\sum_k \rho_{ik}\left(\hat{N}_{ik}^{\text{Trips}}\right)^\sigma \left[\hat{p}_{ikM}\pi_{ik}^{rev} + 1 - \pi_{ik}^{rev}\right]^\sigma} \\
\Delta p_{ijM} &= p_{ijM}\left(\hat{p}_{ijM} - 1\right) \\
\Delta t_{ijM} &= t_{ijM}\left(\hat{t}_{ijM} - 1\right)
\end{aligned}$$

where

$$
\begin{aligned}
\pi^{\lambda}_{ij} &= \frac{(1-\delta^E_{ij})\delta^A_{ij}B^T_{ij}}{(1-\delta^E_{ij})\delta^A_{ij}B^T_{ij}+B_{ij}} \\
\hat{t}_{ijM} &= \pi^T_{ij} + (1-\pi^T_{ij})\left[\pi^W_{ij}\widehat{(1-\beta_{ij})}\hat{t}^W_{ij} + \pi^F_{ij}\hat{\beta}_{ij}\hat{t}^F_{ij}\right] \\
\pi^T_{ij} &= \frac{t^T_{ijM}}{t_{ijM}} \\
\pi^W_{ij} &= \frac{(1-\beta_{ij})t^W_{ij}}{(1-\beta_{ij})t^W_{ij}+\beta_{ij}t^F_{ij}} \\
\pi^F_{ij} &= \frac{\beta_{ij}t^F_{ij}}{(1-\beta_{ij})t^W_{ij}+\beta_{ij}t^F_{ij}} \\
\hat{\beta}_{ij} &= \frac{\hat{\lambda}_{ij}}{\hat{s}_{ijM}} \\
\widehat{(1-\beta_{ij})} &= \frac{1-\beta_{ij}\hat{\beta}_{ij}}{1-\beta_{ij}} \\
\hat{\delta}^E_{ij} &= \frac{t^Q_{ij}}{t^Q_{ij}+t^T_{ij}}\hat{t}^Q_{ij} + \frac{t^T_{ij}}{t^Q_{ij}+t^T_{ij}} \\
\widehat{(1-\delta^E_{ij})} &= \frac{1-\delta^E_{ij}\hat{\delta}^E_{ij}}{1-\delta^E_{ij}} \\
\hat{N}^{\text{Trips}}_{ij} &= \left[\frac{t^Q_{ij}}{t^Q_{ij}+t^T_{ij}}\hat{t}^Q_{ij} + \frac{t^T_{ij}}{t^Q_{ij}+t^T_{ij}}\right]^{-1} \\
\rho_{ij} &= \frac{\left(N^{\text{Trips}}_{ij}\times[p_{ijM}\bar{n}-c_{ij}-\tau_{ij}]\right)^{\sigma}}{\sum_k\left(N^{\text{Trips}}_{ik}\times[p_{ikM}\bar{n}-c_{ik}-\tau_{ik}]\right)^{\sigma}} \\
\hat{p}_{ijM} &= \frac{p_{ijM}-c_{ij}/\bar{n}}{p_{ijM}}\widehat{p_{ijM}-c_{ij}}/\bar{n} + \frac{c_{ij}/\bar{n}}{p_{ijM}} \\
\widehat{p_{ijM}-c_{ij}}/\bar{n} &= \frac{1}{\frac{\hat{t}^Q_{ij}}{\widehat{t^Q_{ij}+t^T_{ij}}}\left[\frac{\theta\gamma(1-s_{ijM})}{\theta\gamma(1-s_{ijM})+\frac{N^Q_{ij}}{1+N^Q_{ij}}\frac{\partial\ln N^Q_{ij}}{\partial p_{ij}}}\widehat{1-s_{ijM}} + \frac{\frac{N^Q_{ij}}{1+N^Q_{ij}}\frac{\partial\ln N^Q_{ij}}{\partial p_{ij}}}{\theta\gamma(1-s_{ijM})+\frac{N^Q_{ij}}{1+N^Q_{ij}}\frac{\partial\ln N^Q_{ij}}{\partial p_{ij}}}\widehat{\frac{N^Q_{ij}}{1+N^Q_{ij}}}\widehat{\frac{\partial\ln N^Q_{ij}}{\partial p_{ij}}}\right]}
\end{aligned}
$$

This system of equations would allow us to use the model to compute full general equilibrium counterfactuals. However, they also elucidate why the model's many interactions restrict its ability to deliver clear comparative statics.

S4.2.9. Testing Profit Equalization Across Treated and Connected Routes

Our empirical results show queue lengths and prices change on treated and connected routes. We use results from the driver survey to argue this is due to route substitution by drivers. This section uses the model to test whether the empirical results on treated and connected routes are consistent with average profit equalization across routes.[54]

Taking logs of equation (4) and substituting in its components, we have

$$\log(V_{ij}^{Q}) = \log(T) - \log(t_{ij}^{T} + t_{ij}^{Q}) + \log(p_{ijM}\bar{n} - c_{ijM})$$

We next take differences and simplify using first order approximations to get

$$\begin{aligned} d\ln(V_{ij}^{Q}) &\approx -\frac{t_{ij}^{Q}}{t_{ij}^{Q} + t_{ij}^{T}} d\ln(t_{ij}^{Q}) + \pi_{ij}^{rev} d\ln(p_{ijM}) \\ &\approx -\frac{t_{ij}^{Q}}{t_{ij}^{Q} + t_{ij}^{T}} \left[\frac{N_{ij}^{Q}}{N_{ij}^{Q} + 1} d\ln(N_{ij}^{Q}) - d\ln(\mu_{ij} s_{ijM})\right] + \pi_{ij}^{rev} d\ln(p_{ijM}). \end{aligned}$$

Here we use that $d\ln(T) = d\ln(\bar{n}) = d\ln(t_{ij}^{T}) = 0$ given that none of minibus workdays, minibus occupancy or road speeds respond to the policy. In the first line we use the approximation $\log(t_{ij}^{T} + t_{ij}^{Q}) \approx \frac{t_{ij}^{Q}}{t_{ij}^{Q}+t_{ij}^{T}} d\ln(t_{ij}^{Q})$ since $d\ln(t_{ij}^{T}) = 0$. In the second line, we take a first order expansion of queue times $t_{ij}^{Q} = \frac{\bar{n}}{\mu_{ij} s_{ijM}}(N_{ij}^{Q} + 1)$ from equation (S5) to deliver the substitution in the second line. Recall also that π_{ij}^{rev} is the ratio of gross to net driver revenue defined in the text.

From Appendix Table S4 we can approximate the baseline time spent in queues $\frac{t_{ij}^{Q}}{t_{ij}^{T}+t_{ij}^{Q}} = 0.424$, and the inverse profit margin $\pi_{ij}^{rev} = 3.36$ based on reported income and expenses broken down by day. The mean queue length of 4.75 yields $\frac{N_{ij}^{Q}}{N_{ij}^{Q}+1} = 0.83$. This yields a model-implied (first order) change in profit of

$$d\log(V_{ij}^{Q}) \approx -0.42\left(0.83 \cdot d\log(N_{ij}^{Q}) - d\log(\mu_{ij} s_{ijM})\right) + 3.36 \cdot d\log(p_{ijM}). \tag{S20}$$

We compute this object for both treated and spillover routes using the specifications in even columns of Table 4 via seemingly unrelated regression. Appendix Table S29 shows the results. The changes in profits are remarkably similar, with a reduction of 0.376 log points on treated routes and 0.334 log points on connected routes. Despite the relative precision with which these objects are estimated, we fail to reject their equality (p-value 0.78). This supports the model mechanism of driver sorting across routes until profits are equalized. This corroborates the reported changes in driver profits in column (4) of Table 3, which reports no significant difference between drivers on treated and connected routes.

[54]We thank Gabriel Kreindler for this suggestion. We also note we are abstracting from the Fréchet adjustment for the idiosyncratic term, and are thus looking at the change in the common component of profits.

Table S29. Testing Equality of the Change in Profits on Treated and Connected Routes

	Treated	Connected
$\Delta \ln(V_{ij}^{Q})$	-0.376**	-0.334***
	(0.177)	(0.090)
P-value T==C	0.777	

Notes: Table reports values for change in profits on treated and connected routes using equation (S20) via seemingly unrelated regression with standard errors clustered by terminal. The last row reports a p-value for the test of equality. * p<0.1; ** p<0.05; *** p<0.01. See Appendix Section S4.2.9 for details.

S4.3. Additional Details on Quantification

Implementing the sufficient statistics approach for the welfare change of commuters and minibus drivers in (9) and (10) requires measuring three sets of objects.

S4.3.1. Measuring Descriptive Statistics

First, we require a collection of descriptive statistics $(s_{ijm}, s'_{ijP}, \pi_{ij}^{rev}, B_i, \rho_{ij}, \mu_{ij})$, which we construct as follows. In keeping with the model, μ_{ij} denotes the number of commuter trips and B_i the number of minibus entrants at a terminal; every other count is written N with a superscript naming what it counts: N_{ij}^{Trips} for daily trips per driver, N_{ij}^{Drivers} for drivers on a route, N_{i}^{Routes} for routes, and N_{ij}^{Q} for buses in the queue.

Mode shares s_{ijm}, s'_{ijP}. Prior to public transit entry there are only two modes. The 2009 Lagos Travel Survey gives a minibus share of all trips of $s_{ijM} = 0.56$, so the outside option share is $s_{ij0} = 1 - s_{ijM} = 0.44$, made up of all non-minibus modes. This is a share of all trips; the 62% minibus share reported in Section 3 is among motorized trips only, and so excludes walking.

In the post-entry period, the share of public transit users is given by

$$s'_{ijP} = 1 - s_{ijM}\hat{s}_{ijM} - s_{ij0}\hat{s}_{ij0} \tag{S21}$$

where hats denote relative changes. In words, public transit users come either from substitution from minibuses or the outside option. To estimate $\hat{s}_{ij0}$, we differentiate the expression for total trips $\mu_{ij} = \mu_{ij}^{0} + \mu_{ij}^{\text{Transit}}$ with respect to public transit entry. Note here that $\mu_{ij}^{\text{Transit}}$ is the number of *total* transit trips across both public transit and minibuses. Since the new public option provides the same travel times as existing minibuses, we assume the total number of trips is unchanged (i.e. the shock is small enough that any change in total trips is negligible) although the share of trips taken using minibuses and the outside option can change. Taking logs and differentiating, and exponentiating the resulting first-order change, then delivers

$$\hat{s}_{ij0} = \exp\left(-\frac{1-s_{ij0}}{s_{ij0}} \frac{d \ln \mu_{ij}^{\text{Transit}}}{d\text{Open}_{ij}}\right)$$

Intuitively, if total trips are unchanged, then the more travelers use transit in the post-entry period, the fewer choose the outside option. Table S7 shows that *total* transit trips increase by 0.149 log points (standard error 0.088) after public transit opens on a route, implying that $\hat{s}_{ij0} = 0.823$. We estimate $\hat{s}_{ijM} = \exp\left(\frac{d \ln Pass_{ijM}}{d\text{Open}_{ij}}\right)$ using our main specification but with log passengers as the outcome and find a semi-elasticity of -0.298 (standard error 0.153).

Substituting these values into (S21), we obtain a post-entry public share of $s'_{ijP} = 0.222$, a post-entry minibus share of $s'_{ijM} = 0.416$, and a post-entry outside option share of $s'_{ij0} = 0.362$. The total share of transit users went from 0.56 prior to public entry to 0.64 after (on routes where public transit is available).

Profit margins π_{ij}^{rev}. We measure π_{ij}^{rev} directly from the driver surveys as the ratio of gross revenue to net income, asking each driver their total revenue and the income they kept after all fees and expenses on their last working day. It averages 3.3.

Total Entrants and Route Choice Probabilities. Our welfare decomposition in (10) requires data only from terminals where the public system operates, as we argue there are no effects on routes where public transit is unavailable at either endpoint. We construct our measures of B_i, ρ_{ij} in two steps: first, by estimating key inputs within our baseline sample of observed routes, and then by extending these inputs to all routes at treated terminals to determine total entrants and driver share distributions.

Total Entrants B_i. Using baseline driver survey data, we regress log daily trips on log trip distance to estimate daily trips per driver per route $N_{ij}^{\text{Trips}} = \exp\left(\hat{\alpha} + \hat{\beta} \ln Dist_{ij}\right)$ as a function of its distance, where hats denote estimated values. We then estimate total minibus daily trips per route by summing bus departures within our three observation windows and applying two adjustments. First, since these windows capture only part of the day, we scale up based on e-ticketing data, which indicates that 40% of weekday trips occur during the observation windows. We assume the same distribution across times of day applies to minibus trips. Second, we account for "sole" trips—those not originating from terminals—estimated at 21% of driver trips based on survey data. Letting $\text{ObservedDriverTrips}_{ij}$ be the number of minibus trips we observe, total daily minibus trips on the route are $\text{TotalDriverTrips}_{ij} = \frac{1}{0.4\times 0.79}\text{ObservedDriverTrips}_{ij}$. These equal the number of drivers times their trips per driver, $\text{TotalDriverTrips}_{ij} = N_{ij}^{\text{Drivers}} N_{ij}^{\text{Trips}}$, so we recover the number of drivers on route ij at baseline from

$$N_{ij}^{\text{Drivers}} = \frac{\text{TotalDriverTrips}_{ij}}{N_{ij}^{\text{Trips}}}. \tag{S22}$$

This measure includes only routes covered in our route observation survey, which may not capture all routes at a terminal. For terminals with public service which are in our observation survey sample, we use our network census to measure the number of treated routes $N^{\text{Routes}}_{i,\text{Treat}}$ and the number of connected routes $N^{\text{Routes}}_{i,\text{Connected}}$. We use the measure of N^{Drivers}_{ij} from above to compute the average number of drivers on treated routes $\bar{N}^{\text{Drivers}}_{i,\text{Treat}}$ and connected routes $\bar{N}^{\text{Drivers}}_{i,\text{Connected}}$, allowing us to compute total entrants as:

$$B_i = N^{\text{Routes}}_{i,\text{Treat}} \times \bar{N}^{\text{Drivers}}_{i,\text{Treat}} + N^{\text{Routes}}_{i,\text{Connected}} \times \bar{N}^{\text{Drivers}}_{i,\text{Connected}}.$$

For terminals with public service not covered in our survey sample, we use an observation we conducted for all routes in the network census immediately following the route mapping to compute the total number of drivers using (S22). For these terminals, $B_i = \sum_j N^{\text{Drivers}}_{ij}$. Ideally, we would measure total entrants for these routes at baseline rather than relying on 2022 network census data. However, the two measures are highly correlated (correlation coefficient 0.96) for terminals included in both datasets, suggesting this approach is reliable.

Distribution of drivers ρ_{ij}. Since treatment effects are constant across treated and connected routes, we only need to measure the fraction of drivers at terminals with public entry who operate on treated routes; the remainder operate on connected routes. For terminals with public service included in our observation survey sample, we compute this directly as $s^{\text{Driver}}_{i,\text{Treat}} = N^{\text{Routes}}_{i,\text{Treat}} \times \bar{N}^{\text{Drivers}}_{i,\text{Treat}} / B_i$. We then regress the share of treated drivers on the share of treated routes for these terminals ($R^2 = 0.71$) and use the fitted values to predict the treated driver share at terminals with public service that are not in our survey sample, based on the share of their routes that are treated (which we observe in the network census). Altogether, this delivers $\rho^{\text{Driver}}_{i,\text{Treat}} = \sum_{j \in \text{Treated}_i} \rho_{ij}$ and $\rho^{\text{Driver}}_{i,\text{Connected}} = 1 - \rho^{\text{Driver}}_{i,\text{Treat}}$. These are the terms we need to estimate the change in driver surplus.[55]

Total number of commuter trips μ_{ij}. To measure the total number of trips μ_{ij}, we start with the identity $\mu_{ij}(1 - s'_{ij0}) = \mu_{ij} s'_{ijP} + \mu_{ij} s'_{ijM}$. The righthand side is total transit trips and is observed directly: public transit trips $\mu_{ij} s'_{ijP}$ come from e-ticketing records, and minibus trips $\mu_{ij} s'_{ijM}$ come from our network census observation survey. We measure these in the period where the two datasets overlap, so that they sum to total transit trips at a single point in time.

However, we need to make three adjustments to our observed transit trips, $\text{ObservedTransitTrips}_{ij}$ (passenger boardings from the e-ticketing records and observation survey above), to get to the total number of transit trips. First, we only observe transit trips during our enumerator observation windows. This accounts for 40% of public trips, and we assume the same distribution applies to minibus trips. Second, we adjust for the 21% of bus trips that are "sole" and do not originate from terminals. Third, we adjust for the fact that some passengers board along the route rather than at the origin terminal. Based on GPS tracking data collected by enumerators during network mapping, 92% of passengers along a route board at the origin terminal.

[55] Since treatment effects are constant on treated and connected routes, the change in surplus at treated terminals is

$$\hat{\Pi}^V_i = \left[\rho^{\text{Driver}}_{i,\text{Treat}} \left(\hat{N}^{\text{Trips}}_{\text{Treat}}\right)^{\sigma} \left[\hat{p}_{\text{Treat}} \pi^{rev}_{ij} + 1 - \pi^{rev}_{ij}\right]^{\sigma} + \rho^{\text{Driver}}_{i,\text{Connected}} \left(\hat{N}^{\text{Trips}}_{\text{Connected},i}\right)^{\sigma} \left[\hat{p}_{\text{Connected},i} \pi^{rev}_{ij} + 1 - \pi^{rev}_{ij}\right]^{\sigma}\right]^{1/\sigma}$$

where $\hat{N}^{\text{Trips}}_{\text{Treat}}$ and $\hat{N}^{\text{Trips}}_{\text{Connected},i}$ are the relative changes in daily trips per driver on treated and connected routes, and $\hat{p}_{\text{Treat}}$ and $\hat{p}_{\text{Connected},i}$ the corresponding fare changes. The connected-route quantities carry the subscript i because they vary across terminals with the number of open public routes.

We therefore inflate our observed transit trips to total transit trips from

$$\mu_{ij}^{\text{Transit}} = \frac{1}{0.4 \times 0.79 \times 0.92} \times \text{ObservedTransitTrips}_{ij}$$

and therefore $\mu_{ij} = \mu_{ij}^{\text{Transit}}/(1 - s'_{ij0})$, where s'_{ij0} was obtained in the previous step.

S4.3.2. Measuring Responses to Public Entry

We require estimates of the changes in minibus prices (Δp_{ijM}, $\hat{p}_{ijM}$), commuter wait times (Δt_{ijM}) and driver daily trips ($\hat{N}_{ij}^{\text{Trips}}$) in response to public entry. These come from our reduced form results.

We find no effects on minibus departures on connected routes, and so define a change in commuter wait times only on treated routes equal to $\Delta t_{ijM} = \bar{t}_{ijM} \times 0.22$, where $\bar{t}_{ijM}$ is the average initial wait time on treated routes and $0.22 = -\frac{\partial \ln Dep_{ij}}{\partial \text{Open}_{ij}}$ is minus one times the semi-elasticity of departures to public entry from column (7) of Table 2. Under Poisson departures, this corresponds to minus one times the semi-elasticity of wait times to public entry, so Δt_{ijM} is the first-order approximation to $\bar{t}_{ijM}(e^{0.22} - 1)$.

We set price changes on treated and connected routes based on the spillover specification results from column (4) of Table 4. On treated routes, $\hat{p}_{ijM} = \exp(-0.108)$, while on connected routes, $\hat{p}_{ijM} = \exp(-0.016 \times N_{i,\text{Connected}}^{\text{Routes}})$, where $N_{i,\text{Connected}}^{\text{Routes}}$ is the number of connected routes at terminal i. The change in prices is then $\Delta p_{ijM} = \bar{p}_{ijM} \ln \hat{p}_{ijM}$, where $\bar{p}_{ijM}$ is the average initial fare, measured separately for treated and connected routes.

We set the relative change in the number of trips per driver to $\hat{N}_{ij}^{\text{Trips}} = \exp(-0.22) = 0.80$, the change in departures from column (7) of Table 2. This is similar to the driver-survey estimate in column (1) of Table 3 ($1 - 1.55/9.58 = 0.84$), but measures trips on the treated route itself rather than trips of drivers who may have switched routes. For connected routes, the coefficient in the continuous spillover specification in Table 5 is highly imprecise, while the dummy spillover specification in Table S24 is more precise and larger in magnitude. Observation surveys indicate no change in departures but an increase in queue lengths on connected routes, suggesting a decline in trips per driver. Given this collective evidence, we allow trips per driver on connected routes to decrease and set the change to 38.2% of the decline on treated routes, based on the ratio of coefficients on connected and treated routes in Panel B of Table S24 ($0.382 = -0.452/-1.183$).

S4.3.3. Measuring Commuter and Driver Preference Parameters

Estimation of commuter preference parameters θ, η, γ is described in the main text.

We estimate the driver route choice parameter σ using the responsiveness of route choice to the profit shock represented by public entry. The fraction of drivers choosing route ij in terminal i depends on its profits relative to other routes from the terminal:

$$\rho_{ij} = \left(V_{ij}^{Q}/\Pi_{i}^{V}\right)^{\sigma}$$

We know that

$$\sigma = \frac{1}{1-\rho_{ij}} \frac{\partial \ln \rho_{ij}}{\partial \ln V_{ij}^Q} = \frac{1}{1-\rho_{ij}} \frac{\partial \ln \rho_{ij} / \partial \text{Open}_{ij}}{\partial \ln V_{ij}^Q / \partial \text{Open}_{ij}},$$

and therefore estimate σ by combining our reduced-form semi-elasticities that tell us how public entry changed profits and entry decisions on different routes.

To compute the expression in the numerator, we pool routes in each treated terminal into two groups: those which receive public service by the end of the driver survey and those which do not. We do this to increase the number of drivers in each cell. We then look, within each terminal, at the change in drivers choosing routes which receive public service through the regression

$$\ln \rho_{ict} = \gamma_{ic} + \delta_{it} + \beta \text{Open}_{ict} + \epsilon_{ict}$$

where i is a terminal, c is a category (those which receive public service by the end of the survey for which $\text{Open}_{ict} = 1$ in the endline survey), and $t = 0, 1$ indicates the baseline and endline surveys respectively. This semi-elasticity is therefore identified from the increase in route switching for drivers who drive treated routes at baseline, between the baseline and endline surveys (to try and capture the long-run response). We find an estimate of $\beta = \partial \ln \rho_{ij} / \partial \text{Open}_{ij}$ of -0.377 (0.131) when estimating in logs and -0.277 (0.082) when estimating using PPML.

To compute the expression in the denominator, we differentiate equation (4) to get

$$\frac{\partial \ln V_{ij}^Q}{\partial \text{Open}_{ij}} = \frac{\partial \ln N_{ij}^{Trips}}{\partial \text{Open}_{ij}} + \pi_{ij}^{rev} \frac{\partial \ln p_{ij}}{\partial \text{Open}_{ij}}$$

As in the prior step, we use our within-terminal results from Table 3 to measure these semi-elasticities, using $\frac{\partial \ln N_{ij}^{Trips}}{\partial \text{Open}_{ij}} = \ln((9.58 - 1.55)/9.58) = -0.1765$ and set $\frac{\partial \ln p_{ij}}{\partial \text{Open}_{ij}} \approx 0$ from that table. We also measure that $E\left[\frac{1}{1-\rho_{ij}}\right] = 1.298$.

Putting these results together, we arrive at

$$\sigma = E\left[\frac{1}{1-\rho_{ij}}\right] \frac{\partial \ln \rho_{ij} / \partial \text{Open}_{ij}}{\partial \ln V_{ij}^Q / \partial \text{Open}_{ij}} = 1.298 \times \frac{-0.377}{-0.1765} = 2.77,$$

using our OLS estimate of $\partial \ln \rho_{ij} / \partial \text{Open}_{ij}$, and $\sigma = 2.04$ when using the PPML estimate.

S4.4. Alternatives

S4.4.1. Including Not Taking a Trip in the Outside Option

The main model interprets μ_{ij} as the number of travelers on a route, since we calibrate mode shares using the 2009 Lagos Travel Survey, which conditions on making a trip. This section adds the choice of not taking a trip to the outside option, using the 2024 Nigerian Time Use Survey (introduced in Section 3.1) to measure the share of Lagos residents who make any trip on a given day.

Calibration. Including not taking a trip in the outside option leaves every model equation unchanged; only the measurement of the shares in the pre- and post-entry periods is affected. Formally, we now let μ_{ij} be the number of *potential* travelers on route ij: those who travel, plus those who could but choose not to. Let $q_{ij} \in (0, 1]$ be the share of these potential travelers who make any trip, so the main calibration is the case $q_{ij} = 1$.

In the pre-entry period, the 2009 travel survey gives a minibus share of 0.56 among those who travel, so across all potential travelers

$$s_{ijM} = q_{ij} \cdot 0.56, \qquad s_{ij0} = 1 - q_{ij} \cdot 0.56. \tag{S23}$$

A lower travel rate shrinks the minibus share and expands the outside option, which now also contains the non-travelers.

In the post-entry period, we compute the public transit share s'_{ijP} from equation (S21) as before, now evaluated at the pre-entry shares from equation (S23).

Lastly, we recover the number of potential travelers from $\mu_{ij} = \mu_{ij}^{\text{Transit}}/(1 - s'_{ij0})$, as in Online Appendix S4.3.1, evaluated at the new value of s'_{ij0}.

Relative to the main calibration, $q_{ij} < 1$ raises the outside option share, lowers the transit shares, and raises the implied number of potential travelers.

Measurement. We measure q_{ij} from the 2024 Nigerian Time Use Survey, whose diary records each individual's main activity (and its location) in 24 hourly slots over a day. In the Lagos sample of $1{,}554$ valid diaries, 87.8% of respondents left home at least once on the diary day, so we set $q_{ij} = 0.878$ on every route.

Results. Table S30 compares the main calibration (row 1, $q_{ij} = 1$) with the calibration that includes not taking a trip in the outside option (row 2, $q_{ij} = 0.878$). The lower travel rate raises the outside option share and lowers both transit shares (columns (1)-(3)). It also raises the implied number of potential travelers (column (4)): with a smaller fraction of people traveling, more potential travelers are needed to generate the same number of transit trips. The aggregate gain is essentially unchanged (it falls by 0.76%), but it is spread among more people, because the post-entry public transit share and the pre-entry minibus share both fall. Surplus per person falls by 13.4% on treated routes and 12.2% on connected routes.

Table S30. Consumer Surplus when the Outside Option Includes Not Traveling

Travel rate q	s_0	s_M	s'_P	Implied μ	Treated ($/day)	Connected ($/day)	Aggregate ($mn/month)
$q = 1$	0.434	0.566	0.222	1.00×	$0.234	$0.0282	$1.93
$q = 0.878$	0.503	0.497	0.197	1.14×	$0.204	$0.0248	$1.91

Notes: Consumer surplus from public transit entry when the outside option includes not traveling, following Appendix S4.4.1.

S4.4.2. Nested Logit

Setup. Throughout this section we drop the route index ij for convenience. We now assume ϵ is drawn from a nested T1EV distribution that places the outside option in one nest and the transit modes $\mathcal{T} = \{M, P\}$ in another,

$$F(\boldsymbol{\epsilon}) = \exp\left(-\left[\exp(-\theta\epsilon_0) + \left(\sum_{m\in\mathcal{T}} \exp\left(-\lambda\epsilon_m\right)\right)^{\theta/\lambda}\right]\right).$$

Here θ governs substitutability between the outside option and the transit nest, and λ substitutability between the two transit modes; we expect $\lambda > \theta$.

Computing the change in welfare from removing public transit, as in the multinomial logit case, delivers

$$\Delta U = \frac{1}{\theta}\ln\left(s_0 + s_T\left[\frac{1}{1 - s'_{P|T}}\exp\left(-\lambda\eta\Delta t_m - \lambda\gamma\Delta p_m\right)\right]^{\theta/\lambda}\right).$$

where s_T is the share of users choosing either transit mode, and $s_{P|T}$ is the share of transit users who choose public transit. The expression has the same structure as in the multinomial logit: the welfare change depends on the post-entry public share, here the within-nest share $s'_{P|T}$, and on the change in minibus travel times and fares. The two coincide when $\lambda = \theta$, since the nested logit nests the multinomial logit.

Estimation.

Data. In March 2024 we ran the commuter survey introduced in Section 3.1, covering commuters living within 2.5km of major public transit terminals to learn about their travel patterns and their perceptions of public transit and minibuses, with a second round in November 2024. We surveyed households living near 10 of the 15 major terminals and targeted regular transit users, restricting eligibility to those who had taken at least two transit trips in the prior week.

To estimate the nested logit model, we use each participant's longest trip on the last weekday they traveled. We treat a respondent as traveling on a route where public transit is available if both the origin and destination of their trip fall within 2km of a public route. This assumes commuters will walk up to 2km, about a 20-25 minute walk, to reach or leave a direct public route. This leaves us with 540 commuters; we vary the 2km threshold in robustness checks. We observe public transit prices and wait times directly from the e-ticketing network, and observe minibus prices and wait times for commuters who chose the minibus. For commuters who chose public transit, we lack contemporaneous data on the minibus alternative they did not take, so we impute it.[56]

Estimating Equations We estimate the within-nest substitution parameter λ from the choices of commuters who take a transit mode on trips where both public transit and minibuses are available. In the nested logit, conditional on choosing transit, the choice probabilities follow a standard logit:

$$s_{P|T} = \frac{\exp\left(\lambda\left(\alpha_P - \eta t_P - \gamma p_P\right)\right)}{\sum_{m\in\{P,M\}} \exp\left(\lambda\left(\alpha_m - \eta t_m - \gamma p_m\right)\right)}.$$

When estimating this equation we impose the value of time $\eta/\gamma = 18.94$ from the wait time experiment, as the theory implies.[57]

We then identify the between-nest substitution parameter θ from the same moment as the price quasi-experiment. In the multinomial model, we use that the semi-elasticity of public transit trips to its price is $\frac{\partial \ln N_P}{\partial p_P} = -(1 - s_P)\theta\gamma$. Under nested logit, this semi-elasticity becomes

$$\frac{\partial \ln N_P}{\partial p_P} = -\gamma\left[\lambda(1 - s_{P|T}) + \theta s_{P|T}(1 - s_T)\right]. \tag{S24}$$

Given λ, data on shares, the trip semi-elasticity, and γ from the wait time experiment, this expression pins down a unique θ. We invert it for θ using the estimate of λ obtained above. Requiring $0 < \theta < \lambda$ imposes tight bounds on λ, which our estimates satisfy.

Parameter Estimates Table S31 presents the results. The first two rows report the coefficients from the conditional logit, $\hat{\beta}_{\text{wait}} = -\lambda\eta$ and $\hat{\beta}_{\text{price}} = -\lambda\gamma$. The next two report the implied within- and between-nest substitution parameters $\hat{\lambda}$ and $\hat{\theta}$.

[56] We assign these commuters the average minibus wait time from Table 1 and a fare predicted from a regression of fare on distance among contemporaneous minibus users.

[57] Relaxing this constraint moves λ from 0.892 to 1.054 in our benchmark specification (column (1)). Because the feasible region $0 < \theta < \lambda$ is narrow, this pushes the implied θ below zero through equation (S24), so we use the constrained estimate.

Table S31. Nested Logit Mode-Choice Estimates

	(1)	(2)	(3)	(4)
$\hat{\beta}_{\text{wait}}$	-0.0414***	-0.0413***	-0.0469***	-0.0473***
	(0.0043)	(0.0044)	(0.0049)	(0.0052)
$\hat{\beta}_{\text{price}}$	-0.0022***	-0.0022***	-0.0025***	-0.0026***
	(0.0002)	(0.0002)	(0.0003)	(0.0003)
$\hat{\lambda}$	0.892	0.890	1.011	1.020
$\hat{\theta}$	0.775	0.787	0.075	0.021
N	540	540	465	465
Demographic controls		X		X
Narrower route definition			X	X

Notes: See Online Appendix S4.4.2 for details. Rows 1 and 2 report the wait time and price coefficients from a conditional logit, with standard errors in parentheses. The estimation imposes the value of time $\eta/\gamma = 18.94$ from the wait time experiment. Column (1) runs the main specification without controls. Column (2) adds controls for age and gender, interacted with mode-specific constants. Columns (3) and (4) repeat the specifications, but counts public transit as available when a trip's origin and destination both lie within 1.5km of a public route rather than 2km. Robust standard errors reported in parentheses. * p<0.1; ** p<0.05; *** p<0.01.

Table S32. Nested Logit Mode-Choice Estimates with Mode Characteristics

	(1)	(2)	(3)	(4)
$\hat{\beta}_{\text{wait}}$	-0.0414***	-0.0407***	-0.0414***	-0.0415***
	(0.0043)	(0.0043)	(0.0044)	(0.0044)
$\hat{\beta}_{\text{price}}$	-0.0022***	-0.0022***	-0.0022***	-0.0022***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)
$\hat{\beta}_{\text{index}}$		0.2421**		0.4707***
		(0.0945)		(0.1653)
$\hat{\lambda}$	0.8915	0.8770	0.8928	0.8939
$\hat{\theta}$	0.7753	0.8602	0.7677	0.7616
N	540	540	540	540
Demographic controls			X	X

Notes: See Online Appendix S4.4.2 for details. Rows 1 and 2 report the wait time and price coefficients from a conditional logit, with standard errors in parentheses as in Table S31. Row 3 reports coefficients on an index of mode characteristics, defined as the average of mode-by-round mean ratings for comfort, safety, and reliability across all respondents, with standard errors in parentheses. The estimation imposes the value of time $\eta/\gamma = 18.94$ from the wait time experiment. Column (1) replicates column (1) in Table S31. Column (2) adds the mode-characteristics index to column (1). Columns (3) and (4) include controls for age and education, interacted with an indicator for the public-transit alternative. Robust standard errors reported in parentheses. * p<0.1; ** p<0.05; *** p<0.01.

In our benchmark specification, column (1), we estimate $\hat{\lambda} = 0.892$ and $\hat{\theta} = 0.775$, satisfying $0 < \theta < \lambda$. Column (2) adds controls for age and gender interacted with mode-specific constants (to allow tastes for modes to vary by demographics). Columns (3) and (4) repeat these two specifications but tighten the threshold for treating a route as served by public transit from 2km to 1.5km. The estimates of λ are fairly stable across the four columns, while θ is considerably more sensitive, reflecting the nonlinearity of equation (S24).

Because our household survey is small and non-representative, we compare our estimate against one from another developing-country setting, Tsivanidis (2026). That paper estimates a nested logit of transport demand with a similar structure, grouping the transit modes in a single nest. Although the Colombian context differs, applying the same trip semi-elasticity condition (equation (S24)) to its estimate implies $\lambda = 0.994$, within the range we find. We use both values in the welfare analysis below as a robustness check.

Welfare Results. Table S33 reports the consumer surplus impacts from Table 7 under each model of travel demand. Panel A reproduces the multinomial logit benchmark. Panels B and C use the nested logit, with λ and θ estimated from our household survey (Panel B) and calibrated from Tsivanidis (2026) (Panel C).

The benefits to commuters from public transit entry are smaller under nested logit demand, though qualitatively similar. With the Lagos estimates the total gain falls by 3.1% ($1.87mn/month against $1.93mn/month), and with the Bogotá values it falls by 14%. Nested logit affects the direct and indirect components differently, so we take each in turn.

First, consider the direct benefit of the new mode. This benefit depends on how commuters substitute across modes, and so the different substitution patterns implied by each model matter for the welfare change. We estimate $\lambda > \theta$, so the two transit modes are closer substitutes for each other than for the outside option. A given post-entry public transit share then implies a smaller utility gain than under multinomial logit, because more of the take-up is drawn from the minibus than from non-travel, and riders who switch from the minibus already had a close substitute. The sign of this effect is intuitive, but its magnitude turns out to be small with our estimates, because λ is only modestly above θ; with the Bogotá calibration, where λ is close to one, the direct benefit falls by 18%.

Second, the indirect effects that operate through the minibus response barely depend on these substitution patterns. To a first order, in fact, the impacts are equal across models. To see this, write consumer surplus under multinomial and nested logit as

$$\begin{aligned} CS^{\text{MNL}} &= \frac{1}{\gamma\theta}\ln\left(1+\sum_m \exp\left(\theta\left(\alpha_m-\gamma p_m-\eta t_m\right)\right)\right) \\ CS^{\text{NL}} &= \frac{1}{\gamma\theta}\ln\left(1+\left[\sum_m \exp\left(\lambda\left(\alpha_m-\gamma p_m-\eta t_m\right)\right)\right]^{\theta/\lambda}\right). \end{aligned}$$

For a small change in a minibus attribute, say $dt_M > 0$, the change in welfare under either model is

$$\begin{aligned} dCS^{\text{MNL}} &= -\frac{1}{\gamma\theta}s_M \times \theta\eta dt_M = -\frac{\eta}{\gamma}s_M dt_M \\ dCS^{\text{NL}} &= -\frac{1}{\gamma\theta}s_M \times \frac{\theta}{\lambda} \times \lambda\eta dt_M = -\frac{\eta}{\gamma}s_M dt_M. \end{aligned}$$

This equivalence reflects the same envelope logic behind the Hulten theorem. When a minibus attribute changes, the inframarginal riders who stay with the minibus are affected directly, while some riders switch away. However, the switchers are marginal, i.e. indifferent between modes, and so experience no first order change in welfare. Only the inframarginal riders, a fraction s_M, bear the change, and for wait times they value it at η/γ. The substitution parameters (θ under multinomial logit, θ and λ under nested logit) govern how many riders switch and where they go, but to first order that reallocation does not affect welfare.[58]

[58] The pre-entry level of surplus differs substantially across models because it is inversely proportional to θ. Before public transit enters there is no nest, only the incumbent minibus versus the outside option, so λ drops out and expected consumer surplus is $CS_0 = \frac{1}{\gamma\theta}\ln\frac{1}{1-s_T}$, with s_T the pre-entry transit share.

Table S33. Consumer Surplus under Multinomial and Nested Logit

		Individuals ($/individual/day)		**Aggregate** ($mn/month)
		Treated	Connected	
Panel A: Multinomial Logit	*Number:*	**240,353**	**877,066**	**1,117,419**
Baseline surplus from private transit		0.83	0.83	22.06
Effect of introducing public transit		+0.23	+0.03	+1.93
Direct benefit of public transit		+0.25	0	+1.43
Additional impact from private Δdepartures		-0.04	0	-0.25
Additional impact from private Δprices		+0.03	+0.03	+0.75
Panel B: Nested Logit (Estimated λ=0.892)	*Number:*	**240,353**	**877,066**	**1,117,419**
Baseline surplus from private transit		0.94	0.94	24.88
Effect of introducing public transit		+0.22	+0.03	+1.87
Direct benefit of public transit		+0.24	0	+1.39
Additional impact from private Δdepartures		-0.05	0	-0.29
Additional impact from private Δprices		+0.03	+0.03	+0.77
Panel C: Nested Logit (Calibrated λ=0.994)	*Number:*	**240,353**	**877,066**	**1,117,419**
Baseline surplus from private transit		5.29	5.29	140.67
Effect of introducing public transit		+0.19	+0.03	+1.66
Direct benefit of public transit		+0.21	0	+1.17
Additional impact from private Δdepartures		-0.05	0	-0.26
Additional impact from private Δprices		+0.03	+0.03	+0.74

Notes: Consumer surplus from public transit entry under the multinomial logit of the main text (Panel A) and under the nested logit of Appendix S4.4.2, using the estimated λ from column (1) of Table S31 (Panel B) and a calibrated value of λ (Panel C) from Tsivanidis (2026). Each panel follows the decomposition of Table 7. Only point estimates are provided for this comparison.

Appendix S5. Additional Empirical Results on Response of Private Transit

This section provides additional analysis of the response of private transit to public transit entry, building on Section 5.

S5.1. Placebo Tests

We conduct a falsification test by applying the same strategy to estimate effects on two sets of planned routes that never became operational. Using the dates the government planned to open these routes, we construct hypothetical Open_{rt} dummies which turn on when these routes were scheduled to be operational. If the main treatment effects are driven by public entry itself, rather than pre-existing trends on targeted routes, these planned opening dummies should have no impact on outcomes.

The EndSARS Protests and Fire at Oyingbo Terminal. The new Oyingbo terminal was slated to open and host 8 new public bus routes starting in late October and early November 2020. However, weeks before opening, a protest against police brutality swept Nigeria, illuminating abuses by a police unit called the Special Anti-Robbery Squad (SARS). Following a fatal police shooting at Lekki Tollgate on October 20, groups of protesters burned a television station and the Oyingbo public bus terminal (see Appendix Figure S9), which canceled the planned routes.

Figure S9. Oyingbo Public Terminal: Fire and Aftermath

A. Fire

B. Aftermath

Evolving Opening Plans. The government's rollout plans evolved over time. We digitized 22 versions of the rollout plans shared with us by the government over 2020 and 2021, allowing us to observe routes that were planned to be opened during this period but were suspended or delayed due to operational considerations.

Results. Table S17 shows the results for the base specification. The first three columns report results for departures. Column (1) repeats the baseline specification (column (7) in Table 2). Column (2) adds the event-time indicator that turns on when a route at Oyingbo would have opened, had it not been cancelled after the fire. Column (3) does the same for planned openings that were canceled due to changes in government deployment plans. Columns (4-6) repeat these specifications for fares, and columns (7-9) for queues.

We do not see evidence of effects after planned but canceled public transit entry. Across all three outcomes, effects are both economically small and statistically insignificant. Over all outcomes, four out of the 6 coefficients go in the opposite direction of the main effect on actually opened routes. These results support the notion that our main treatment effects are due to entry itself, and not differential trends on routes the government planned to enter.

In Appendix Table S18 we repeat the placebo regressions for the spillover specification, finding that neither type of canceled route yields phantom effects in the spillover specification.

S5.2. Departure Observations: Dynamics

To assess dynamics, we take the base regression specification (equation (11)) and replace the binary treatment indicator with indicators for whether public buses have been running for less than 3 months (associated coefficient β_1), 3 to 6 months (β_2), or longer than 6 months (β_3) on route r by survey round t. This results in estimating equation

$$\begin{aligned} Y_{rt\tau} = & \beta_1 \mathbb{I}\{\text{Open}_{rt} < 3 \text{ months}\} + \beta_2 \mathbb{I}\{\text{Open}_{rt} \text{ 3-6 months}\} + \\ & \beta_3 \mathbb{I}\{\text{Open}_{rt} > 6 \text{ months}\} + \gamma_{m(r\tau)} + \delta_{i(r)t} + \boldsymbol{\kappa}_t{}' X_{rt} + \epsilon_{rt\tau}. \end{aligned} \tag{S25}$$

Results are shown in Appendix Table S34, repeating the specifications of Table 2. There is evidence of a slight decline in departures within 3 months, which grows stronger and persists after 3 months, although the estimate loses precision under wild-bootstrap inference, as shown in Panel A. Fare effects are always negative but are imprecisely measured and do not show a clear time trend, shown in Panel B. Queues decline much more after 3 months, shown in Panel C. A delay between the adjustment of departures and queues would be consistent with drivers adjusting to reduced passenger flows after a delay. We present dynamic results visually in Appendix Figure S5, showing event studies for our main results using both two-way fixed effects and Sun and Abraham (2021) estimates.

Table S34. Effect of Public Transit on Private Transit: Dynamic

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: log(Departures)							
Open < 3 months	-0.079	-0.119	-0.113	-0.115	-0.117	-0.149	-0.084
	(0.039)	(0.001)	(0.016)	(0.021)	(0.043)	(0.040)	(0.275)
	[0.183]	[0.065]	[0.100]	[0.114]	[0.153]	[0.119]	[0.369]
Open 3-6 months	-0.133	-0.241	-0.237	-0.259	-0.252	-0.284	-0.287
	(0.055)	(0.004)	(0.021)	(0.008)	(0.015)	(0.006)	(0.022)
	[0.481]	[0.211]	[0.290]	[0.208]	[0.234]	[0.166]	[0.257]
Open >6 months	-0.152	-0.217	-0.218	-0.226	-0.225	-0.257	-0.355
	(0.256)	(0.046)	(0.065)	(0.066)	(0.047)	(0.099)	(0.003)
	[0.391]	[0.216]	[0.242]	[0.241]	[0.211]	[0.270]	[0.106]
Panel B: log(Fare)							
Open < 3 months	-0.045	-0.064	-0.047	-0.055	-0.053	-0.055	-0.017
	(0.302)	(0.251)	(0.394)	(0.310)	(0.343)	(0.287)	(0.697)
	[0.360]	[0.288]	[0.455]	[0.375]	[0.416]	[0.347]	[0.771]
Open 3-6 months	-0.019	-0.028	-0.016	-0.040	-0.049	-0.047	-0.014
	(0.637)	(0.552)	(0.760)	(0.397)	(0.299)	(0.175)	(0.649)
	[0.732]	[0.640]	[0.824]	[0.563]	[0.475]	[0.356]	[0.726]
Open >6 months	-0.022	-0.056	-0.047	-0.074	-0.082	-0.086	-0.047
	(0.810)	(0.544)	(0.615)	(0.429)	(0.362)	(0.280)	(0.535)
	[0.843]	[0.662]	[0.713]	[0.580]	[0.516]	[0.464]	[0.636]
Panel C: log(Queue)							
Open < 3 months	-0.064	-0.080	-0.063	-0.048	-0.042	-0.047	-0.001
	(0.593)	(0.583)	(0.606)	(0.694)	(0.742)	(0.648)	(0.993)
	[0.625]	[0.643]	[0.657]	[0.730]	[0.771]	[0.691]	[0.993]
Open 3-6 months	-0.442	-0.486	-0.459	-0.473	-0.459	-0.490	-0.514
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
	[0.021]	[0.060]	[0.013]	[0.015]	[0.017]	[0.011]	[0.019]
Open >6 months	-0.240	-0.374	-0.364	-0.393	-0.371	-0.400	-0.475
	(0.215)	(0.016)	(0.020)	(0.009)	(0.016)	(0.027)	(0.005)
	[0.410]	[0.181]	[0.177]	[0.131]	[0.132]	[0.151]	[0.094]
Route × Period FE	X	X	X	X	X	X	X
Day of Week × Survey Round FE	X	X	X	X	X	X	X
Hour of Dep × Survey Round FE	X	X	X	X	X	X	X
Terminal × Survey Round FE		X	X	X	X	X	X
Trip Dist Controls × Survey Round FE			X	X	X	X	X
Dep. Plan × Survey Round FE				X	X	X	X
CBD Controls					X		
O & D Lat-Lon Poly × Survey Round FE						X	
O & D LGA × Survey Round FE							X

Notes: P-values clustered by route and terminal are reported in parentheses. Wild p-values computed using Cameron, Gelbach, and Miller (2008) are reported in square brackets. N and R^2 rows omitted for brevity.

S5.3. Testing for the Presence of Higher-Order Spillovers

The spillover specifications in the paper remain valid only if our measure of connections captures all spillovers from treated to untreated routes. In other words, there should be no "higher-order" spillovers beyond interactions at origin or destination terminals.

To test this, we run an augmented version of our spillover regression

$$\begin{aligned} Y_{rt\tau} = \beta \mathbb{I}\{\text{Open}_{rt}\} + \alpha \text{ Connections to public transit}_{rt} \\ + \psi \text{Unconnected overlap with public transit}_{rt} + \gamma_{m(r\tau)} + \boldsymbol{\kappa}_t{}' X_{rt} + \epsilon_{rt\tau}. \end{aligned}$$

This specification includes a term that measures the fraction of each control route ("unconnected" with public transit at both endpoints) that overlaps with a public transit route at each survey date. To construct this measure, we intersect the shapefile of each minibus route with a 50-meter buffer around each public transit route that opens during our sample period. We then use the opening dates of each public transit route to calculate the fraction of each minibus route that overlaps a public route at each survey date. We construct this measure only for routes that are in the control group of the spillover specification, i.e. those which do not share any connection with public transit at either endpoint; for treated and connected routes it is set to zero.

Our goal is to test whether minibus routes in the control group are affected by the treatment. We do this in two ways. First, we test whether impacts on control routes are the same regardless of their overlap with the public system ($\psi = 0$). Second, we examine whether the treatment effect estimate (β) changes when this overlap measure is included. If treatment affects control routes through the overlap measure we construct, and this overlap influences outcomes in the control group, this SUTVA violation would introduce an omitted variable bias that would cause the estimate of β to change.

The results are presented in Appendix Table S35. Column (1) repeats the main spillover specification, column (2) adds controls for the fraction of each route overlapping with different road types (motorway, main, secondary) since new public routes are more likely to enter on busier roads, and column (3) includes the overlap variable for unconnected routes.

Table S35. Testing Higher-Order Spillovers through Public Transit Overlap on Control Routes

	(1)	(2)	(3)
Panel A: log(Departures)			
Open	-0.163*	-0.179*	-0.178*
	(0.090)	(0.101)	(0.106)
Number of connected routes	0.002	0.005	0.005
	(0.013)	(0.012)	(0.013)
Unconnected overlap with public bus route			0.023
			(0.310)
N	21,284	21,284	21,284
R^2	0.63	0.63	0.63
p-val Overlap==0			0.940
p-val Open (2)==(3)			0.939
Panel B: log(Fare)			
Open	-0.108**	-0.091*	-0.095**
	(0.052)	(0.047)	(0.046)
Number of connected routes	-0.016***	-0.014**	-0.014**
	(0.006)	(0.005)	(0.006)
Unconnected overlap with public bus route			-0.059
			(0.106)
N	23,067	23,067	23,067
R^2	0.94	0.94	0.94
p-val Overlap==0			0.580
p-val Open (2)==(3)			0.600
Panel C: log(Queues)			
Open	-0.163*	-0.203*	-0.207*
	(0.093)	(0.109)	(0.112)
Number of connected routes	0.029*	0.036*	0.035*
	(0.017)	(0.018)	(0.019)
Unconnected overlap with public bus route			-0.076
			(0.377)
N	22,290	22,290	22,290
R^2	0.55	0.55	0.55
p-val Overlap==0			0.841
p-val Open (2)==(3)			0.843
Road Type Intersections X Survey Round FE		X	X

Notes: Table reports regression from Appendix Section S5.3. Road Type Intersections measure the fraction of each minibus road that overlaps different road types (motorway, main, secondary). Column (3) adds the fraction of each control route that overlaps the public system at each survey round. P-val Overlap==0 tests whether the coefficient on this overlap variable is zero. P-val Open (2)==(3) tests whether the coefficient on Open is the same in columns 2 and 3. Standard errors clustered by route and terminal are reported in parentheses. * p<0.1; ** p<0.05; *** p<0.01.

First, we fail to reject $\psi = 0$ for each of the three outcomes. However, the estimates are relatively noisy, and for two outcomes—fares and queues—the point estimates are non-negligible compared to the main treatment effect. To further investigate, we turn to the second test: examining whether the coefficient on the treatment variable (Open_{rt}) changes when the overlap variable is included. We find that the coefficients are statistically indistinguishable between these specifications, and the point estimates are remarkably consistent across outcomes.

Taken together, we interpret this evidence to suggest that the primary effects from new public transit openings occur at endpoints. This supports the validity of our spillover specifications and suggests that their estimates do not seem influenced by possible violations of SUTVA.

Appendix S6. Price Sensitivity

This section estimates the substitution parameter θ by estimating trip responses to unexpected public bus fare changes. If public buses follow the same indirect utility structure as minibuses (equation (3)), the commuter model implies we can recover $\theta = -\frac{\partial \ln N_{rt}}{\partial p_{rt}} \frac{1}{\gamma \mathbb{E}[(1-s_{ijP})]}$ using the semi-elasticity of the number of trips to price (estimated here), combined with our field experiment's estimate of γ and the observed share $\mathbb{E}[(1 - s_{ijP})]$.

S6.1. Background on Price Changes

The government implemented several changes to public bus prices in 2023 and 2024, in response to pressure arising from rising fuel prices and macroeconomic conditions. On August 2, 2023, the government announced a 50% subsidy on public transit. Alongside this first price change, the government said it had arranged with minibus operators to reduce prices by 25%, but it has limited control over minibus fares so it is not clear whether minibus fares changed. We will treat this first price change separately from changes that followed. On November 7 the public transit subsidy was reduced to 25%.[59] On January 29, 2024, the public transit subsidy was eliminated, and on February 26, 2024, it was reinstated at 25%. This provides four changes in price for public buses, all of which were unexpected.

We illustrate the price changes by route in Figure S10. Each panel shows a different price change date, plotting the price two weeks prior (x-axis) and two weeks after (y-axis), with a point for each combination of route and entry point. The size of the point is scaled to the number of boardings over those days. The plots include two reference lines: one corresponding to the announced price change (dashed blue) and the 45° line corresponding to no price change (dotted red). The observed prices broadly coincide with the announced price changes; however in some cases, prices were rounded to

[59] On November 6, the public transit subsidy was eliminated for a single day. We will include fixed effects in our specifications for this particular day.

convenient currency multiples, and some smaller routes did not follow the overall price changes. As a result of these discrepancies, the first change in particular lowered prices by less than 50% on average.

Figure S10. Public Bus Price Changes

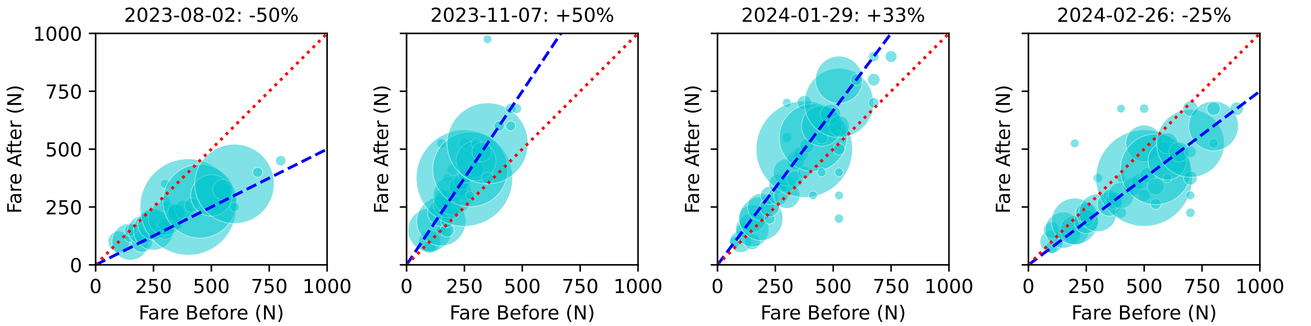

Note: Plots show the fares the two weeks before (x-axis) and after (y-axis) designated price changes on public bus routes. Each point represents a combination of an initial fare and final fare, with size scaled to the number of boardings on routes and boarding locations facing those corresponding fares during those days. Dashed blue line corresponds to the announced price change; dotted red line corresponds to no price change.

We next show descriptively the trends in prices, ridership, and number of public buses in Figure S11. The dates of the price changes are designated with dashed blue lines. The sharp changes in average fare paid are evident in the top panel. These changes coincide with changes in the total number of trips on public transit, as shown in the second panel. However, price is not the only factor affecting the number of trips. Some changes in trips coincide with changes in service, as shown in the final panel, which shows the number of unique buses used each week. We note that although service is relatively stable around price changes, there is a slight increase in buses following the first price change (which would reduce wait times for passengers). For this reason, in addition to the coincident announcement about minibus fares, we exclude the first price change from our preferred estimates, but report results for all changes. Our specifications will focus on relatively narrow windows around each price change to minimize the risk of other minibus fare and supply adjustments. Demand will end up responding quickly, suggesting that despite these narrow windows, we can still capture the full adjustment.

Figure S11. Public Transit Bus Trends 2023-2024

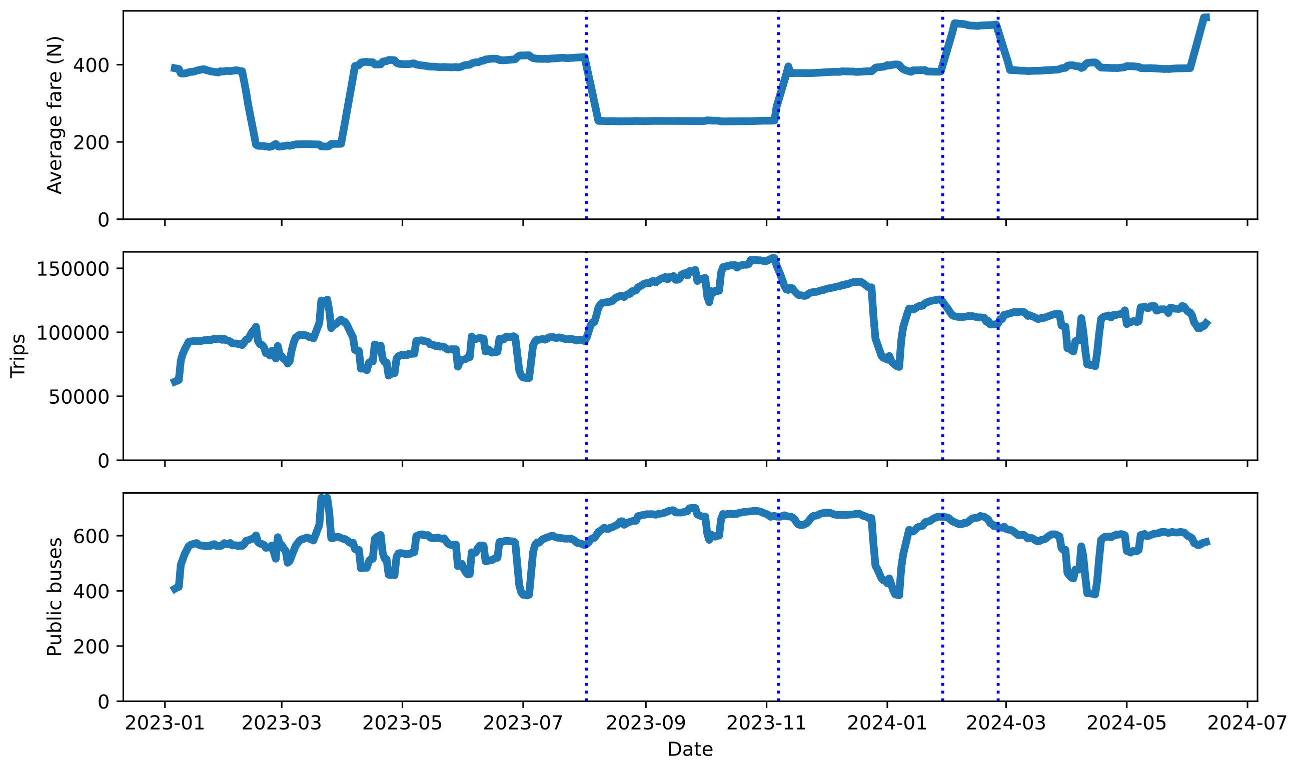

Note: Panels show the trend in average fares, total trips, and number of buses in the public bus system, per the electronic ticketing data, with a 7 day rolling average. Dotted lines indicate dates of fare changes. The major dropoff in travel and supply around 2023-12 to 2024-01 begins around Christmas.

Overall these results suggest an empirical strategy that focuses on the variation in ridership around the sharp fare changes, accounting for slight deviations in the implemented fare changes between routes.

S6.2. Estimation

We estimate the impact on prices around the discontinuities, and relate this to the corresponding impact on ridership. This allows us to account for the deviations in implementations. We assemble a dataset from e-ticketing at the route-day (rt) level, and include only observations within a designated number of days of a price change.[60]

Dynamics. To assess what time window is relevant for measuring the change, we first estimate the response of trips after different windows of time. We focus on the second price change (7 November 2023), because it is furthest from changes in the supply of buses and from other price changes. We jointly

[60]There are some route-days that have no boardings. Many of these appear to be idiosyncratic, suggesting they result from missing data. For that reason, we omit route-days that contain no boardings. We compute the mean price for each route-day, to account for minor variation in the recorded fare. In the log price specifications we omit a small number of observations that have a recorded price of zero.

estimate the system of seemingly unrelated linear regressions on price and log ridership,

$$p_{rt} = \boldsymbol{\zeta} Window_t + \phi_r + \psi_{DoW(t)} + \nu_{rt}$$

$$logN_{rt} = \boldsymbol{\alpha} Window_t + \mu_r + \xi_{DoW(t)} + \epsilon_{rt}$$

where $Window_t$ is a vector of indicator variables, which is 1 if t lies within a given window of days of the price change and 0 otherwise. We include fixed effects for route and day of week ($DoW(t)$). The price semi-elasticity for window k is then given by $\frac{\partial \ln N_{rt}}{\partial p_{rt}} = \frac{\alpha_k}{\zeta_k}$ (in log trips per ₦).

We also estimate a detrended version to account for possible trends in ridership. Using data from only the period before the price change, we estimate the trend equation

$$logN_{rt} = \beta \cdot t + \tilde{\mu}_r + \tilde{\xi}_{DoW(t)} + e_{rt}.$$

We then predict the ridership for all periods $\hat{logN}_{rt}$, and estimate the detrended regression

$$logN_{rt} - \hat{logN}_{rt} = \boldsymbol{\alpha} Window_t + \mu_r + \xi_{DoW(t)} + \epsilon_{rt}$$

We estimate standard errors by bootstrapping the entire joint estimation procedure, resampling routes with replacement. We include 28 days before the price change (before which there was a temporary supply shock) and 42 days following the second price change (which coincides with December 18), as travel demand drops off around Christmas (December 25).

Results of the two specifications are shown in corresponding columns of Table S36. In both specifications, we see that trips adjust quickly in response to the price change: the price adjustment in later periods is similar to that in the first 14 day period, and if anything lower. Based on this, we move forward with 14 day windows, which allows us to use variation from more price changes.

Table S36. Response to a Price Change in Public System: Dynamic

	(1)	(2)
Price Impact ≤14 days (₦: ζ_1)	91.0832***	91.0832***
	(4.925)	(5.5768)
Price Impact 15-28 days (₦: ζ_2)	89.6402***	89.6402***
	(5.1617)	(5.6865)
Price Impact 29-42 days (₦: ζ_3)	88.9993***	88.9993***
	(5.2662)	(5.7557)
Log Trip Impact ≤14 days (α_1)	-0.2327***	-0.2569***
	(0.0394)	(0.0549)
Log Trip Impact 15-28 days (α_2)	-0.2115***	-0.2534***
	(0.0367)	(0.0714)
Log Trip Impact 29-42 days (α_3)	-0.1632***	-0.223***
	(0.0317)	(0.0853)
Price Sensitivity ≤14 days (log trips/₦: $\frac{\alpha_1}{\zeta_1}$)	-0.0026***	-0.0028***
	(0.0005)	(0.0007)
Price Sensitivity 15-28 days (log trips/₦: $\frac{\alpha_2}{\zeta_2}$)	-0.0024***	-0.0028***
	(0.0005)	(0.0009)
Price Sensitivity 29-42 days (log trips/₦: $\frac{\alpha_3}{\zeta_3}$)	-0.0018***	-0.0025**
	(0.0004)	(0.001)
Trips detrended		X
N	13466	13466
System R^2	0.9402	0.9402

Notes: Fare and log trips effects estimated jointly in a seemingly unrelated regression around the 7 November 2023 price change, including route and day of week fixed effects. Ratio of effects computed from these individual estimates. Specifications also include an indicator for 6 November 2023 (the single day where the subsidy was entirely removed). In the detrended specification, a linear time trend in the trips measure is estimated on the pre-period; this trend is then removed from all periods. Standard errors in parentheses estimated via bootstrapping the entire procedure, resampling routes with replacement. * p<0.1; ** p<0.05; *** p<0.01.

Pooled estimates. We produce our main estimates with an instrumental variables specification. We estimate price sensitivity with a specification of the form

$$logN_{rt} = \beta^p p_{rt} + \mu_r + \boldsymbol{\lambda} AroundDiscontinuity_t + \epsilon_{rt}$$

We instrument for price using the first stage specification

$$p_{rt} = \boldsymbol{\zeta} AfterDiscontinuity_t + \phi_r + \tilde{\boldsymbol{\lambda}} AroundDiscontinuity_t + \nu_{rt}$$

where $AfterDiscontinuity_t$ is a vector of indicator variables for which the k'th is 1 after the k'th price change and zero otherwise, and $AroundDiscontinuity_t$ is a vector of indicator variables for which the k'th is 1 within 14 days of the k'th price change and zero otherwise. We include only observations within 14 days of a price change.[61]

Estimated price sensitivities are shown in Table S37. The first column shows the estimate pooled over all price changes. The second column shows the estimate pooled over all but the first price change, since that change coincides with a slight supply increase and the announced change on minibus prices. The last 4 columns report the estimates from each price change individually. Estimated sensitivities are fairly similar regardless of which price change is used (save for the last, which is small and insignificant). This grants confidence that the shifts in ridership are not driven by extraneous factors. We use the estimate from column (2) as our value of $\frac{\partial \ln N_{rt}}{\partial p_{rt}} = -\theta\gamma\mathbb{E}[(1 - s_{ijP})]$.

Table S37. Trips Price Sensitivity in Public System

	(1) All changes	(2) All but first	(3) 2023-08-02	(4) 2023-11-07	(5) 2024-01-29	(6) 2024-02-26
Fare	-0.0022***	-0.0017***	-0.0032***	-0.0025***	-0.0018***	-0.0004
	(0.0003)	(0.0003)	(0.0005)	(0.0004)	(0.0004)	(0.0004)
N	10199	7942	2257	2637	2677	2628
N_{routes}	158	148	120	129	132	130
Window (days)	14	14	14	14	14	14
First stage F stat	2501.9	2146.0	7100.1	2714.8	3781.9	3998.5
R^2	0.79	0.801	0.839	0.82	0.847	0.846

Notes: Instrumental variables estimates of effect of fare on log trips based on price changes on designated dates. Instruments are indicator variables for the periods between price changes, with sample restricted to 14 days from the nearest price change, plus an indicator for 6 November 2023 (the single day where the subsidy was removed). Regressions include fixed effects by route, day of week, and for being within the window of each price change. Bootstrapped standard errors in parentheses, resampling routes with replacement. * p<0.1; ** p<0.05; *** p<0.01.

As a robustness check, we verify that our estimate is not arising from a supply shift, in Appendix Table S38. We estimate the same specifications but with the outcome set to the log number of buses. The estimated effects are an order of magnitude smaller, have mixed signs, and the estimate on the 'All but first' specification is not statistically significantly different from zero. This suggests our procedure is not

[61] We also include an indicator for the date 2023-11-6 (the single day where the subsidy was entirely removed) in both stages, for regressions that span that date.

simply capturing shifts in supply.

Table S38. Does Supply Respond to Price Changes in Public System?

	(1) All changes	(2) All but first	(3) 2023-08-02	(4) 2023-11-07	(5) 2024-01-29	(6) 2024-02-26
Fare	-0.0002***	-0.0001	-0.0005***	-0.0006***	-0.0002	0.0009***
	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
N	10199	7942	2257	2637	2677	2628
N_{routes}	158	148	120	129	132	130
Window (days)	14	14	14	14	14	14
First stage F stat	2501.9	2146.0	7100.1	2714.8	3781.9	3998.5
R^2	0.869	0.881	0.919	0.918	0.927	0.917

Notes: Instrumental variables estimates of effect of fare on log buses based on price changes on designated dates. Instruments are indicator variables for the periods between price changes, with sample restricted to 14 days from the nearest price change, plus an indicator for 2023-11-6 (the single day where the subsidy was removed). Regressions include fixed effects by route, day of week, and for being within the window of each price change. Bootstrapped standard errors in parentheses, resampling routes with replacement. * $p<0.1$; ** $p<0.05$; *** $p<0.01$.

Appendix S7. Wait Time Experiment

This section provides more details on the wait time experiment described in Section 6.2.

S7.1. Implementation

Enumerators had bright yellow hats, and waited near bus stops. Figure S8 shows a participant checking in with one of our enumerators.

The design of the experiment had several features to improve performance under poor network connectivity, and to minimize the chance of fraud. The sequence of random numbers was drawn in advance and so only needed to be loaded once on the enumerator's phone (it did not require an active internet connection). Numbers were different from day to day so one could not text the previous day's code. The sequence was saved in a separate file and designed to be difficult to parse even if one viewed the source code. The website required logging in with a valid enumerator phone number, and all usage was logged to detect suspicious usage or sharing. We did not find evidence of suspicious usage.

All participants responded to a baseline survey during onboarding. We cross-randomized several aspects of the design:

- Bus stops were randomly selected to have the game proceed for either 3 or 5 weeks
- Each participant received a random checkin offer $s_n^{checkin}$ which was then held constant throughout the experiment

- Upon checking in at the bus stop, each participant received a wait offer draw $(s_{nt}, \Delta t_{nt})$ which differed by day:
  - Participants were randomly allocated to one of two onboarding experiences:
    * For one group, on the first checkin at the bus stop, participants received a reward of $s_{nt} =$₦2000 for zero wait ($\Delta t_{nt} = 0$). On their next two checkins, participants received a random draw $(s_{nt}, \Delta t_{nt})$ which was generous (high payment per wait).
    * For the other group, on the first *two* checkins at the bus stop, participants received a reward for zero wait ($\Delta t_{nt} = 0$; $s_{nt} =$₦2000 on the first checkin and ₦500 on the second). On their following checkin, participants received a random draw $(s_{nt}, \Delta t_{nt})$ which was generous (high payment per wait).
  - On each checkin after their third, all participants received their own random draw $(s_{nt}, \Delta t_{nt})$ from the stable distribution.
- Starting July 24, 2023, bus stops were randomly selected to either (1) transition to a different set of 'scheduled' wait time offers which was announced by text message, (2) be sent a reminder message about the game (to act as a placebo) and then continue as normal, or (3) continue as normal with no reminder message. We omit the scheduled user-days from the analysis in this paper.

Near the end of the game, we followed up with an endline survey. For some participants this took place during the final week of the game; for some it took place after the game had concluded.

S7.2. Deviations from Pre-Analysis Plan

This section describes deviations from our pre-analysis plan (PAP).

Compliance rates were lower than we expected, so we revised our design to (1) use a selection correction in our estimation, (2) record intentions to travel as well as actual check ins, and (3) randomize the checkin offer $s_n^{checkin}$ to act as an instrument for checkins that is excluded from the wait decision. We treated initial data as pilot data, and revised the PAP accordingly. When implementing the theory in the paper, we corrected a mistake in the error terms in the selection correction, and also simplified the model so that ν_{nt}^{wait} is unknown at the point of checking in.

We refined a few aspects of the experiment while it was under way. Our main specifications include only users exposed to the final design. Table S27 column (4) shows that results are similar if we include earlier observations while the design was changing.[62]

The experiment was occasionally affected by network downtime. If the network prevented the sending of text messages, we informed participants to resend the messages in the same sequence when network became available (which would result in them being paid out as normal, since our system

[62] On June 9, 2023, to increase takeup we increased the payment upon registration from ₦500 to ₦1000 and increased the first day offer amount from ₦500 to ₦2000. We also began randomly allocating participants to either receive one or two days of zero wait offers to measure anticipation. Early recruits were given the information about the distribution of offers (as described in footnote 35) via text message on June 16; recruits after that were provided this on a handout upon intake. Prior to June 18 we varied the distribution of offers as we were learning the range that attained variation in acceptance; from June 19 onward, we used a single stable distribution for our main treatment. Our main sample includes only participants who began on or after June 19.

validates waits based on the random codes, not on when they are sent). When systematic downtime affected many users, we followed a protocol to send apology text messages and transfer airtime once the network was up.

S7.3. Modeling Waiting in Main Commuter Model

For simplicity, our main commuter choice model in equation (1) focuses on the choice between modes. This section generalizes that model to allow for commuters to also make waiting choices at the bus stop. From this model, we derive the estimating equation used in the wait time experiment, equation (12), which conditions on mode choice.

Commuter model with timing. Each day, each commuter n engages in a two-step decision process. First, she observes her iid idiosyncratic preference for each mode ϵ_{ijmn} and selects a mode $m \in \mathcal{M}$ (which could be the outside option). If she takes a nontransit mode, the game ends. If she takes transit, there is one additional step. She leaves home, and arrives at the bus stop at an exogenous time t_0. After she is within the vicinity of the bus stop, she observes her idiosyncratic preference for starting to wait at different times, ν_{nt}, which are mean zero and iid, and decides when to start waiting $t = t^{start}$ from a finite set of times $\mathcal{T}^{start} \geq t_0$ for some $R \in \mathbb{N}$. (She may choose to delay starting, for example, to pick up a newspaper or an item from a vendor at the bus stop.) After she starts waiting, a bus arrival is realized after $\tilde{\omega}_{ijmt}$ minutes, charging price p_{ijmt} (both drawn iid). After boarding the bus, she obtains realized utility

$$\tilde{u}_{ijmnt} = \tilde{\alpha}_m - \gamma \tilde{p}_{ijmt} - \eta(w_{ijm} + t - t_0 + \tilde{\omega}_{ijmt} + t^T_{ij}) + \epsilon_{ijmn} + \nu_{nt}$$

where travel time is given by the sum of time spent getting to the bus stop (w_{ijm}), planned waiting at the bus stop ($t - t_0$), additional wait until the bus arrives ($\tilde{\omega}_{ijmt} \geq 0$), and in-vehicle travel time (t^T_{ij}).

Working backwards, during the second step she will select when to start waiting (t) to maximize

$$u_{ijmnt} = \tilde{\alpha}_m - \gamma p_{ijm} - \eta(w_{ijm} + t - t_0 + \omega_{ijm} + t^T_{ij}) + \epsilon_{ijmn} + \nu_{nt}$$

where $p_{ijm} = \mathbb{E}[\tilde{p}_{ijmt}]$ and $\omega_{ijm} = \mathbb{E}[\tilde{\omega}_{ijmt}]$.

During the first step, she anticipates that for minibus travel she will select the optimal time to start waiting, $t^{start*} = \arg\max_{t \in \mathcal{T}^{start}}[\nu_{nt} - \eta t]$, yielding travel time $t_{ijM} = w_{ijM} + \mathbb{E}\left[t^{start*}\right] - t_0 + \omega_{ijM} + t^T_{ij}$. She selects the mode m to maximize

$$u_{ijmn} = \alpha_m - \gamma p_{ijm} - \eta t_{ijm} + \epsilon_{ijmn}$$

where the constant term now differs, because for minibuses it includes the expected taste shock at the selected start time, $\alpha_M = \tilde{\alpha}_M + \mathbb{E}\left[\nu_{nt^{start*}}\right]$. This equation coincides with the aggregation of the problem used in the commuter mode choice problem, equation (1). Note that all parameters are identical except the constant, regardless of the distribution of ν_{nt}.

Wait time experiment. The wait experiment introduces a step 1.5. The commuter checks in with us at the bus stop at time t_0. We offer payment s_{nt} if the commuter delays the start and waits an additional

Δt_{nt} minutes.

Then in step 2 the commuter selects $t = t^{start}$ to maximize the incentivized utility

$$u^{WT}_{ijmnt} = \tilde{\alpha}_m - \gamma(p_{ijm} - s_{nt} \cdot 1_{t \geq t_0 + \Delta t_{nt}}) - \eta(w_{ijm} + t - t_0 + \omega_{ijm} + t^T_{ij}) + \epsilon_{ijmn} + \nu_{nt}$$

which includes a payment if the commuter selects $t \geq t_0 + \Delta t_{nt}$.

Then the decision to wait in equation (13) is defined by taking the difference $u^{WT}_{ijmnt_0+\Delta t_{nt}} - u^{WT}_{ijmnt_0}$. In that section, we assume that the differenced error $\nu^{wait}_{nt_0} = \nu_{nt_0+\Delta t_{nt}} - \nu_{nt_0}$ is normally distributed.[63]

The wait experiment conditions on mode choice, and also includes a decision of whether to check in, with associated error $\nu^{checkin}_{nt}$, modeling ν^{wait}_{nt} and $\nu^{checkin}_{nt}$ as jointly bivariate normally distributed.

S7.4. Likelihood

Let F^x indicate the marginal CDF along error x. Then the average log likelihood is given by

$$\begin{aligned}
&\bar{l}(\gamma, \eta, \sigma^{checkin}, \rho | \mathbf{C}, \mathbf{D}, \mathbf{s}^{checkin}, \mathbf{s}, \mathbf{\Delta t}) = \\
&\qquad \frac{1}{N} \sum_{nt} \Big[C_{nt} D_{nt} \log \Big(1 - F^{\nu^{wait}}(-\underline{D}_{nt}) - F^{\nu^{checkin}}(-\underline{C}_{nt}) + F(-\underline{D}_{nt}, -\underline{C}_{nt}) \Big) + \\
&\qquad\qquad C_{nt}(1 - D_{nt}) \log \Big(F^{\nu^{wait}}(-\underline{D}_{nt}) - F(-\underline{D}_{nt}, -\underline{C}_{nt}) \Big) + \\
&\qquad\qquad (1 - C_{nt}) \log F^{\nu^{checkin}}(-\underline{C}_{nt}) \Big]
\end{aligned} \tag{S26}$$

where participation thresholds are given by $\underline{D}_{nt}(\gamma, \eta) = \gamma s_{nt} - \eta \Delta t_{nt}$ and $\underline{C}_{nt}(\gamma, \eta, s^{checkin}_n) = \gamma s^{checkin}_n - \eta t^{checkin} + \mathbb{E}_{\mathcal{O}}[\gamma s - \eta \Delta t + \nu^{wait}_n | \gamma s - \eta \Delta t + \nu^{wait}_n > 0] \cdot \Pr(\gamma s - \eta \Delta t + \nu^{wait}_n > 0)$, the number of total observations is N.

S7.5. Additional Results

We present descriptives from our participants in Appendix Table S39. The rows are divided into three panels. The first panel presents characteristics measured in the baseline survey. The second panel shows settings and behavior in the game itself. The final panel shows responses in the endline.

[63] A normally distributed difference between shocks could be rationalized if the shocks themselves, ν_{nt}, were normally distributed, and commuters selected between two times to start waiting, $t = t_0$ and $t = t_0 + \Delta t_{nt}$.

Table S39. Wait Time Game Descriptives

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
	min	mean	std	median	max	count	corr(gameLength)	corr(zeroWaitDays)	corr(sCheckin)	corr(sWaitMean)	corr(tWaitMean)
Baseline											
Gender	0.0	0.47	0.5	0.0	1.0	640.0	-0.01	0.04	-0.06	0.03	-0.06
Age	18.0	36.88	12.75	35.0	78.0	640.0	0.01	-0.02	-0.02	0.04	-0.06
Education Category	1.0	5.12	0.94	5.0	7.0	640.0	0.01	-0.01	-0.07	-0.0	-0.01
Monthly Income (₦)	12500.0	61289.43	79331.78	37500.5	750000.0	640.0	0.06	-0.05	0.04	-0.02	0.01
Game											
Game Length (weeks)	3.0	3.46	0.67	3.0	5.0	640.0	1.0	0.01	-0.01	-0.28	0.16
Number of Initial Zero Wait Days	1.0	1.48	0.5	1.0	2.0	640.0	0.01	1.0	-0.06	0.06	-0.1
$s_n^{checkin}$ (₦)	200.0	550.94	290.05	400.0	1000.0	640.0	-0.01	-0.06	1.0	-0.04	-0.01
$\overline{s}_n$ (₦)	69.12	285.75	80.14	285.98	640.0	640.0	-0.28	0.06	-0.04	1.0	-0.11
$\overline{\Delta t}_n$ (min)	4.4	8.72	1.59	8.65	13.85	640.0	0.16	-0.1	-0.01	-0.11	1.0
$\overline{C}_n$	0.0	0.59	0.34	0.69	1.0	640.0	-0.13	-0.03	0.1	0.41	-0.23
$\overline{D}_n$	0.0	0.41	0.28	0.39	1.0	640.0	-0.21	-0.05	0.01	0.43	-0.23
Endline											
Expected $\overline{s}_n$ (₦)	10.0	1411.06	4959.46	1000.0	100000.0	575.0	-0.01	-0.08	-0.06	-0.1	0.05
Expected $\overline{\Delta t}_n$ (min)	1.0	10.63	7.85	10.0	59.0	575.0	-0.05	-0.08	-0.02	-0.03	0.03

Notes: Individuals selected an income category; we designate an individual's monthly income as the median income for the category chosen. 'corr' reports the correlation between the designated column and row variables. Game length (weeks) and the number of initial zero-wait days were randomly assigned. Variables include check in offer $s_n^{checkin}$, average wait offer payment $\overline{s}_n$, average wait offer duration $\overline{\Delta t_n}$ (in minutes), proportion of days checked in $\overline{C}_n$, and proportion of wait offers accepted $\overline{D}_n$. Endline asks about participants' expectations about wait offers (both payment $\overline{s}_n$ and duration $\overline{\Delta t_n}$).

Table S39 columns (1-6) report summary statistics of these measures across each individual to provide a description of our sample. The remaining columns (7-11) assess the success of our randomization. They report the correlation between the variables listed in the column and row, at the individual level. The correlation of randomly selected variables with baseline characteristics are all close to zero. Some of these variables are correlated with game outcomes downstream of treatment (for example, the checkin offer, $s_n^{checkin}$, is correlated with the proportion of days checking in, $\overline{C}_n$).

In the endline we asked about the wait time and offer they expected to obtain in the study.[64] The expected wait time (mean 10.6 minutes; median 10) lines up closely with the distribution we drew from (mean 10), as well as the empirical average (mean 8.7; median 8.7). However, participants appeared to struggle with answering the monetary offer; there were many outliers in these responses, and the mean participants reported is roughly 5 times the empirical mean, suggesting the question was not understood. (Participants may have answered how much they would have liked to be paid, rather than what they expected the system's offer would be.)

S7.6. Robustness

S7.6.1. Sample Definitions

We assess robustness to alternate sample definitions in Table S27. In column (2) we use survey data to remove observations for days in the first week that participants told us they did not plan to travel (based on the baseline), and days in the last week that participants told us they did not travel (based on the endline). Column (3) excludes the high offers from the first days of checking in. Column (4) includes observations during the pilot portion of the experiment where some design features were changing.[65] Estimates are similar in all cases.

S7.6.2. Arrival Time Adjustment

Another possible concern is that participants may anticipate a return to waiting, and so arrive at the bus stop early. If there were consistent evidence of this, our model could be adjusted to account for the scheduling cost. We assess the degree to which arrival times change in response to wait offers in Table S40 by regressing check in times at the bus stop (in minutes since midnight) on covariates. We perform three types of comparisons. Column (1) compares the time of the first checkin (omitted category) to later checkins, broken down by offer type, over the entire sample. Offer types include the no-wait second checkin (randomly assigned), the generous offers in the third and second checkin (for those not assigned a no-wait second checkin), and the stable distribution which was offered from the fourth checkin onward. For all of these offers, estimates are similar: participants check in on average 10.3 to 11.7 minutes earlier for the first checkin. This effect is the opposite one would expect if they solely factored in the wait time: they may have arrived earlier on the first check in to have enough time to resolve uncertainty about the experimental setup. Column (2) adds to this specification a control for days elapsed. If it takes time for participants to adjust their schedules to arrive earlier, this coefficient

[64] We asked, 'How long did you think we would ask you to wait at the bus stop before taking the bus?' and 'How much did you think we would offer you to wait that long?', respectively.

[65] Described in Section S7.2.

would be negative. However, it is a small and noisy zero. Column (3) restricts the sample to only the second checkin, for which participants were told that they would receive a payment either with or without a wait (randomly assigned). Here we see some evidence of adjustment: those assigned the no-wait offer arrive on average 6 minutes later; however the estimate is extremely noisy: despite having 535 observations we fail to reject a difference with zero (p-value: 0.50). Altogether, we do not see consistent evidence of schedule adjustment in response to the wait time experiment, and so have currently left that out of the model.

Table S40. Wait Experiment Robustness: Check in Time

	(1)	(2)	(3)
Offer: Second checkin no wait	10.424*	10.286*	6.149
	(5.732)	(5.758)	(9.046)
Offer: Second/third checkin generous	11.712***	11.489***	
	(4.175)	(4.266)	
Offer: Stable distribution	10.305***	9.330*	
	(3.920)	(4.786)	
Day		0.120	
		(0.343)	
Sample	All checkins	All checkins	Second checkin
N	5494	5494	535
R^2	0.609	0.609	0.038
Day of Week FE	X	X	X
Individual FE	X	X	
Start Date FE			X

Notes: The outcome is check in time, in minutes since midnight. In the first two columns, the omitted category is the first day offer with no wait. In the third column, the omitted category is the generous offer. The third column omits 11 cases who received the stable distribution rather than the generous distribution on their second day. * p<0.1; ** p<0.05; *** p<0.01.

S7.6.3. Arrival Dependent Waiting

We asked our enumerators to stand away from bus stops, not in direct line of sight. However, it is possible that participants monitored bus arrivals, and decided whether to accept our wait offer depending on whether a bus arrived in the waiting time interval. As a robustness check we estimate a likelihood for which waiting depends on arrivals, using data we gathered on the frequency of minibus departures.

Table S41. Wait Experiment: Arrival Dependent Waiting

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Main	λ=150.0 dep/30 min (Headway Q1)	λ=23.7 dep/30 min (Headway Q10)	λ=6.7 dep/30 min (Headway Q50)	λ=2.1 dep/30 min (Headway Q90)	λ=1.5 dep/30 min (Headway Q95)	λ=0.8 dep/30 min (Headway Q99)	λ=0.1 dep/30 min
γ (utils/₦)	0.0025***	0.0025***	0.0025***	0.0026***	0.0028***	0.003***	0.005	0.0037***
	(0.0002)	(0.0007)	(0.0007)	(0.0006)	(0.0006)	(0.0007)	(0.0042)	(0.001)
η (utils/min)	0.0464***	0.0464***	0.0462***	0.0444***	0.0559***	0.0702***	0.2401	1.3901
	(0.0031)	(0.0032)	(0.0032)	(0.0034)	(0.0064)	(0.0126)	(0.2719)	(1.1064)
$\frac{\eta}{\gamma}$ (₦/min)	18.9357***	18.9357***	18.7826***	17.3052***	20.0311***	23.6811***	47.8999***	373.1481*
	(1.7123)	(4.8738)	(4.7944)	(3.8734)	(3.1467)	(3.0303)	(14.8144)	(202.1312)
$\sigma^{checkin}$	14.3662***	14.3662	14.4355	15.4013	17.0815	18.0341	28.6195	12.4344
	(3.901)	(37.0444)	(37.3864)	(41.0838)	(44.2656)	(49.8159)	(162.9444)	(15.6073)
ρ	0.5163***	0.5163**	0.5082**	0.4078**	0.279*	0.2547*	0.0766	-0.1498
	(0.0588)	(0.2162)	(0.2151)	(0.1953)	(0.1537)	(0.1507)	(0.4023)	(0.1376)
N	8640	8640	8640	8640	8640	8640	8640	8640
Log Likelihood	-8720.18	-8720.18	-8719.66	-8722.11	-8725.67	-8722.87	-8688.93	-8769.74

Notes: In columns after the first, participants also accept the wait offer if no bus arrives during the wait time. We simulate bus departures following a Poisson process with average rate λ departures per half hour. Standard errors clustered at the user level. Estimation fixes $\sigma^{wait} = 1$. * p<0.1; ** p<0.05; *** p<0.01.

Let $\Pr(A)$ be the probability that a bus arrives within a timespan of Δt_{nt} minutes. We assume that the number of buses within that period, X, follows a Poisson arrival process, so that the probability of at least one arrival is given by $\Pr(A) = \Pr(X \geq 1) = 1 - e^{-\Delta t_{nt} \lambda}$ where λ is the average rate at which buses arrive.

Now let's refine the effect of the wait offer on travel time, building on equation (12). Let t_m be the status quo travel time and $t_m(\Delta t_{nt})$ be the departure time if the person accepts a wait offer of Δt_{nt} minutes, by checking out. Conditional on checking in, the utility of accepting the offer has two cases. If no bus arrives within the interval Δt_{nt} (case $N = 1 - A$), then the offer entails no additional wait and the person will check out as long as

$$\gamma s_{nt} > \epsilon_{mnt} - \epsilon_{mnt+\Delta t_{nt}}.$$

On the other hand, if a bus arrives in that interval (case A) then the person will skip the bus to accept the offer only if they value the payment more than the wait. How long is the additional wait? Since the Poisson arrival process is memoryless, $\mathbb{E}[t_m(\Delta t_{nt}) - t_m] = \Delta t_{nt}$. This leads to the original condition,

$$\tilde{\alpha}_m - \gamma(p_m - s_{nt}) - \eta(t_m + \Delta t_{nt}) + \epsilon_{mnt+\Delta t_{nt}} > \tilde{\alpha}_m - \gamma p_m - \eta t_m + \epsilon_{mnt}.$$

Now, decompose the idiosyncratic error into the mode component and time component, $\epsilon_{mnt} = \epsilon_{mn} + \nu_{nt}$. Then the consumer will wait in the two cases if

$$D_{nt}^{*N} = \gamma s_{nt} + \nu_{nt}^{wait} > 0$$
$$D_{nt}^{*A} = \gamma s_{nt} - \eta \Delta t_{nt} + \nu_{nt}^{wait} > 0.$$

where we define $\nu_{nt}^{wait} = \nu_{nt+\Delta t_{nt}} - \nu_{nt}$.

Perfect Compliance. For now, assume that compliance is perfect so participants always check in. Then, if there is no arrival then the probability of waiting is given by $1 - F^{\nu^{wait}}(-\underline{D}_{nt}^{N})$ for threshold $\underline{D}_{nt}^{N}(\gamma) = \gamma s_{nt}$. If there is an arrival, the probability of waiting is given by $1 - F^{\nu^{wait}}(-\underline{D}_{nt}^{A})$, for $\underline{D}_{nt}^{A}(\gamma, \eta) = \gamma s_{nt} - \eta \Delta t_{nt}$. Combining these and the probability of arrival yields the average log likelihood

$$\begin{aligned} \bar{l}(\gamma, \eta, \sigma^{wait}, \rho | \mathbf{D}, \mathbf{s}, \boldsymbol{\Delta}\mathbf{t}) = \frac{1}{N} & \\ \sum_{nt} D_{nt} \log \Big(1 - \Pr(1-A) F^{\nu^{wait}}(-\underline{D}_{nt}^{N}) - \Pr(A) F^{\nu^{wait}}(-\underline{D}_{nt}^{A})\Big) + & \\ (1 - D_{nt}) \log \Big(\Pr(1-A) F^{\nu^{wait}}(-\underline{D}_{nt}^{N}) + \Pr(A) F^{\nu^{wait}}(-\underline{D}_{nt}^{A})\Big) & \end{aligned} \tag{S27}$$

Note that when $A = 1$ this corresponds to the main model with perfect compliance. As A decreases, that increases the likelihood of accepting the offer ($D_{nt} = 1$) since the threshold for accepting the offer is lower since it involves no wait ($\underline{D}_{nt}^{A} \leq \underline{D}_{nt}^{N}$ so $F^{\nu^{wait}}(-\underline{D}_{nt}^{A}) \geq F^{\nu^{wait}}(-\underline{D}_{nt}^{N})$).

Imperfect Compliance. Now, let's consider compliance. The utility of checking in is given by

$$
\begin{aligned}
C^*_{nt} = \gamma s_n^{checkin} - \eta t^{checkin} \qquad & + \Pr(1-A)\mathbb{E}_{\mathcal{O}}[D^{*N}_{nt} | D^{*N}_{nt} > 0] \Pr(D^{*N}_{nt} > 0) \\
& + \Pr(A)\mathbb{E}_{\mathcal{O}}[D^{*A}_{nt} | D^{*A}_{nt} > 0] \Pr(D^{*A}_{nt} > 0) \\
& + \nu_{nt}^{checkin}
\end{aligned}
$$

The average log likelihood is given by

$$
\begin{aligned}
\bar{l}(\gamma, \eta, \sigma^{wait}, \sigma^{checkin}, \rho | \mathbf{C}, \mathbf{D}, \mathbf{s}^{checkin}, \mathbf{s}, \Delta\mathbf{t}) = \frac{1}{N} \\
\sum_{nt} C_{nt} D_{nt} \log \Big(1 - F^{\nu^{checkin}}(-\underline{C}_{nt}) + \Pr(1-A) \left[-F^{\nu^{wait}}(-\underline{D}^N_{nt}) + F(-\underline{D}^N_{nt}, -\underline{C}_{nt}) \right] \\
+ \Pr(A) \left[-F^{\nu^{wait}}(-\underline{D}^A_{nt}) + F(-\underline{D}^A_{nt}, -\underline{C}_{nt}) \right] \Big) + \\
C_{nt}(1 - D_{nt}) \log \Big(\Pr(1-A) \left[F^{\nu^{wait}}(-\underline{D}^N_{nt}) - F(-\underline{D}^N_{nt}, -\underline{C}_{nt}) \right] \\
+ \Pr(A) \left[F^{\nu^{wait}}(-\underline{D}^A_{nt}) - F(-\underline{D}^A_{nt}, -\underline{C}_{nt}) \right] \Big) + \\
(1 - C_{nt}) \log F^{\nu^{checkin}}(-\underline{C}_{nt})
\end{aligned}
\tag{S28}
$$

where participation thresholds are given by $\underline{D}^N_{nt}(\gamma, \eta) = \gamma s_{nt}$, $\underline{D}^A_{nt}(\gamma, \eta) = \gamma s_{nt} - \eta \Delta t_{nt}$ and $\underline{C}_{nt}(\gamma, \eta, s_n^{checkin}) = \gamma s_n^{checkin} - \eta t^{checkin} + \Pr(1-A)\mathbb{E}_{\mathcal{O}}[\gamma s_{nt} + \nu^{wait}_{nt} | \gamma s_{nt} + \nu^{wait}_{nt} > 0] \cdot \Pr(\gamma s_{nt} + \nu^{wait}_{nt} > 0) + \Pr(A)\mathbb{E}_{\mathcal{O}}[\gamma s_{nt} - \eta \Delta t_{nt} + \nu^{wait}_{nt} | \gamma s_{nt} - \eta \Delta t_{nt} + \nu^{wait}_{nt} > 0] \cdot \Pr(\gamma s_{nt} - \eta \Delta t_{nt} + \nu^{wait}_{nt} > 0)$, the number of total observations is N, and bold symbols represent vectors. Note that when $A = 1$ this corresponds to the main model. We estimate parameters using maximum likelihood, and cluster standard errors at the commuter level.

Results are shown in Table S41 for different bus arrival rates λ. Estimates are very similar even under extreme arrival rates. We trace out estimates that result under the arrival rates corresponding to various quantiles of realized headway between individual minibuses observed in departure observations. Note that these are observations of headways between individual minibuses, not average headways, so will have a more extreme distribution than any average λ commuters would expect. When buses are very frequent, the estimates coincide with the main estimate; they remain very similar up to the 90th percentile of observed individual headways which corresponds to 2.1 arrivals per half hour. The implied value of time only starts to deviate substantially in the extremes. It is 25% higher at the 95th percentile of headway, and at the 99th percentile of headway (entailing 0.8 buses per half hour), this model would expect commuters to accept almost all wait offers; the fact that they reject some would then suggest a high value of time ($\frac{\eta}{\gamma} = 48$ ₦/min). The implied value of time only begins to differ dramatically for very extreme values of λ (0.1 buses per half hour, corresponding to a departure every 5 hours, leads to a more extreme estimate of $\frac{\eta}{\gamma} = 373$ ₦/min). Since average headways are not this extreme, we interpret the results to suggest that arrival dependent waiting is unlikely to have enormous effects on our estimates of the value of time.