

Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta *

Gabriel Kreindler, Arya Gaduh, Tilman Graff, Rema Hanna, Benjamin A. Olken[†]

June 9, 2023

Abstract

Designing public transport networks involves tradeoffs between coverage, service frequency, and direct service. We use the expansion of the bus system in Jakarta, Indonesia, to study these tradeoffs. We analyze how new direct connections, changes in bus travel time, and wait time reductions affect bus ridership and aggregate flows, and estimate a transit network demand model by matching the route launch events. Commuters in Jakarta are 2-4 times more sensitive to wait time than bus time, and inattentive to long routes. We develop a flexible framework to characterize optimal networks. A less concentrated network would increase ridership and commuter welfare.

*We are grateful to TransJakarta for sharing their data. We thank David Atkin, Emily Breza, Dave Donaldson, Allan Hsiao, Myrto Kalouptsi, Natash Pillai, Felix Tintelnot, and Aaron Smith for useful conversations, Nick Tsivanidis and Matthew Turner for useful discussions of this paper, and participants at the NBER Summer Institute urban meeting and seminar participants for useful feedback. Lolita Moorena, Nadia Rayhanna, Nadia R Avianti, and Nikhil Kumar provided exceptional research assistance. We thank Paolo Adajar, Aaron Berman, Robert Dulin, Laya Gollapudi, Deivy Houeix, Helen Hu, Drew Johnston, Moulinrouge F Kaspar, Ariba Khan, Van Anh Le, Lauren Li, Anna Lingyi Ma, Kalyan Palepu, Jackson C Phifer, Xiaocun Qiu, Margaret Sands, San Singh, Yulu Tang and Hoi Wai Yu for research assistant work at different stages of the project. We are grateful to Kyle Schindl, Analysia Watley, Aditi Chitkara, and Nikhil Kumar for their help in developing the bus GPS headway engine. The project has IRB approval from MIT COUHES, protocol 1810556772. This material is based upon work supported by the National Science Foundation under Grant No. 2049784, by the Australian Government, the International Growth Centre, the Harvard Asia Center, and the Harvard Data Science Initiative.

[†]Kreindler: Harvard University (gkreindler@g.harvard.edu), Gaduh: University of Arkansas (agaduh@walton.uark.edu), Graff: Harvard University (tgraff@g.harvard.edu), Hanna: Harvard Kennedy School (rema_hanna@hks.harvard.edu), Olken: MIT (bolken@mit.edu).

1 Introduction

Governments are investing heavily in public transport in the megacities of low- and middle-income countries (LMIC), giving city governments a growing role in planning and operations. For example, the number of LMIC cities with a Bus Rapid Transit (BRT) system went from 14 cities in 2000 to 100 in 2020 (BRT+ Centre of Excellence and EMBARQ, 2023). Bus networks are particularly important because they account for a large share of ridership. For example, Delhi’s public bus system carried 3.3 million passengers per day in 2017, much more than its more famous metro system (2.7 million, Hindustan Times 2018). Ensuring that these transport systems are well-designed is key to sustainable and efficient mobility.

Designing bus networks involves tradeoffs. Because the fixed cost of a bus route is relatively small (compared to, say, building a subway or light-rail line), there is great latitude in designing a system. For a given number of buses, one can choose between having a more *direct* network (i.e., more direct routes, rather than a hub-and-spoke system that requires more changes between routes); a more *intensive* route network (i.e., more frequent service to key locations); and a more *expansive* route network (i.e., many routes serving more destinations).

Fundamental to these tradeoffs is an understanding of public transport preferences and demand, which in turn affects bus ridership: how commuters value bus travel time, wait times in stations, and transfers between routes. Different preference parameter configurations may result in very different-looking optimal transportation networks. Estimating these preferences is challenging, however, because it requires plausibly exogenous variation in *system-level* attributes.

In this paper, we study the expansion of the TransJakarta bus system in Jakarta, Indonesia. TransJakarta operates a public Bus Rapid Transit (BRT) network, which is the largest such network in the world (Institute for Transportation and Development Policy, 2019). From 2016-2020, TransJakarta launched a total of 93 routes across the city in a staggered fashion for both BRT routes and non-BRT (feeder) routes. These new routes changed the attractiveness of bus travel in different ways for different potential journeys, which we use to estimate preferences and study the implication for network design.

Our paper proceeds in three steps. First, we estimate the reduced form impact on bus ridership and overall commuting flows from improvements to the TransJakarta network. To do so, we show how the 93 new bus routes create different types of improvements in the availability of direct connections, travel times, and wait times for different origin-destination pairs, and analyze these ‘events’ using a differences-in-differences framework. Second, we estimate a model of commuter travel behavior, which allows us to back out the underlying

preference parameters that match the reduced-form estimates from step 1. Third, we simulate what optimal transportation networks would look like given the preferences from step 2, and contrast their characteristics with that of the existing network.

For the first step, we define three types of ‘events’ that correspond to different types of service improvement for origin-destination (o, d) pairs. The first two events focus on new direct routes for (o, d) pairs that were previously served only via transfer connection. The first event type (hereafter, Event 1) considers cases where the total time on the bus is unchanged with the new direct route. This introduces variation in the degree to which the network is *direct*. Event 2 considers new direct routes where the total time on the bus also falls, creating additional variation in service *speed* in addition to becoming more direct. Event 3 considers (o, d) pairs that receive a new direct route that overlaps with the old direct route(s) from o to d . This leads to more frequent service, varying the degree to which service is *intensive*, without substantially changing time spent on the bus or the directness of the route.

We use several new datasets to analyze these events. To measure (o, d) ridership flows on the bus network, we use TransJakarta’s detailed administrative data for both BRT routes and non-BRT feeder routes. TransJakarta collects fares using a tap-to-pay smart card system and our data includes every tap — over 500 million taps in the study period — along with an anonymized card identifier.¹ Second, we measure bus travel times and the distribution of wait times that passengers face using GPS data that tracks the position of every TransJakarta bus every 5-10 seconds. Third, we measure overall commuting flows on each (o, d) pair (regardless of transportation mode) using anonymized smartphone location data. We aggregate the outcome data at the weekly level and at the level of 1km-wide hexagonal grids, and show robustness to other grid shapes and sizes.

Using a unified difference-in-differences design, we find that bus ridership is responsive to each of these three types of service quality improvements.² For BRT routes, a new direct route between an existing (o, d) pair increases ridership by 0.16 log points. Event 2, which substantially reduced travel times in addition to a new direct route, increased ridership by 0.27 log points. Event 3, which reduced wait times by 0.32 log points (through adding busses via an incidental new route), increased ridership by about 0.09 log points, implying that ridership also responds substantially to wait times. Event-study graphs show clear jumps in ridership immediately after the launch of new routes, and no pre-trends. The results for

¹Since the system enforces tap-in everywhere in the network, but only uses tap-out at a subset of BRT stations, we use an algorithm to infer each trip’s likely destination. We validate this approach using the subset of trips for which we directly observe both the origin and destination.

²We include a large set of “never-treated” origin-destination pairs, which largely alleviates the class of problems with two-way fixed effects specifications (see, e.g., [De Chaisemartin and D’Haultfoeuille 2022](#)).

non-BRT routes are typically larger in magnitude in proportional terms and similar in levels than the BRT routes. Our results are stable when varying grid cell size from 500 meters to 2 kilometers, implying that the events lead to new bus ridership rather than displacing passengers from neighboring TransJakarta stations.

We next look at the impact of the same set of events on the aggregate volume of trips between a pair of locations (regardless of whether people take the bus or an alternative), measured using the smartphone location data. We do not find observable increases in all trips for any of the events in our benchmark specification. In fact, we typically reject moderate positive effects. In short, the results imply substantial increases in ridership when service frequency improves on all three dimensions — directness, speed, and wait times — and with this being driven by substitution towards bus trips, rather than by completely new journeys.³

With our reduced form estimates in hand, in the second step of our paper, we develop a model of commuter travel behavior to estimate underlying preference parameters. The model has a new formulation of how commuters choose routes within a geographically realistic transit network, in the presence of stochastic wait times. A commuter traveling from o to d chooses between direct and transfer bus options that differ in terms of type (BRT or non-BRT), total travel time on the bus, and the necessity of a transfer. Buses on each route arrive according to a Poisson process. A commuter who decides to take the bus gets a draw of wait times for all routes in her choice set and selects the best option overall. Because the idiosyncratic term in the model comes from random wait times given by Poisson processes, the model is invariant to aggregating identical routes, a property that allows us to bypass spurious gains from variety when we return to the optimal network design problem.⁴

Importantly, commuters sometimes face large choice sets with dozens of options within TransJakarta’s network, many of which are too slow to consider. Therefore, we allow for (and estimate the degree of) partial inattention to bus route options that have high travel time on the bus, relative to the quickest bus option in the choice set (Larcom et al., 2017)

In a higher nest, commuters choose between using the bus network, taking expectations over the wait times they may face, or a private transportation outside option (e.g. motorbike, car), based on a logit specification.⁵

³Converting road traffic lanes to dedicated bus lanes can increase road traffic congestion (Gaduh et al., 2022), yet the TransJakarta network expansion we study here did not reduce road traffic lanes. The only major change to BRT infrastructure in this period was the launch of corridor 13, which runs along an elevated busway so did not reduce road traffic lanes.

⁴In typical discrete choice models such as logit, splitting a bus route into two identical routes with half the original bus allocation is the source of an expected utility gain because the new routes have independent idiosyncratic errors; this is the famous ‘red-bus, blue-bus’ problem discussed by McFadden (1974). We show that the general version of our model, which appears in Daganzo (1979) under the name “negative exponential distribution model,” does not have this issue.

⁵While it is possible to further embed the model into a destination-choice nest or in an urban general

We estimate the model for the discretized Jakarta metropolitan area using 1km hexagonal grid cells. We use smartphone data to estimate commuting flows across all possible combinations of these grid cells. We then use the model to compute predicted ridership over each (o, d) pair and examine how it changes over time as the bus network changes, by running the same reduced form analysis as in the actual data. This allows us to estimate the preference and attention parameters that match the actual event-study estimates.

Using this approach, we find a high disutility for waiting time relative to time traveling on the bus. Specifically, we estimate that travelers find wait time $2.4\times$ more costly than travel times for BRT, and $4.2\times$ most costly for non-BRT. Commuters dislike bus options that involve transfers due to the additional wait time and typically longer time on the bus. However, there is no additional transfer penalty, above and beyond the wait time and time on the bus costs. Finally, we find that commuters pay significantly less attention to bus options with bus travel time more than 34%-44% longer than the quickest option in their choice set.

What do these preferences mean for the optimal design of the network? In the third part of the paper, we introduce a theoretical and computational framework that allows us to characterize optimal networks, compare them with the current network, and perform comparative statics with respect to structural model parameters.

We ask how the network of bus routes should be configured across the city and how the existing set of buses should be allocated to these routes to maximize a utilitarian welfare objective informed by our estimated demand model. This involves balancing tradeoffs across reach (having enough bus stops to cover origins and destinations for which people would like to travel), wait times (putting more buses on a given route), and the topology of the network (which stations should be connected directly versus by transfer). We hold fixed TransJakarta’s current BRT infrastructure, which allows for faster travel times along the 13 BRT corridors, but we allow routes to enter and exit the BRT corridors, to connect them in new ways, and to travel on existing streets. We study this problem using realistic geography of Jakarta at the levels of 2km-wide square grids.

Our problem displays divergent substitution and complementarity forces at the same time, hence violating the type of complementarity property that enables the characterization of optimal allocations in other settings (Jia, 2008; Arkolakis and Eckert, 2021). This is also a very high-dimensional problem, and the number of possible configurations is extremely large, making it difficult to identify the (generically unique) welfare-maximizing network.⁶

equilibrium model (Tsivanidis, 2022), the null reduced form effect on aggregate trips suggests that over the period we study, the preferences we estimate over wait time, transfers, and travel time are not biased by changes in aggregate trip patterns.

⁶With 418 grid cells and 1,536 possible edges, there are $2^{1,536}$ configurations where each edge is either

Therefore, we set up a slightly modified problem: the social planner chooses a network N that maximizes utilitarian welfare $W(N, \theta)$ from our estimated demand model — average expected utility for all commuters — plus an idiosyncratic shock ϵ_N for each possible network. The shocks capture network-level factors that the social planner cares about, such as cost shocks or benefits from certain networks, outside of our demand model. As such, in our framework, the social planner has a discrete choice problem over the space of all possible networks. A network with high model welfare carries a large probability of being chosen, because small idiosyncratic shocks may tip the planner into preferring it. We parametrize the problem as a multinomial logit with parameter β .

This approach allows us to transform a global optimization problem into a problem of sampling from a large-scale multinomial logit distribution. Several algorithms based on the Metropolis-Hastings algorithm have theoretical guarantees and good practical performance in sampling from such distributions.⁷ We derive formulas for properties of optimal networks and comparative statics. We can estimate these objects using their sample analogs. Our framework is general and may be used in other settings where a social planner’s (or an agent’s) problem is high-dimensional.⁸

Using this approach, we find that optimal networks, based on the preference parameters that we estimated, are substantially more extensive than the actual network. For example, we estimate that a typical optimal network should cover about 66 percent of all grid cells, compared to 42 percent with the actual TransJakarta network. 91% of all trips by Jakarta residents would have access to the bus network, compared to 73% in the current network.⁹ Bus frequency in the city center is slightly lower in the optimal network, so commuters in those areas need to wait longer for a bus than in the current network. While our estimated parameters indicate that commuters are highly sensitive to wait times, we find that the current TransJakarta network nevertheless concentrates too many buses in the city center, where wait times are already very short.

In our last exercise, we explore how certain characteristics of the estimated optimal networks would change if preferences differed. For example, doubling the wait time cost leads to more concentrated networks connecting only 23 percent of origin-destination pairs,

connected or not – i.e. about 10^{500} possibilities. This is a lower bound of just the route design problem, and does not even account for the number of buses on each route.

⁷We use a version of the simulated annealing (SA) algorithm, which we run multiple times to obtain a sample of networks drawn (asymptotically) from the planner’s distribution.

⁸While the details of the implementation of the sampling algorithm will be application-specific, the theoretical framework can be applied to problems such as optimal spatial policy (Fajgelbaum and Gaubert, 2020), optimal transportation infrastructure Allen and Arkolakis (2022); Fajgelbaum and Schaal (2020), and firms’ global sourcing decisions Antràs et al. (2017). The setting is appropriate when the welfare difference (or ratios) can be computed for pairs of counterfactuals, for example, if exact hat methods are applicable.

⁹Our results do not suggest a tension between distributional and efficiency concerns in this setting.

41% less than before. Increasing the transfer penalty by the equivalent of 15 minutes of wait time increases the share of connected location pairs that have a direct connection from 12% to 16%. This type of comparative static exercise connects micro parameters in the model to macro properties of optimal bus networks, which helps uncover which parameters matter for network design and to assess sensitivity of our results to the parameters to estimate.

Our project connects to several literatures. First, we build on classic questions in the transportation economics literature: travel demand estimation and mode choice in particular (McFadden, 1974; Ben-Akiva et al., 1985), and increasing returns in public transport (Mohring, 1972).¹⁰ We build on the literature investigating how commuters value trip attributes, see e.g. Wardman (2004); Currie (2005); Abrantes and Wardman (2011) and Small et al. (2007, page 53). Studies in this literature tend to focus on high-income countries and often infer stated preferences from respondents' choices between hypothetical alternatives. In contrast, we estimate preferences using natural experiments. In that respect, our study is closest to Kreindler (2023), which uses a field experiment to estimate the underlying preference parameters required to calculate optimal congestion pricing in Bangalore, India. Here, by contrast, we use a variety of real-world changes in the attractiveness of the network given by different types of network expansion for identification. Our study is also related to Buchholz et al. (2020) and Goldszmidt et al. (2020), both of which estimate the value of wait time for taxis and ride-hail from choices made on ride-hail platforms in Europe and the United States.¹¹

Second, by embedding travel demand preferences into a model of optimal bus route network design, our project also contributes to the growing trade-inspired literature that tackles questions of how to design a transport network (Fajgelbaum and Schaal, 2020; Allen and Arkolakis, 2022; Balboni, 2021; Santamaria, 2022; Alder, 2023). This literature has so far focused on road infrastructure (typically inter-city), rather than urban transit.¹² Our detailed microdata allows us to estimate fine parameters that are usually hard to incorporate in these studies, and shows how they are related to features of optimal networks.

Our paper is also related to the literature on the impact of transit systems in LMIC. Gaduh et al. (2022) studies the impact of the initial launch of the TransJakarta system in 2002 to show that at that point in time, the system had no impacts on public transport use but the conversion of road lanes into dedicated lanes for the BRT network worsened road

¹⁰See also Coulombel and Monchambert (2019) for evidence on dis-economies of scale.

¹¹Hörcher et al. (2017) estimate subway crowding costs in Hong Kong with a revealed preferences approach.

¹²Fajgelbaum and Schaal (2020) study a neoclassical trade model where the planner's problem is globally concave. Allen and Arkolakis (2022) study a tractable general equilibrium model with routing and congestion. Balboni (2021) and Santamaria (2022) compare different road investment plans in dynamic settings. Alder (2023) introduces a heuristic algorithm for finding the globally optimal highway network in the context of a quantitative trade model for India.

traffic congestion. Majid et al. (2018) study the impact of a new BRT transit line in Lahore, Pakistan. Tsivanidis (2022), Zárate (2023), and Balboni et al. (2021) measure the impact of the BRT systems in Bogotá, Mexico City, and Dar Es Salaam, respectively, and quantify general equilibrium effects through the lens of quantitative urban models. Tsivanidis, for example, focuses on estimating the welfare and inequality effects of the system and how it affects the organization of the city. All these papers take the public transport system design as given. By contrast, we tackle the complementary problem, focusing on estimating underlying preference parameters, and solving for the optimal system design.

Meta-heuristic algorithms, including the simulated annealing algorithm that we use here, have been used extensively in the transportation and operations research literatures to *approximate* optimal designs in urban transit network design problems (Wei and Machemehl, 2006; Iliopoulou et al., 2019). Our contribution is to link in an asymptotic sense the output of such algorithms to a social planner’s discrete choice problem over all possible networks. Instead of approximating an optimal network, we sample multiple networks, which allows us to estimate optimal network characteristics and comparative statics.¹³

The rest of this paper is organized as follows. Section 2 describes the TransJakarta bus network, its expansions, and our data. Section 3 presents the reduced-form event-study results that show how ridership responds to new routes and improved bus frequency. Section 4 introduces the model, describes how we estimate the model using the reduced form moments, and presents estimation results. Section 5 explores the implications for optimal network design.

2 Setting and Data

2.1 The TransJakarta Bus Network

Setting Description. TransJakarta operates an integrated bus system serving the central urban districts in Greater Jakarta — a metropolitan area of about 18 million people. Its routes are concentrated primarily within the city of Jakarta, though some extend to surrounding districts. Established in 2004, TransJakarta serves hundreds of thousands of passengers daily, with daily ridership peaking around 1 million in early 2020.

As of August 2019, TransJakarta has over 139 bus routes, which are a mix of Bus Rapid Transit (BRT) and non-BRT routes. BRT routes operate on dedicated bus lanes along designated stations. Passengers pay at turnstiles to enter the BRT station and board the bus. It has a network length of more than 120 miles of BRT corridors, the longest such

¹³Daganzo (2010), Fielbaum et al. (2016) characterize optimal transit networks in idealized environments.

system in the world (Institute for Transportation and Development Policy, 2019). Non-BRT feeder routes operate partially on normal city streets, and partially along the BRT corridors, where they also stop in BRT stations. These non-BRT routes connect locations that are further away from the BRT system.

TransJakarta costs a flat fare of IDR 3,500 (USD 0.25) per trip, regardless of distance. Payment is collected by tap-in smart cards, either at BRT stations or through fare machines on non-BRT busses; tap-out is enforced at a small number of BRT stations as well.¹⁴ Free transfers are allowed at BRT stations.

TransJakarta is the primary public rapid transit system in the city. A commuter rail system (KRL CommuterLine) with about 1 million ridership serves outlying areas for longer trips that rarely overlap with TransJakarta trips. In 2019, a single 16 km subway line was opened to serve about 80,000 riders per day. In 2018, privately provided shared transport, using minibuses, accounted for 2.9% of all trips in the larger Jabodetabek region (JUTPI2, 2019). The primary alternative to TransJakarta is private transport, consisting of a mix of motorcycles, private cars, motorcycle taxis (*ojek*), and automobile taxis.

TransJakarta Network Expansions. We study the large expansion of the TransJakarta network, which added more than 93 BRT and non-BRT routes between January 2016 and February 2020. Only one entirely new BRT corridor was launched during this period.¹⁵ The newly launched BRT routes run along existing BRT corridors, sometimes connecting two corridors or running an express route on certain sections. The new non-BRT feeder routes stop both in some BRT stations as well as in roadside non-BRT bus stops, connecting one or multiple BRT corridors to other areas of the city.

The number of operating buses more than doubled during this period, from about 700 to more than 1,600. Figure 1 shows the expansion of the system. These expansions — for both BRT and non-BRT routes — took place at different times throughout the city.

New lines were launched based on a mix of external constraints and discretion. All non-BRT route launches were chosen from a list of routes that were pre-approved in 2016 by the Jakarta transportation department (*Dishub*). For BRT routes, TransJakarta created new routes by connecting existing BRT stations along the existing BRT dedicated lanes. For both BRT and non-BRT, TransJakarta chose routes and launch periods based on new bus fleet delivery dates, bus availability based on bus-operator contracts, inputs from field reports, and other factors. Appendix Table A.2 shows that the order of BRT route launch is balanced with respect to the geographic characteristics of the routes.

¹⁴Cash payments are allowed on non-BRT busses, accounting for around half of all transactions in 2019.

¹⁵Corridor 13 is an elevated busway, thereby not affecting the number of lanes available to road traffic.

2.2 Data

We use three different datasets: **TransJakarta ridership** data from administrative data on TransJakarta’s smart card transactions; origin-destination **aggregate-trip flows** data from anonymized smartphone location data; and **bus locations** data, which we obtain by processing detailed GPS data on every TransJakarta bus throughout our period.

We focus on the central urban districts of Greater Jakarta.¹⁶ We use consistently defined geographical environments throughout our analysis. We divide space into identical grid cells and aggregate stations at the grid cells level, and trips at the grid cell pair level.¹⁷ Throughout the reduced form and demand estimation analysis, We use a regular hexagonal tiling with 1766 grid cells where adjacent grid-cell centroids are 1,000 meters apart, and then show robustness to using 500-meter and 2000-meter square grid cells. For the optimization section, we use 2,000-meter square grid cells to reduce the problem’s dimensionality. We aggregate outcomes (e.g. ridership, trips, bus supply) at the weekday or work-week level, focusing on the 5AM to 10PM interval.

Ridership Data. We use the administrative ridership data that is captured electronically by smart cards to construct a highly granular TransJakarta origin-destination (o, d) ridership matrix at each point in time since 2016 (since 2017 for non-BRT). In BRT stations, passengers tap smart cards at turnstiles to enter the boarding area. In 36 percent of BRT stations, passengers also tap smart cards upon exit. In non-BRT bus stops, passengers board the bus directly from the street and tap to pay for the bus ticket upon boarding. We observe the time of each tap, as well as an anonymized identifier for the smart card used, which allows us to link transactions from the same smart card over time.¹⁸

While we observe each passenger’s trip origins, we only observe destinations (d_{it}) for the trips that end at stations where tap-out is enforced. In all other cases, we construct a proxy destination station based on an algorithm that uses the origin station of the next trip for the same smart card, or another frequent origin for that card. We validate this algorithm using the existing tap-out data.¹⁹ We describe the algorithm and ridership data cleaning process and coverage in Appendix A.3.4.

¹⁶This consists of the Special Capital Region of Jakarta (DKI Jakarta) and all the surrounding urban districts: Tangerang, South Tangerang, Depok, and Bekasi (see Figure A.1).

¹⁷We collapse the TransJakarta network to the level of grid cells. For transfer options, we consider that origin o and destination d are connected by a transfer option by routes r_1 and r_2 if these two routes intersect in an actual station.

¹⁸Individuals can purchase new smart cards, but do so infrequently. The median smart card in our data is active for over 4 months, and the median tap belongs to a smart card that is active for over 20 months.

¹⁹For the set of destination stations where we have actual exit transactions, bivariate regressions of imputed and actual daily ridership shares have $R^2 = 0.85$ (Appendix Figure A.19).

Aggregate Trip Flows. We augment our administrative data with anonymized historical smartphone location data to measure overall trip behavior, regardless of whether a given commuter uses the TransJakarta system or not. We have daily smartphone location data from March 2018 through March 15, 2020 from Veraset, a private data provider. These data cover 35 million weekday trips that belong to 2.3 million unique devices in our study area. (For reference, our study area had 14 million individuals over 15 years old at the 2020 census.) We use these data to: (1) create a panel of trip flows separately for each week in the data as an outcome variable when studying the impact of the network’s expansion, and (2) construct a cross-section of typical (o, d) trip flows throughout the period in order to compute choice probabilities for using TransJakarta. See Appendix A.3.5 for details.

The smartphone location data captures rich travel patterns within Jakarta (Figure ??). Figure A.21 shows that the Veraset data is broadly representative across areas with different levels of poverty and density.

Bus Location and Allocation, and Network Expansion. We have data on the planned and realized daily bus allocations for each route, as well as exact GPS locations for every bus every five to ten seconds for the near-universe of routes. These data cover approximately 1,800 busses and 16,000 bus trips per day since January 2017. We process this raw GPS data to compute median bus travel times on all routes, between every origin and destination stations, and to measure the empirical distribution of wait time for each route. See Appendix A.3.2 and A.3.3.

To reconstruct the TransJakarta network at different points in time, we combine data from TransJakarta and data from a mobility planning app available in Jakarta during the study period. We cross-check launch dates inferred from bus allocation data with route- and date-specific TransJakarta ridership data, as well as directly with TransJakarta staff.

3 Reduced Form Results: the Impact of Service Quality

We begin by estimating the ‘reduced form’ impacts of improved service quality on TransJakarta ridership and on all trips (measured using smartphone location data). We use the fact that the launch of a new bus route changes different dimensions of how attractive it is to take TransJakarta for different (o, d) pairs, based on how it affects the choice set of direct and transfer bus routes between those locations.

We focus on three types of ‘events’ induced by new route launches. The first two capture the launches of the first *direct* route between two locations that are already connected by *transfer*. In Event 1, the new direct route is not faster than the fastest existing transfer

connection, in terms of total time on the bus. In Event 2, the new direct route is faster. Event 3 captures the addition of busses between two already directly connected locations because the new route overlaps with the existing direct route(s) between o and d . This increases the bus arrival rate and lowers wait times for those traveling from o to d . We examine these event types separately for BRT and non-BRT connections.

3.1 Estimation Framework

For each of the three event types described above, for both BRT ($E \in \{1B, 2B, 3B\}$) and non-BRT network expansions, ($E \in \{1N, 2N, 3N\}$), we estimate the following equation:

$$\ln(\mathbb{E}Y_{odt}) = \alpha^E Post_{odt}^E + \alpha_{-10}^E L_{\leq -10,odt}^E + \alpha_{10}^E L_{\geq 10,odt}^E + \mu_{od}^E + \nu_{ot}^E + \xi_{dt}^E + \varepsilon_{odt}^E \quad (1)$$

Our dataset is at the origin (o) \times destination (d) \times week (t) level, so we include all two-way fixed effects. Specifically, μ_{od}^E are origin \times destination fixed effects — i.e. fixed effects for every combination of start and end grid cell — and ν_{ot}^E and ξ_{dt}^E are origin \times time and destination \times time fixed effects. These fixed effects flexibly capture differences in ridership across origin-destination parts, as well as arbitrary time-related shocks for each origin and destination.

The key variable of interest is $Post_{odt}^E$, a dummy variable for the event having taken place on the od route in the previous 10 months.²⁰ We control for $L_{\leq -10,odt}^E$ and $L_{\geq 10,odt}^E$, which are dummies for whether an event between o and d takes place 10 or more months in the future, and in the past, respectively. The coefficient α^E in equation 1 thus captures the overall effect of an event of type E for the pair (o, d) in the first 10 months after it occurs, relative to the 9 months prior to the event.

We estimate (1) using robust Poisson Pseudo Maximum Likelihood (PPML). We cluster standard errors two-way by both origin and destination, which allows for arbitrary serial correlation over time for each origin and each destination, as well as arbitrary correlation among destinations for a given origin, and vice-versa (Cameron et al., 2011).

The sample includes only odt observations that are connected at time t within the Trans-Jakarta network, implying that both o and d contain TransJakarta stations before each type of event. The sample of pairs od for event type E is restricted to all origins treated at least once ($Post_{od't'}^E = 1$ for some d', t') and to all destinations treated at least once. Most of the origin-destination pairs included in the sample are never treated. (For example, for BRT event type 1 this number is 78.4 percent).²¹

²⁰For Event 3, we focus on the first event if there are multiple events for a given o, d pair.

²¹The fact that most of our od pairs are never treated largely alleviates the class of problems with two-way

We use two outcome variables Y_{odt} . First, we examine the total TransJakarta ridership between grid cells o and d during week t , summing up over all station pairs in o and d . Second, we also measure impact on all trips between o and d in week t , computed based on smartphone location data.

We use the full time period of data that we have for each type of event, but since the BRT, non-BRT, and smartphone trip data all begin on different dates, the data start at different times.²² All data end in mid-March 2020, prior to the COVID-pandemic shutdown.

For each of the events described below, we present the overall effects from estimating the corresponding version of equation 1. We also present event study graphs, where we add month-by-month leads and lags of the key explanatory variables.

In Events 1 and 2, a route launch creates the first direct connection between o and d , a location pair that already has a transfer connection. The two event types differ only in terms of whether the new direct route has faster travel time on the bus compared to the fastest pre-existing (transfer) option.

In Event 3, grid-cell pairs o and d that are already directly connected get more busses from a new route because it overlaps with the existing route for the portion between o and d . Specifically, $Post_{odt}^3$ is a dummy for the first event of an additional direct route launched between o and d taking place, in the ten months before week t .

3.2 “First Stage” Impacts on Travel and Wait Times

We first study how the three different events affect the attractiveness of riding the TransJakarta network from o to d (Table 1).

We first estimate the impact on the log minimum travel time on the bus between o and d . We use the bus GPS data to measure bus travel times, and at each time t between 2016 and 2020 we take the minimum over all available bus options (direct or single-transfer). This measure captures the pure travel time on the bus — i.e. it does not include any time spent waiting for a bus to arrive or any time waiting for a connecting bus.

Adding a new direct line between o and d reduces travel time, but only when the new line is faster than the existing transfer connection (Event 2), yielding an average effect of about 0.29 log points with a standard error of 0.027 (column 2). While the fact that travel time reductions are limited to Event 2 is, of course, mechanical to the way we define the events,

fixed effects specifications that have been highlighted recently (see, e.g., De Chaisemartin and D’Haultfoeuille 2022), which are primarily a concern when using previously-treated observations as a comparison group.

²²For BRT events, the sample comprises January 2016 to mid-March 2020. For non-BRT events, we further restrict the sample to the period after mid-January 2017, when our non-BRT data begins. When we use smartphone trips as an outcome variable, the time periods span from March 2018 to March 2020. We always exclude May-July 2018 due to missing BRT ridership data.

we show that the differences are quantitatively large and statistically quite meaningful.²³ Event 3 has a very small effect on log travel time.

Next, we examine the log number of busses arriving per hour at the origin, over all the direct and transfer options that connect the origin and destination. This is a proxy of waiting times at the origin station. Event 3 increases the bus arrival rate by 0.32 log points (Table 1, column 6). Events 1 and 2 have small effects on the bus arrival rates at the origin. Note that these effects do not capture the reduction in wait time for transfer connections.

Results for non-BRT network expansion are broadly similar and typically larger in magnitude (Table 1, panel B). Importantly, baseline bus arrival rates for treated origin-destination pairs is 2–4 times lower for non-BRT routes compared to BRT (columns 4–6). This means that the absolute reduction in wait time is much larger for non-BRT events compared to BRT events. Results are qualitatively similar when using 500-meter or 2000-meter square grids as the base geography (Tables A.3 and A.4).

The key point from Table 1 is that the three events affect the desirability of using the bus system in different ways: adding a new and faster direct line reduces travel time *and* reduces waiting times, in addition to eliminating the need to transfer, whereas additional busses mainly reduces waiting times.²⁴ This variation will allow us to identify preference parameters in Section 4.

3.3 The Impact of New Route Events on Bus Ridership

Travel Times and Direct Connections: Impact of New Routes. The impacts of the three events on bus ridership, estimated using equation (1), are presented in columns 7–9 in Table 1 for BRT and non-BRT. Figure 2 shows event-study versions of each of these equations with monthly lags and leads from the date of the event.

Ridership is highly responsive to improvements in service quality. On average, adding an additional direct BRT route between o and d without a travel time improvement (Event 1) leads to an increase in ridership of 0.16 log points over the 10 months following the introduction of the new route (Table 1, column 7). When the new direct route also decreases travel time (Event 2), which column 2 showed led to a 0.29 log point reduction in travel time, this leads to an increase in ridership of 0.27 log points.

²³The first type of event in fact has a small positive effect on log travel time. This is to some degree mechanical based on how we separate events 1 and 2. New direct routes are categorized under event 2 if the new direct route’s travel time is strictly smaller than that of existing transfer routes, and event 1 otherwise.

²⁴Both “first stage” variables are imperfect proxies for the attributes of interest. For example, the second outcome counts bus arrivals even for transfer options between o and d that have very long travel time, which in practice will be valued less by commuters. The minimum travel time between o and d shown in Table 1 might depend on a very infrequent bus line. The model in section 4 will lay out how commuters value these attributes jointly for all bus options in their choice set.

The event-study version of equation (1), shown in Figure 2, Panel (a), left column, shows no pre-trends before the first and second events, and a discrete uptick in ridership at the time of the new route launch. The increase is notably higher for Event 2. The event-study graphs also show that riders take some period to adjust to the new route.

The non-BRT results paint a qualitatively similar picture of the impact of bus service improvements (right-side of Figure 2 and Table 1, Panel B). While for all events, the point estimates are larger than those for BRT, the implied *level* effects are much more similar. For Event 1, the BRT effect implies that a treated origin, destination pair has 19.1 ($= 111.3 \cdot (e^{0.158} - 1)$) more riders per week on average in the post period. This compares with an effect of 33.8 additional riders per week for non-BRT Event 1. For Event 2, the BRT and non-BRT level effects are 23.6 and 20.6 additional riders per week, respectively.

Wait Times and the “Mohring Effect.” We next investigate the addition of direct busses between o to d caused by the launch of a new route that overlaps with existing direct routes (Event 3). Adding additional busses leads to a 0.09 log point increase in ridership (Table 1, column 9). Figure 2, Panel (b) shows that there is a large, discrete uptick in ridership exactly when the new route is introduced.

Given the first-stage results discussed above, we can interpret Event 3 as being almost entirely about the effect of wait times on ridership. Thus, combining the estimates from columns 6 and 9 of Table 1, we get an implied elasticity of ridership with respect to wait times of -0.29 , i.e. a 10 percent decrease in wait times leads to a 2.9 percent increase in ridership.

These estimates speak to the so-called ‘Mohring effect.’ [Mohring \(1972\)](#) argued that if demand for public transit is responsive to wait times, then there is an externality from riding the bus — more bus ridership allows the bus operator to add more busses to the route, decreasing wait times for other riders. Our estimates show that, indeed, ridership is sensitive to bus frequency, suggesting that this effect operates in this case and that the optimal public transit subsidy is likely positive for this reason.

An extreme form of the Mohring effect is when the elasticity of ridership with respect to wait times is greater than 1 in absolute value over some range. In this case, the planner may have multiple local optimal levels of service frequency in ridership: a low-ridership, high-wait-time regime, and a high-ridership, low-wait-time regime. For BRT, we can reject an elasticity of -1 at conventional significance levels.

For non-BRT routes, however, we find a much larger implied elasticity of ridership with respect to wait times, of -1.05 ($= 0.450/0.425$, Table 1, panel B). This suggests that adding more non-BRT frequency on some non-BRT routes could increase ridership enough to main-

tain (or even increase) *average* ridership per bus.

Similar to the first two types of events, the *level* impacts of Event 3 are similar for BRT and non-BRT. Specifically, the BRT effect implies that a treated origin, destination pair has 20.3 ($= 210.4 \cdot (e^{0.092} - 1)$) additional riders per week on average in the post period, relative to an effect size of 14.77 for non-BRT.

Alternative Aggregation Levels. One possible concern is that our results are capturing displacement effects within the TransJakarta network. This would be the case if, for example, some commuters would switch from using a certain origin bus station located in a grid cell o' to one of the treated origin grid cells o after an event.

To examine this, we consider different aggregation levels, where we re-estimate all our results using both 500-meter squares (smaller) and 2000-meter squares (larger). If substitution was a central concern, we would find much larger effects when we use smaller units of aggregation than when we use larger units. Figure A.5 compares the magnitude of coefficients for all three events, separately for BRT and non-BRT, placing estimated $Post_{odt}$ coefficients and their 95% confidence intervals side by side.²⁵ There is no systematic pattern of decreasing coefficients and confidence intervals overlap significantly in all graphs. This implies that our treatment effects are not significantly contaminated by displacement effects.

Heterogeneity Based on Poverty. One important question is whether preference parameters depend on local poverty levels. Table A.6 shows the impact of all six events on bus ridership when we include interactions for whether the origin grid cell has above-median poverty rate in the estimation sample; these regressions also control for interactions with log population. We use poverty and population data at the kelurahan (urban villages) from SMERU (2014) and the PODES 2010 village survey, respectively. The results show no clear patterns based on income levels. It is important to note the lack of heterogeneity here when thinking about optimal network design.

3.4 Impacts on Overall Trips

Having shown how the TransJakarta ridership responds to these events, we next turn to the Veraset smartphone-based trip data to examine whether *aggregate trip volume* — regardless of whether passengers take TransJakarta or not — is responsive to the TransJakarta network expansions.

We re-estimate equation (1) for the total number of trips from o to d at time t from Veraset (Table 2 and Figure 3). Our measure of aggregate trips is derived from individual-

²⁵The detailed 500 and 2000 meter results are shown in Tables A.3 and A.4 and Figures A.2 and A.3.

level smartphone traces, and it covers all types of trips, not only commuting trips between home and work. Since our Veraset data only begins in March 2018, for comparability we also re-estimate the effects on TransJakarta ridership for the same time period.

The key result is that we do not see positive and significant impacts of the three types of events on aggregate travel volumes between pairs of 1km hexagon locations. For example, for BRT Event 1, we find a coefficient of -0.008 with a standard error of 0.051 (column 10). This allows us to reject at the 95% level a positive impact of approximately $+0.091$, compared to the precise 0.11 effect on bus ridership in the same sample (column 7). Thus, we can rule out moderately positive impacts on all smartphone location trips. Figure 3 shows no clear patterns before and after the events. We find qualitatively similar results for most event types, and they are robust to analyzing the data at the level of smaller, 500-meter square grids (Table A.5).

These results suggest that our main results on bus ridership reflect an immediate and large mode substitution towards TransJakarta, without an increase in total trips between the affected origin-destination pairs. Of course, we cannot rule out that over a longer time period compared to the 10 months we focus on here, the pattern of trips will also change. However, based on these results, in our model we will focus on mode choice and hold (origin and) destination trip choices fixed.

Overall, the reduced form results suggest that riders are responsive to the three dimensions of service quality we consider: wait times, on-the-bus ride times, and direct connections. We do not find evidence that these changes affect the aggregate volume of trips. Since the events we study affect multiple dimensions of service quality, we next estimate a model of commuter demand, which allows us to jointly infer the underlying preference parameters that best match the responsiveness we observe in the data.

4 Model and Estimation

We now set up a model of demand for public transportation that describes how commuters choose bus routes within the TransJakarta network, and at a higher level, whether they use the TransJakarta network or an “outside option” that captures private transport modes.²⁶ We consider the problem of a commuter traveling between an origin grid cell and a destination grid cell who chooses between direct bus routes that link the origin and destination grid cells

²⁶While it is possible to further embed these decisions into a model of destination choice, or a general equilibrium urban equilibrium model (Tsivanidis, 2022), given our null results on aggregate trips in this context, we hold these decisions fixed in the model.

(we assume one bus station per grid cell), single-transfer bus connections, and an outside option. The exact choice set for bus options will depend on the state of the public transport network at a point in time.

How do commuters decide which bus routes to take? The core of our model is a new formulation for how commuters choose routes in a transit network. The model highlights the importance of wait times, which affect decisions in two key ways. First, overlapping routes between an origin and a destination effectively decrease wait times, because the traveler can take the first bus that arrives among these bus routes. (This is exactly the mechanism that we study empirically in Event 3.) Second, travelers sometimes forego short wait times for other route characteristics. For example, a traveler may decide to not take the first bus that arrives at the station if this option involves a transfer and instead wait longer for a direct (and shorter) route.

In the model, a traveler’s utility from a given option is additive in wait times and several deterministic factors: travel time and the necessity to make a transfer. For each route and station, bus arrivals follow a Poisson arrival process. Commuters draw a vector of random wait times for the different routes in their choice set. This wait-time randomness induces idiosyncratic heterogeneity in route choices.

Our network routing model has four desirable properties. First, the model is invariant to aggregation of identical routes. Second, we confirm using our bus GPS data that the distribution of wait times is approximately exponential, as assumed in our model (Appendix Figure A.18). Third, our model has realistic predictions: overlapping routes lead to shorter wait times, and commuters sometimes choose to wait for a later bus arrival in order to use a better route. Fourth, we obtain tractable expressions for expected utility and choice probabilities.

We add to this baseline formulation the possibility that commuters are partly inattentive to certain options in their choice set. We assume that a route’s arrival rate may be attenuated by an attention factor that depends on the travel time on the bus of options using that route, relative to the minimum travel time on the bus in the choice set. The model with full attention is a special case.

In the upper nest, commuters make a logit choice between using the private mode or the TransJakarta network, where the benefit of the latter is given by the risk-neutral expected utility over wait times of choosing from among available TransJakarta options. The private option in the model includes private vehicles, such as motorcycles, taxis, motorcycle taxis, and ride-hail apps, as well as other types of shared transport, such as mini-buses outside the TransJakarta network. We assume that these modes are infinitely elastic to meet demand at a given cost, but their prices and attributes do not vary over time. We also assume that

the TransJakarta network does not affect road traffic congestion.²⁷

We estimate the model parameters leveraging the TransJakarta network expansion, focusing on commuters’ values for travel time, wait time, and transfers, as well as attention parameters, which we allow to differ for BRT and non-BRT routes. We use a classical minimum distance approach whereby we search for the parameter vector that allows the model-predicted ridership to replicate the reduced form analysis documented above. In addition to the six moments (three events for BRT and non-BRT), we also use a seventh moment, described below, to help pin down the attention parameters.

4.1 Bus Route Choice Model

A Random Utility Model with Waiting Times. The core part of our model is a flexible, static discrete choice over bus options, where idiosyncratic heterogeneity is given by exponentially distributed random wait times. [Daganzo \(1979\)](#) calls this the “negative exponential distribution model.” We first describe the general structure and key properties of the model, before delving into the details that are specific to bus network routing.

Consider an agent making a static discrete choice among a finite set of options $k \in K$. Each option has a deterministic utility v_k and a random component $w_k \geq 0$, which measures wait time. The wait time for each option is governed by an independent Poisson process with arrival rate λ_k , which can differ by option. This means that the wait time w_k is drawn independently from an exponential distribution $\Pr(w_k > w) = \exp(-\lambda_k w)$. The agent observes the wait times for all options and picks the option with the highest combined utility $u_k = v_k - \alpha_{\text{wait}} w_k$, where α_{wait} governs the agent’s preferences over wait times.

We incorporate partial inattention by assuming that the agent notices each arrival from option k with independent probability p_k . In other words, with probability $1 - p_k$ the agent fails to notice the first arrival for option k . This leads to an “effective” arrival rate for option k of $\tilde{\lambda}_k = \phi(p_k)\lambda_k$ with $0 \leq \phi(p) \leq 1$ (see [Appendix A.5](#)).

This model has a convenient invariance property: it is unchanged when options that have the same indirect utility are combined into a single option with the sum of arrival rates.

Proposition 1. *Assume that the choice set contains two options $k = 1, 2$ with $v_1 = v_2$ and arrival rates λ_1 and λ_2 . The model where options 1 and 2 are replaced by a single option 3 with $v_3 = v_1 = v_2$ and $\lambda_3 = \lambda_1 + \lambda_2$ is isomorphic to the original model in terms of choice probabilities and expected utility.*

²⁷BRT routes run on dedicated lanes separate from other road traffic, and new BRT routes during our study period were launched along existing BRT corridors, without changes in the supply of lanes for other road traffic. We assume that non-BRT routes did not affect road traffic congestion due to their low frequency and relatively low mode share.

Proof. This follows from the fact that the sum of two independent Poisson count processes with arrival rates λ_1 and λ_2 is a Poisson count process with arrival rate $\lambda_1 + \lambda_2$. \square

This property implies that the demand model is unchanged if a route is split into two identical routes, with the total number of buses (which, as we will see, determine arrival rates) split between the two new routes. In other words, our model does not feature the ‘red bus, blue bus’ problem discussed by McFadden in the context of the multinomial logit and other random utility models. This will prove to be an attractive invariance property for the planner’s optimization problem that we will later set up in section 5.

Choice probabilities and expected utility are as follows:

Proposition 2. *Assume that options are ranked $v_1 < v_2 < \dots < v_N$. The probability π_k to choose option k is given by*

$$\lambda_k^{-1} \pi_k = \sum_{i=1}^k e^{-M_i} \frac{e^{v_i \alpha_{\text{wait}}^{-1} \Lambda_i} - e^{v_{i-1} \alpha_{\text{wait}}^{-1} \Lambda_i}}{\Lambda_i},$$

where $\Lambda_i = \sum_{\ell=i}^N \lambda_\ell$ and $M_i = \sum_{\ell=i}^N v_\ell \lambda_\ell$, and $v_0 = -\infty$ by convention. Expected utility is

$$\mathbb{E} \max_k u_k = v_N - \pi_N \frac{\alpha_{\text{wait}}}{\lambda_N},$$

where π_N is the choice probability of the option with the highest deterministic utility.

The proof uses algebraic manipulations of the exponential distribution (Appendix A.6). The first part of this result ensures that computing choice probabilities is computationally tractable. The second part shows that expected utility has a particularly simple expression. Note that if the commuter only had the option N in their choice set, expected utility would be $v_N - \alpha_{\text{wait}}/\lambda_N$. In general, the influence of other options $k \neq N$ on expected utility is summarized by the probability π_N to choose the top option.

Bus Route Choice Model Setup. The travel demand model in the context of the TransJakarta network is as follows. Consider a commuter i traveling from a given origin grid cell o to a destination grid cell d . They have a choice set with a finite number of bus options $k \in M_{odt}$, where the choice set depends on the TransJakarta network at calendar date t , which in our application will range between 2016 and 2020. Each option has a utility level $u_k = v_k - \alpha_{\text{wait}} T_k^{\text{wait}}$, where v_k is the deterministic component that depends on characteristics such as travel time and transfer terms (in cases where the bus option k involves a transfer), and T_k^{wait} is the wait time at the origin determined by a Poisson arrival process that we describe in the next section.

The bus route choice set M_{odt} contains all direct and single-transfer bus connections between o and d available at t . The bus connections must start in the origin grid cell o , end in the destination grid cell d , and transfer in an intermediate TransJakarta station m .

The utility for a *direct* public transit option k is:

$$u_k = \underbrace{-\alpha_{\text{time}} T_k^{\text{time}}}_{v_k} - \alpha_{\text{wait}} T_k^{\text{wait}} \quad (2)$$

where T_k^{time} and T_k^{wait} are travel time on the bus and wait time, respectively.

If passengers decide to take a *transfer* option $k \in M_{odt}$ that includes connecting in an intermediate station, we assume that they get the utility from the first leg k_1 of the route up to the intermediate station, and the expected utility from the second leg of the route, taking expectations over the best direct option k_2 connecting the intermediate station to the destination:

$$u_k = \underbrace{-\alpha_{\text{time}} T_{k_1}^{\text{time}} + \mathbb{E} \max_{k_2} [-\alpha_{\text{time}} T_{k_2}^{\text{time}} - \alpha_{\text{wait}} T_{k_2}^{\text{wait}}]}_{v_k} + \mu_{\text{transfer}} - \alpha_{\text{wait}} T_{k_1}^{\text{wait}}. \quad (3)$$

We assume that second-leg wait times are not known at the time when the commuter makes the initial decision, so the consumer needs to take expectations over the best route k_2 that they will take from the intermediate station to the destination. μ_{transfer} captures the pure transfer penalty of taking a transfer *above and beyond* the travel time and the (expected) utility for the second leg.

Preference parameters are allowed to differ by BRT/non-BRT. For clarity, we do not track this distinction in notation now and discuss it when we set up the model's estimation.

We assume the following timing for the commuter's decision. Conditional on deciding to use the TransJakarta network, the commuter goes to the (unique) station in her origin grid o and observes wait times T_k^{wait} for all possible routes for options k in her choice set. For transfer routes, she observes the wait time for the first leg and forms expectations about the wait time for the second leg. She then chooses the bus option that maximizes utility.²⁸

Poisson Bus Arrival Process and Exponential Wait Time Distributions. Busses on a route r arrive at the origin grid o according to a Poisson count process with arrival rate λ_r . The arrival rate is given by $\lambda_r = N_r / RTT_r$, the ratio between the number of busses allocated to that route, and the return travel time needed for one bus to make a full loop on

²⁸During our study period, TransJakarta posted actual arrival times at each BRT station – see Figure A.17 for an example. Users could also look up wait times from home using an app, but these may change between the time the user leaves home and arrives at the station.

route r .²⁹ We assume that the Poisson processes for different routes are independent.

Agents in the model are “non-planning” in that they cannot synchronize their trip departure time with a specific bus given an arrival schedule. This is realistic for TransJakarta, which does not publish a schedule and where very infrequent service is rare. Under this assumption, the wait time at o for the next bus from a given route r is exponentially distributed. Using bus GPS data, we show that exponentially distributed wait times are a good approximation for small and moderate wait times (Figure A.18, panels A and B). The exponential distribution in the model slightly over-estimates long wait times, suggesting that our estimate of α_{wait} captures steeper-than-linear costs for large wait times.

The mean and variance of wait time for a route follow a tight relationship, which is similar for BRT and non-BRT routes (Figure A.18, panel C). Without independent variation in these moments, we note that our estimate of α_{wait} will capture the combined preference for the wait time distribution.

Large Choice Sets and Partial Inattention. The model outlined so far assumes that commuters consider all possible bus route combinations. However, the TransJakarta bus network is sufficiently interconnected and complex that some choice sets are very large. While the median choice set has four bus options, a quarter of origin-destination pairs have choice sets with at least 14 options, and 10 percent of choice sets have at least 28 bus options.

Many choice set options — those with very long travel times — are ex-ante unattractive. For the choice set at the 75th percentile, the slowest option in the choice set (by total travel time on bus) is 2.8 times slower than the fastest option in the choice set. For 10 percent of choice sets, this ratio is at least 3.6.

We allow commuters in the model to be partially inattentive to options in their choice set that have a long bus travel time. This is a simple heuristic for reducing the size of the consideration set, bearing some similarity to models of rational inattention (Gabaix, 2019).³⁰

To capture partial inattention, we consider for each route an effective arrival rate $\tilde{\lambda}_{rod} = \lambda_r \cdot \phi_{rod}$, where $\phi_{rod} \in [0, 1]$. This is equivalent to assuming that for any bus on route k , the commuter does not observe it with some probability, independently of other buses. Note that under this model, the commuter still draws a waiting time for each route, but this may be longer for some routes. In our parametrization, the ϕ_{rod} terms will be close to zero or one, so essentially, the commuter will either fully ignore or pay full attention to the route.

We parameterize the attention factor as a function of travel time on the bus. Specifically, let T_{rod}^{\min} denote the shortest travel time on the bus among bus options that include route r ,

²⁹Bus allocation and bus speeds are mostly constant throughout the day (Appendix A.3.2).

³⁰Larcom et al. (2017) use a short-term disruption due to a strike to show that some commuters in London are using suboptimal routes.

and T_{odt}^{\min} the shortest time in the entire choice set between o and d at time t . We define

$$\phi_{rod t} = \phi\left(\frac{T_{rod t}^{\min}}{T_{odt}^{\min}}\right),$$

where $\phi(T_{\text{ratio}}) = \frac{\exp(\gamma(\eta - T_{\text{ratio}}))}{1 + \exp(\gamma(\eta - T_{\text{ratio}}))}$ has a decreasing logistic shape in the travel time ratio T_{ratio} . We will estimate the parameter η , which is an attention cutoff. Commuters pay significantly less attention when the travel time ratio is above η , and this function approaches a step function around η as the shape parameter γ grows to infinity.

Partial inattention leads to important differences in how commuters respond to changes in the TransJakarta network. Consider the addition of a faster route on top of a set of existing bus options, as in Event 2. With full attention, the new route can significantly decrease the waiting time at the origin, which increases bus network ridership. This effect is particularly strong when time on the bus is not very costly (low α_{time}) because in that case, all routes are good substitutes for each other. With partial inattention, the commuter will pay less attention to pre-existing bus options after the new faster route is launched, because of the travel time difference. Partial inattention can rationalize low responsiveness to this type of change in the choice set.

4.2 The Logit Choice Between Bus Network and Private Option

At the higher decision nest, a commuter traveling between o and d first decides between using the TransJakarta network or using an outside option, a catch-all for private modes (private motorcycle, for hire motorcycles, car, other private minibusses, etc.). The utilities for the two options for a trip i between o and d at time t are:

$$u_{it}^{\text{bus}} = \underbrace{\left(\mathbb{E} \max_{k \in M_{odt}} u_k\right)}_{v_{odt}^{\text{bus}}} + \epsilon_{it}^{\text{bus}} \quad (4)$$

$$u_{it}^{\text{private}} = \zeta_{od}^{\text{private}} + \epsilon_{it}^{\text{private}}$$

where the term in large brackets captures expected utility over different realizations of wait time vectors. The term $\zeta_{od}^{\text{private}}$ captures all time-invariant factors that make the private option more attractive for that specific origin-destination pair.

The terms $\epsilon_{it}^{\text{bus}}$ and $\epsilon_{it}^{\text{private}}$ are Gumbel-distributed error terms with parameter β , giving rise to logit probabilities.

4.3 Model Estimation

Our estimation strategy is based on finding the vector of preference parameters that allows the model to match the reduced form results documented in section 3. We estimate

$$\theta = (\alpha_{\text{time}}, \alpha_{\text{wait}}^{\text{BRT}}, \mu_{\text{transfer}}^{\text{BRT}}, \eta^{\text{BRT}}, \alpha_{\text{wait}}^{\text{non-BRT}}, \mu_{\text{transfer}}^{\text{non-BRT}}, \eta^{\text{non-BRT}}).$$

That is, we estimate a single cost of travel time on the bus, and separate BRT and non-BRT costs of waiting, transfer shifters, and attention cutoffs. Waiting in non-BRT stations and transfers between non-BRT and BRT bus lines may differ compared to the costs of using the BRT network. For example, the non-BRT bus stations are on the side of the road and differ significantly from the BRT stations, which are covered and where commuters pay at turnstiles to enter. However, we impose that travel time on the bus is valued similarly, given that the physical buses used on non-BRT and BRT routes are similar. Because we lack price variation to pin down the level of utility, we normalize the variance of logit shocks in (4).³¹ We set the attention shape parameter $\gamma = 30$, which means that attention drops quickly to close to zero above η^{BRT} and $\eta^{\text{non-BRT}}$. We estimate the origin-destination private option attractiveness terms ζ_{od} to match the average bus ridership between o and d over time.

Event Study Moments. We match the following seven moments. First, we match the six “Post” coefficients α^E on bus ridership for all three events, separately for BRT and non-BRT, from Table 1.

To help pin down the attention parameters, we add a seventh moment focused on trip duration. We estimate the impact of BRT Event 2 on the logarithm of the total trip duration from o to d .³² We first measure BRT trip duration (inclusive of waiting time) in the data using the specific tap-in and tap-out times for the same smart card at the BRT entry and exit station. To do this, we restrict to stations that are tap-out compliant, defined as when tap-out transactions are at least 30% of all taps at that station. Thirty-six percent of all stations (92 stations) are tap-out compliant according to this definition.

Travel times between o and d go down sharply after the new direct and quicker route is launched between o and d (Figure 4 and Table 3). The magnitude is 0.06 log points, which is smaller than the impact on log minimum travel time (-0.28). This reflects a mix of two effects. Travel time on the bus indeed falls to the extent that more commuters substitute to

³¹We normalize the logit parameter $\beta = \frac{\bar{D}}{D_{od}}$ where D_{od} is the straight line distance between grids o and d , and $\bar{D} = 8.5\text{km}$ is the average distance. Using a constant β would lead logit choices that are nearly random for o, d pairs that are close because the travel time component of utility is generally increasing in straight line distance. Our normalization compensates for this mechanical effect.

³²We focus on Event 2 because it significantly changes travel time on the bus.

the new direct route. However, realized wait times increase as commuters are more likely to wait for the new direct route. The model estimation will help disentangle the attention and value of travel time parameters that give this result.

Computing Event Study Moments in the Model. For a parameter vector θ , we compute model-predicted ridership for all origin-destination pairs and all time periods. Using the 1km hexagonal grids, we compute ridership for 33,880 o, d pairs (covering a maximum number of 350,493 bus route options), for all versions of the TransJakarta network between January 2016 and March 2020.

Given θ , for any o, d, t we compute choice probabilities

$$\pi_{odt}(\theta) = \frac{\exp(\beta v_{odt}^{\text{bus}})}{\exp(\beta v_{odt}^{\text{bus}}) + \exp(\beta \zeta_{od}^{\text{private}})},$$

which depend on the bus route choice model through the expected utility term v_{odt}^{bus} . Combined with the (cross-sectional) smartphone commuting flow data V_{od} , this yields model-predicted ridership for each o, d, t . We also compute expected trip duration. The choice set between o and d at t includes all TransJakarta routes already launched at t that stop in both o and d . Transfer options consist of a first-leg route r_1 that connects the o to an intermediate station m , such that there is at least one direct route between m and d .³³

We use GPS data to measure travel times on the bus within the network. (Appendix A.3.2 shows that these are stable within days and across our entire study period.) In the model, we assume that each route has a bus allocation that is fixed over time, equal to the average bus allocation in the data. This determines the route’s bus arrival rate.³⁴

Given θ , we estimate $\zeta_{od}^{\text{private}}$ to match average bus ridership over time between o and d .³⁵

We next run the reduced form analysis on model ridership, which we will then match to the empirical moments from Section 3. Given the high-dimensional fixed effects in equation (1), estimating PPML regressions at each iteration of the model is computationally prohibitive. To address this challenge, we proceed in three steps. First, we approximate the result of PPML estimation of the level of model-predicted ridership $R = R_{odt}^{\text{model}}(\theta)$ with a weighted least-squares regression on $\log(R)$.³⁶ Second, we obtain a significant computa-

³³Transfer options require a connection at the level of a real TransJakarta station m , not only a grid cell.

³⁴We use a “mechanical” bus allocation to avoid issues of reverse causality in model-predicted ridership.

³⁵We compute the bus share as the ratio of bus ridership from TransJakarta data to total commuting estimated using smartphone location data. When this ratio is above 1, we use a Bayesian procedure to estimate the posterior over $(0, 1)$. We then find ζ_{od} to match the mean of the posterior ratio.

³⁶In the presence of heterogeneous treatment effects, running PPML with outcome R and OLS on $\log(R)$ will weight treatment effects differently (Tyazhelnikov and Zhou, 2021). That paper shows that using R or a proxy of R as weights leads to a weighting of heterogeneous treatment effects that is similar to PPML. In our

tional speed improvement by pre-computing a vector x^E for each specification E , such that the coefficient of interest α^E from WLS is given by the dot product $\alpha^E = (x^E)'Y_{odt}$ for any outcome vector Y_{odt} .³⁷ Third, we use a slightly coarser set of fixed effects: origin-destination, origin-quarter-of-the-year, destination-quarter, and week fixed effects. We confirm that the reduced form results are very similar to this specification relative to (1), and after estimation, we report model fit results using the original, finer fixed effects.

Estimation and Inference. We use classical minimum distance estimation (CMD) to match the vector of seven numbers $\hat{m} = (\hat{\alpha}^{1B}, \hat{\alpha}^{2B}, \hat{\alpha}^{3B}, \hat{\alpha}^{1N}, \hat{\alpha}^{2N}, \hat{\alpha}^{3N}, \hat{\alpha}^{2B, \text{duration}})$, where $\{\hat{\alpha}^{1B}, \hat{\alpha}^{2B}, \hat{\alpha}^{3B}, \hat{\alpha}^{1N}, \hat{\alpha}^{2N}, \hat{\alpha}^{3N}\}$ are the empirical estimates for events 1, 2, and 3 for both BRT and non-BRT estimated in Section 3 above and $\hat{\alpha}^{2B, \text{duration}}$ is the travel time moment discussed above. We find the parameter vector θ that minimizes the objective function

$$\min_{\theta} (m(\theta) - \hat{m})' \widehat{W} (m(\theta) - \hat{m}), \quad (5)$$

where $m(\theta)$ is the vector of model moments, and $\widehat{W} = \widehat{\Omega}^{-1}$ is the optimal weighting matrix given by the inverse variance-covariance matrix of the moments \hat{m} . The estimate $\widehat{\Omega}$ comes from estimating all reduced form event coefficients “stacked” in a seemingly unrelated regression framework. We cluster standard errors two-way by origin grid and by destination grid, which introduces dependence between the different regressions. To reduce the risk of finding a local minimizer θ of (5), we repeat the optimization routine starting from 20 randomly selected initial conditions and confirm that all converge to the same vector.

To obtain confidence intervals for $\hat{\theta}$, we repeat estimation 100 times where we match the moment vector $\hat{m}^k = \hat{m} + \varepsilon^k, k = 1, \dots, 100$, where $\varepsilon^k \sim \mathcal{N}(\mathbf{0}, \widehat{\Omega})$ are draws from a multivariate normal distribution centered at zero with covariance matrix $\widehat{\Omega}$ (i.i.d. over k). We use the resulting $\hat{\theta}^k$ estimates to construct confidence intervals.

Identification. Our classical minimum distance strategy has two main advantages. First, the reduced form moments that we established in Section 3 correspond to model comparative static exercises of exogenously varying wait time, travel time, and directedness. Second, these moments are also directly relevant for our counterfactuals, where we will be considering

case, we use the measured TransJakarta ridership R_{odt}^{data} as weights, as a proxy for $R_{odt}^{\text{model}}(\theta)$. Unlike model ridership, which depends on θ , actual ridership is pre-determined before estimation, which is an advantage for the procedure discussed in the next paragraph.

³⁷The vector x^E is a row of the WLS matrix $(X'WX)^{-1}X'W$ where X is the matrix of covariates, including all fixed effects, and W is the weighting matrix. We compute this matrix inversion only once before estimation, which relies on using pre-determined weights. To further reduce size, we apply the Frisch-Waugh-Lovell theorem to partial out the fixed effect with the largest number of categories.

alternate networks that differ in terms of these attributes.

To shed light on how model parameters and the event study moments used in estimation are linked, we calculate the Jacobian matrix, which measures how, in the model, each moment’s value changes when each of the parameters changes. Table A.7 reports the matrix of derivatives $dm_i(\theta)/d\theta_j$, each measuring how moment m_i responds to a marginal change in parameter θ_j , and Figure A.9 plots these relationships for larger parameter changes. For clarity, both exhibits focus on BRT moments and parameters. The value of travel time on the bus α_{time} significantly and positively affects Event 2, with muted impacts on the other events, which is expected given that travel time does not change significantly for Events 1 and 3. The transfer shifter μ_{transfer} affects Events 1 and 2 negatively, with an impact on Event 3 that is almost zero. Again, this is as expected, as treated (o, d) pairs in the third event are already directly connected.

All parameters are estimated jointly, yet given this approximately upper-triangular matrix format, we can describe intuitively how parameters are identified in the following recursive manner. First, Event 3 depends almost only on the value of wait time, so its value pins down the value of wait time. Next, Event 1 depends on wait time and the transfer shifter, so its value pins then down the transfer shifter. Finally, with these two parameters in hand, Event 2 pins down the value of travel time on the bus.

4.4 Estimation Results

Table 4 shows the estimated parameters and 95 percent confidence intervals. We normalize some parameters relative to travel or wait time to ease interpretation. We first estimate the model only for BRT parameters in column 1, using the three BRT events as moments, as well as Event 2 for log trip duration.

Commuters in Jakarta view time spent waiting for the bus as more costly compared to travel time on the bus — the estimated value of wait time $\alpha_{\text{wait}}^{\text{BRT}}$ is 2.4 times larger compared to the value of time on the bus $\alpha_{\text{time}}^{\text{BRT}}$, and we can reject at 95% significance level that the two parameters are equal.³⁸ The wait time parameter governs the degree to which people are sensitive to frequent service.

There are several reasons why commuters may dislike wait time more than time on the bus. In our model, α_{wait} implicitly captures the importance of the entire wait time distribution, including wait time uncertainty. Waiting conditions may be less comfortable than

³⁸By comparison, existing studies of time valuation in public transport contexts focus on Western countries, and most are based on hypothetical choices. The few that use a revealed preferences approach typically use observational variation in attributes, unlike the quasi-random variation we use here. Papers in this literature typically find that the ratio of wait time to in-vehicle time is above one, but with significant dispersion (Wardman, 2004).

riding the bus (for example, bus stations are not air-conditioned). In addition, commuters may mis-perceive wait time as being longer than it actually is.³⁹

Turning to the pure transfer penalty, we cannot reject the hypothesis that $\mu_{\text{transfer}}^{\text{BRT}}$ is equal to zero. The 95% confidence interval is relatively tight around zero, allowing us to reject a cost equivalent to 2 minutes of wait time (which is the same as 4.3 minutes of travel time) or a positive benefit of 4.5 minutes of wait time.

Recall that μ_{transfer} captures preferences for transfer routes *above and beyond* the cost of transfer routes in terms of additional wait time (e.g. when making the connection) and any additional time on the bus. Our results suggest that there is no *additional* utility cost of transfers associated with, for example, having to give up one’s seat on the first bus and move around. The main factors that matter are time spent waiting and traveling on the bus.⁴⁰

We find that commuters exhibit inattention to the longer bus options in their choice set. Specifically, a commuter with estimated preferences is less attentive to bus options that are more than 34% longer than the quickest option in the choice set. This cutoff parameter is precisely estimated between 1.25 and 1.45.

In the second column, we jointly estimate BRT and non-BRT parameters, adding the moments for non-BRT events 1,2, and 3. The general pattern of results is similar. The ratio of the cost of wait time to travel time is 4.2 for non-BRT, which is statistically significantly larger than the ratio for BRT. This means that waiting for buses in non-BRT stations is even costlier than in BRT stations. The higher wait cost is perhaps not surprising given that non-BRT stations are just stops by the side of the road, not enclosed stations as with the BRT. We cannot reject that the transfer shifter is zero. However, the point estimate is now negative, and we can only reject at 5% level a transfer cost equivalent to 5.3 minutes of wait time, which is the same as 22 minutes of travel time. The attention parameter is $\eta^{\text{non-BRT}} = 1.44$, slightly higher but still broadly similar compared to BRT.

We next explore the model fit. Both models in Table 4 are exactly identified and the model perfectly matches the data moments at the estimated parameters. However, we can learn more about model fit by comparing the time series patterns, which are not explicitly matched. Figure A.6 shows these results, replicating the event study graphs from Figure 2 using model-predicted ridership. In general, for all six moments, the model does a good job of replicating the lack of pre-trends and increase in ridership after the launch of a new

³⁹For example, Fan et al. (2016) conducts a survey experiment in Minnesota and finds that people mis-recall wait times, reporting stated wait times that are larger than what they actually experienced.

⁴⁰As in the case of wait time, there is a lack of estimates of the pure transfer penalty based on quasi-random variation, as in this paper. Currie (2005) reviews existing studies, concentrated in Western countries and based mostly on hypothetical choice or, to a lower extent, observational variation, and reports a pure transfer penalty of 22 minutes of bus travel time on average over studies, ranging from 5 to 50 minutes.

route. The model also does a good job to fit aggregate ridership trends, including after an origin-destination pair first becomes connected by non-BRT (Figures A.7 and A.8), with the mention that ridership adjustments happen more gradually in the data.⁴¹ Overall, these results help build confidence in the model predictions for ridership and commuter welfare of counterfactual routes, which we consider in the last section of the paper.

5 Optimal Network Design

We now use the preference parameters estimated in the previous section to study the implications for the bus transit network design problem in Jakarta. We ask how the planner can optimally re-allocate bus routes and the bus allocation to routes, holding the total number of busses and the BRT infrastructure fixed, to improve welfare given the estimated preference parameters. We focus on two key exercises. First, we compare the current network used by TransJakarta to the planner’s solution. Second, we study how the shape of the network solution changes when structural preference parameters change.

This problem is discrete and high-dimensional. It does not appear to have an analytical solution and it resembles NP-hard problems such as the traveling salesperson problem. A key challenge is that it exhibits both substitution and complementarity forces, so we cannot use results that rely on complementarity properties (Jia, 2008; Arkolakis and Eckert, 2021).

We introduce a general framework that allows us to characterize optimal allocations and perform comparative static exercises in our setting.⁴² The key idea is to allow the planner to care both about model welfare as well as idiosyncratic shocks defined for every possible network. We assume these shocks are unknown to the researcher. The key implication is that for two networks N, N' with similar model welfare $W(N), W(N')$, small shocks may tip the planner into preferring one network over the other. This “smoothes” the planner’s problem and leads to a *probability distribution* over networks, where networks with higher model welfare have higher probabilities of being chosen by the social planner as the ‘optimal’ network. We assume that shocks are type-I generalized extreme value (Gumbel) distributed, giving a multinomial logit distribution.

We can then characterize certain properties of optimal networks by taking expectations over the planner’s distribution. We derive analytic formulas for expected optimal properties and their local comparative static with respect to model parameters. These objects can

⁴¹While these aggregate moments are not directly targeted, the terms $\zeta_{od}^{\text{private}}$ are estimated to match average bus ridership over the estimation sample time period.

⁴²The methods we develop may also be applied to quantitative models where optimal policies are high-dimensional, discrete, and analytically intractable, especially when exact hat methods make it possible to compare welfare between two counterfactual policies.

be easily estimated with a sample of networks independently drawn from the planner’s distribution.

In this setting, the objective is to sample from the planner’s multinomial logit distribution. We show that established algorithms, such as Metropolis-Hastings, parallel tempering, or a modified version of simulated annealing (SA), asymptotically sample from the planner’s distribution. In our application, we run multiple independent SA simulations to obtain an i.i.d. sample of optimal networks.

We find that, through the lens of the model, optimal networks are more expansive than the current TransJakarta network. Increasing the cost of wait time leads to less expansive networks, while increasing the transfer penalty makes the network slightly more direct.

5.1 Bus Network Optimization Environment

The planner’s problem that we consider here is to design a bus network N that maximizes average commuter expected utility. This involves deciding the number of bus routes, where bus routes run, and how to allocate a fixed total number of busses to these routes. Formally, a network is a tuple $N = (K, (r_1, \dots, r_K), (b_1, \dots, b_K))$ where K is the number of routes, r_k is a route, and b_k is the number of busses on route r_k .

We divide the greater Jakarta area into 418 $2\text{km} \times 2\text{km}$ square grid cells. Each grid cell is connected to its eight adjacent grids, including diagonals. A bus route is a non-self-intersecting path on this graph with 1,536 edges. Busses travel in both directions along each route and stop at each grid station. We construct predicted bus travel times for all possible edges using the existing bus travel time data, separately for BRT and non-BRT (see Appendix A.7.1 for details).

The total number of busses and BRT infrastructure are held fixed.^{43,44}

We use the travel demand model that we estimate in section 4. A commuter in the model going from o to d chooses among available transportation options given the proposed network N and estimated preferences $\hat{\theta}$, including inattention parameters. Commuting flows between every (o, d) pair are fixed, computed using the smartphone location data (see section 2.2).⁴⁵ The counterfactual networks we consider here comprise intensive and extensive margin changes that resemble the historical changes that our model fits or is estimated to match,

⁴³In counterfactual networks, each route has the BRT travel time on edges that have BRT infrastructure at present, and non-BRT travel times on all other edges.

⁴⁴We assume that bus fuel costs only depend on travel time and hence do not change in counterfactuals.

⁴⁵We re-estimate $\zeta_{od} = \zeta(D_{od})$ as a non-parametric function of distance D_{od} such that the model evaluated at the current TransJakarta network matches the total ridership for each distance bin. This gives us an imputed private option attractiveness for all (o, d) pairs. It also accounts for the difference in grid size relative to demand estimation.

which increases our confidence in the model’s projections.

To account for potentially higher ridership levels in counterfactual networks, we introduce bus capacity constraints. We assume that when a bus has ridership above its capacity, passengers perceive travel time as flowing more slowly. This time-dilation effect grows steeply as ridership per bus exceeds bus capacity. As ridership in excess of capacity is endogenous, we define a *ridership equilibrium* where commuter optimal choices given crowding are consistent with the assumed crowding levels. Appendix A.7.2 describes the model and our computational approach.

The space we are optimizing over is extremely large. Even the number of unique paths is exponential in the number of grid cells, and a network consists of any combination of routes. The number of edge connection configurations is a very rough lower bound. With 418 grid cells and 1,536 possible edges, there are $2^{1,536}$ configurations where each edge is either connected or not – i.e. about 10^{500} possibilities. The allocation of busses to lines adds even more combinatorial dimensions to the optimisation problem, making it infeasible to derive the global welfare-maximizing network through exhaustive search over all possible networks.

5.2 Characterizing Optimal Policies

The Social Planner’s Problem with Idiosyncratic Factors. Consider a social planner who chooses from a finite set of policies (networks, in our setting) $N \in \mathcal{N}$. In a typical application, the set \mathcal{N} is high-dimensional and extremely large. The planner chooses the network N that solves

$$\max_N W(N; \theta) + \epsilon_N,$$

where $W(N; \theta)$ is welfare according to a known, fully specified model, and θ is a vector of parameters. ϵ_N is an idiosyncratic shock for network N , capturing factors not in the model that gives rise to $W(N; \theta)$, such as network construction cost shocks or preferences for specific networks. Assume that the ϵ_N are i.i.d. Gumbel with parameter β .⁴⁶

The probability that network N maximizes social welfare is

$$\pi(N; \theta) = \frac{\exp(\beta W(N; \theta))}{\sum_{N' \in \mathcal{N}} \exp(\beta W(N'; \theta))}. \quad (6)$$

We next show how to characterize optimal networks by taking expectations over π , and how to estimate these objects using a sample of networks drawn from π . We then introduce

⁴⁶In our application, we choose β to be very large, i.e. we assume that the idiosyncratic shocks are small relative to model welfare.

methods that allow us to asymptotically sample from the π distribution.

Optimal Policies: Properties and Local Comparative Statics. Consider a property defined by a function $f(N, \theta)$. For example, this could measure the number of stations for a bus network N , in which case f only depends on N and not on θ . Alternatively, it could measure model-predicted ridership, in which case it depends on both N and θ .

We define the *expected optimal property* as

$$f^*(\theta) = \sum_{N \in \mathcal{N}} \pi(N; \theta) f(N, \theta).$$

This is a weighted average of the property f over all possible networks N , with weights given by the probability that N is optimal. Given a sample of networks N_1, \dots, N_K drawn independently from π , we can estimate $f^*(\theta)$ using the sample counterpart $\hat{f}^*(\theta) \equiv \sum_{k=1}^K f(N_k, \theta)$.

We can compute comparative statics in θ by estimating $\hat{f}^*(\theta')$ for any other value $\theta' \neq \theta$. To compute *local* comparative statics of the expected optimal property, assume that $W(N; \theta)$ is differentiable in θ for all N . Appendix A.7.3 shows that the gradient of the expected optimal property is

$$D_\theta f^*(\theta) = \sum_N \pi(N; \theta) \left(\beta D_\theta W(N; \theta) (f(N, \theta) - f^*(\theta)) + D_\theta f(N; \theta) \right), \quad (7)$$

where $D_\theta W$ and $D_\theta f$ are the gradients of W and f with respect to θ , respectively. The first term captures the change in f^* due to changes in the probabilities $\pi(N; \theta)$, while the second term captures the direct effect on f of changing θ . When f does not depend directly on θ , as for the network statistics we analyze later, the gradient scales linearly with β and with the scale of welfare W . Hence, we focus on the sign of $D_\theta f^*(\theta)$ and compare the local effects of different parameters in θ . Once again, we can estimate $D_\theta f^*(\theta)$ using a sample of networks from the π distribution.

5.3 The Simulated Annealing Algorithm

Building on the Metropolis-Hastings algorithm, a large family of algorithms has been designed to allow sampling asymptotically from probability distributions that are difficult to compute explicitly. The key idea in our case is to construct a Markov chain over the space of all networks, such that the chain’s stationary distribution is exactly (6). A key issue is the speed of convergence (or “mixing”) to the stationary distribution. We use a version of the simulated annealing (SA) algorithm (Nikolaev and Jacobson, 2010; Rothlauf, 2011),

which is related to the Metropolis-Hasting algorithm but has a pre-specified number of steps and begins with “noisier” transitions in the initial phase, which helps ensure that the chain explores widely the space of all networks.

Our algorithm has K steps, an “inverse temperature” parameter that increases exponentially from β_1 to β_K , and begins at an arbitrary network N_1 . Given network N_k at step k , a probabilistic proposal function $\Psi(N' | N_k)$ creates a candidate network N' . This function is user-provided and application-specific. The network N' is accepted with probability:

$$\Pr(N_{k+1} = N' | N_k) = \min \left(1, \frac{\exp(\beta_k W(N')) \Psi(N_k | N')}{\exp(\beta_k W(N_k)) \Psi(N' | N_k)} \right). \quad (8)$$

The candidate network is always accepted if it has higher welfare (with an adjustment for proposal probabilities), and otherwise it is accepted with a probability that is increasing in $W(N')$ and decreasing in β_k .

This algorithm asymptotically samples from the planner’s distribution π .

Proposition 3. *Assume that the final inverse temperature parameter is set to $\beta_K = \beta$, the logit parameter from the social planner’s problem and that the proposal distribution Ψ is irreducible, aperiodic, and satisfies $\Psi(N | N') > 0 \iff \Psi(N' | N) > 0$. As the number of steps $K \rightarrow \infty$, the final state N_K of the simulated annealing algorithm converges in distribution to π , the planner’s distribution over allocations.*

The proof first shows that the stationary distribution of MH with parameter β is exactly π . Second, we sketch the argument that as K grows, the last section of SA approximates an MH chain with parameter β . See Appendix A.7.4 for details.

In the initial phase of the SA algorithm, the low inverse temperature implies that new candidate networks are accepted with high probability. This makes the chain approximately a random walk in the space of networks. This helps the algorithm explore widely the space of networks before converging to the π distribution, thus making SA less likely to get stuck near a local optimum close to the initial network.

The proposal function Ψ is important for the practical success of the SA algorithm. We use small changes such as adding or deleting one stop at the end of a route, local route re-routing, and minor changes in bus allocation between routes, and large changes such as deleting an entire bus route, adding an entirely new bus route, or re-allocating a large share of buses from one route to other randomly chosen routes. See Appendix A.7.5 for details.

Parameters and Implementation. We set $\beta = 10^8$, so idiosyncratic shocks have a small variance and the planner’s choice of network is mostly driven by model welfare. We use

the commuter preference parameters estimated in section 4.3. We use a bus capacity of 84 people per bus based on official TransJakarta data on its bus fleet in early 2020.

We obtain independent asymptotic draws from π using 200 independent SA runs. Each time, we run the SA algorithm for $K = 200,000$ steps.⁴⁷

Since we are sampling from a distribution of networks, we don't expect the algorithm to result in the same network each time. However, given the logit parameter we choose, we expect these networks to be similar in terms of welfare. Figure A.11 shows the evolution of the 200 SA runs, and indeed we find that the distribution of welfare for the final network in each run is concentrated. Results for ridership are similar.⁴⁸

5.4 Results

Figure 5a shows the TransJakarta bus network at the end of our study period. The network is relatively dense in the urban core of the city, with most of the 107 bus lines crossing through the downtown areas. While the core of the city seems well connected by frequent bus services, Jakarta's periphery has few busses passing through, and most locations of the city are not connected at all.

Figure 5c shows one example of a network asymptotically drawn from the planner's distribution. (Figure A.12 shows several other examples.) This network covers more locations compared to the actual TransJakarta network, extending significantly beyond the central Jakarta DKI area. This larger coverage comes at the cost of slightly lower bus frequency in the central area of Jakarta.

Optimal networks are more expansive than the current TransJakarta network. Table 5 reports statistics for the current TransJakarta network in column 1 and for the average over optimal networks in column 2. Panel A reports properties of the network. Optimal networks connect 57% more locations – 66% of all locations in the study area compared to 42% for the current network. The share of location pairs connected by a direct or transfer bus connection also increases considerably, from 12% to 39%. This is achieved with a large network that covers 31% more edges with bus connections, from 544 km to 714 km.⁴⁹

Optimal networks lead to a significant increase in welfare compared to the current net-

⁴⁷The acceptance probability (8) depends on the ratio of proposal kernels Ψ . Calculating this ratio is computationally intensive, so we approximate it by a ratio of 1 when running the simulated annealing algorithm. For most types of network change proposals, this approximation is likely valid. A notable exception is when adding new routes, when this approximation will tend to favor less expansive networks. This effect works against the main result we describe below. See Appendix A.7.5 for details.

⁴⁸Figure A.13 plots the entire distribution of other network statistics.

⁴⁹In our context, efficiency gains do not come with a distributional cost, as coverage also grows for poor areas. 44% of the locations with above-median poverty rate have access to the current TransJakarta network, which grows to 61.6% on average for optimal networks.

work. The average expected utility over all the trips in Jakarta is higher by the equivalent of 21.7 seconds of travel time on the bus, per trip. This may not sound enormous, but recall that this is over all trips, not just bus trips – so given that 1.57% of all trips use the bus with the current network, the equivalent variation improvement is a reduction of 23.1 minutes of bus travel time per initial bus trip. Ridership increases from 1.57% of all trips with the current network to 2.57% on average with optimal networks.⁵⁰ As a consequence, the number of passengers per bus also increases. For example, bus occupancy as experienced by the average passenger doubles. Notably, these numbers are well below the average bus capacity of 84 people per bus (see Figure A.14). Bus trips are longer on average in optimal networks, and optimal networks have a slightly higher share of direct bus trips.

Comparative Statics. We next study how the optimal network shape changes when commuter preferences change (Table 6). This type of exercise helps us assess the sensitivity of our results to the parameters to estimate, and reveals how these micro preference parameters affect optimal network design. We focus on changes in the three key preference parameters that we estimated: the value of wait time α_{wait} , the value of travel time on the bus α_{time} , and the transfer shifter μ_{transfer} . We study the impact of large changes in parameters by sampling 40 optimal networks each for these different parameters using the simulated annealing algorithm.⁵¹ We study the impact of these three parameter changes on three classes of bus network shape measures.

First, we find that network coverage shrinks as wait time costs increase (Table 6, panel A). For example, if commuters dislike wait time twice as much as we estimate, the share of location pairs connected by optimal networks goes from 39% with our baseline parameters to 23%. However, even with double wait time costs – which means that wait time is $4.8\times$ and $8.4\times$ as costly as time on the bus for BRT and non-BRT – optimal networks are still substantially more expansive than the current TransJakarta network.

Second, we consider the speed of bus connections (panel B). For each origin-destination pair with a direct bus connection, we consider the quickest bus connection in the network, and divide its duration by the duration of the quickest possible bus route between the two locations, over all possible paths in the geographic environment. In the current TransJakarta network, direct connections are on average 36% slower than the quickest possible bus connection. The average ratio is similar for the baseline optimal networks, but it goes down

⁵⁰These numbers are based on our estimate of 35.1 million weekday trips between different grid cells, which we obtain from the smartphone location data. This implies a weekday average TransJakarta ridership of 551,036 with the current network.

⁵¹In Table A.9 we report the *local* comparative statics from (7) around the baseline estimated parameters. For most entries, we find the same direction for comparative statics.

significantly, to 26%, when using $2 \times \alpha_{\text{time}}$. This counterfactual also affects the coverage measures in panel A, that is, they lead to more compact networks, although this effect is smaller than that of doubling α_{wait} .

Third, we consider to what extent the network provides direct connections versus offering connections that require a transfer (panel C). In the current TransJakarta network, 21% of all location pairs that are connected by bus also have a direct connection. This number is lower for optimal networks (12%), and it goes up to 16% when we assume that the transfer penalty is higher by the equivalent of 15 minutes of wait time. This counterfactual also increases the (inverse) measure of network speed in panel B, likely because of additional location pairs that are connected by winding direct connections.

None of the parameter vectors we consider leads to network shape parameters that match the current TransJakarta network. However, a higher value of waiting time and a sharper transfer penalty help move the network statistics toward those achieved by the current TransJakarta network.

6 Conclusion

The trend toward centralized urban planning of public transport networks is likely to continue in the coming decades. For example, there are half a dozen BRT systems in African cities, with ten others are under planning or construction ([The World Bank, 2022](#)). As the number of large cities increases, this number will grow, impacting traffic congestion and pollution, commuting and work patterns, and more generally the shape of how cities grow.

In this paper, we studied the design of citywide public transport networks. Launching a public transport route affects the attractiveness of using the network throughout the city. Its impact depends on how the new route increases access, and whether it leads to faster, more frequent, or more direct connections, and how potential riders value those attributes. We used the large expansion of the TransJakarta public bus system between 2016-2020 to separate how specific types of network changes affected bus ridership, and then estimated separate commuter preference parameters for these service quality dimensions by matching the reduced form impacts on ridership.

We find that commuters in Jakarta value bus wait time 2-4 times more than time spent on the bus, transfers in BRT stations do not carry a utility penalty above and beyond the additional wait time and time on the bus, and commuters pay little attention to slow options in their bus network choice set. Characterizing the optimal networks in terms of coverage, speed, and directedness metrics, we found that the optimal bus networks in Jakarta would be more expansive — even holding the resources of the bus company fixed — connecting

more people and location pairs, and boosting overall usage. An optimally designed network that is based on these parameters thus has the potential to dramatically expand ridership and consumer welfare compared to the networks in place today.

References

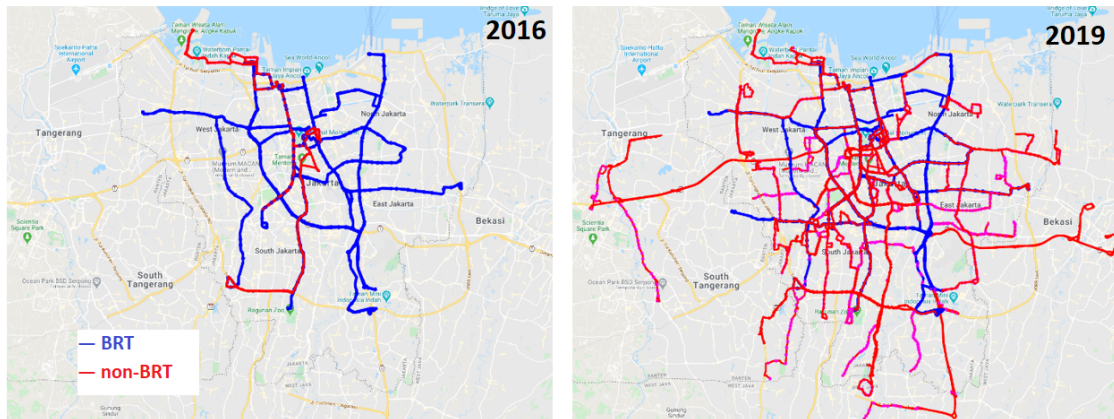
- ABRANTES, P. A. AND M. R. WARDMAN (2011): “Meta-Analysis of UK Values of Travel Time: An Update,” *Transportation Research Part A: Policy and Practice*, 45, 1–17.
- ALDER, S. (2023): “Chinese Roads in India: The Effect of Transport Infrastructure on Economic Development,” *Working Paper*.
- ALLEN, T. AND C. ARKOLAKIS (2022): “The Welfare Effects of Transportation Infrastructure Improvements,” *The Review of Economic Studies*, 89, 2911–2957.
- ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2017): “Measuring the Sensitivity of Parameter Estimates to Estimation Moments,” *Quarterly Journal of Economics*, 132.
- ANTRÀS, P., T. C. FORT, AND F. TINTELNOT (2017): “The Margins of Global Sourcing: Theory and Evidence from US Firms,” *American Economic Review*, 107, 2514–64.
- ARKOLAKIS, C. AND F. ECKERT (2021): “Combinatorial Discrete Choice,” *Working Paper*.
- BALBONI, C. (2021): “In Harm’s Way: Infrastructure Investments and the Persistence of Coastal Cities,” *Working Paper*.
- BALBONI, C., G. BRYAN, M. MORTEN, AND B. SIDDIQI (2021): “Transportation, Gentrification, and Urban Mobility: The Inequality Effects of Place-Based Policies,” *Working Paper*.
- BEN-AKIVA, M. E., S. R. LERMAN, S. R. LERMAN, ET AL. (1985): *Discrete Choice Analysis: Theory and Application to Travel Demand*, vol. 9, MIT press.
- BLANCHARD, P., D. GOLLIN, AND M. KIRCHBERGER (2021): “Perpetual Motion: Human Mobility and Spatial Frictions in Three African Countries,” CEPR Discussion Papers 16661, C.E.P.R. Discussion Papers.
- BUCHHOLZ, N., L. DOVAL, J. KASTL, F. MATĚJKA, AND T. SALZ (2020): “The Value of Time: Evidence from Auctioned Cab Rides,” *NBER Working Paper 27087*.
- CAMERON, A. C., J. B. GELBACH, AND D. L. MILLER (2011): “Robust Inference with Multiway Clustering,” *Journal of Business & Economic Statistics*, 29, 238–249.
- COULOMBEL, N. AND G. MONCHAMBERT (2019): “Diseconomies of Scale and Subsidies in Urban Public Transportation,” *HAL Working Paper, ffhals-02373768v1f*.
- CURRIE, G. (2005): “The Demand Performance of Bus Rapid Transit,” *Journal of Public Transportation*, 8, 41–55.

- DAGANZO, C. (1979): *Multinomial Probit: The Theory and its Application to Demand Forecasting*, Academic Press.
- DAGANZO, C. F. (2010): “Structure of Competitive Transit Networks,” *Transportation Research Part B: Methodological*, 44, 434–446.
- DE CHAISEMARTIN, C. AND X. D’HAULTFOEUILLE (2022): “Two-Way Fixed Effects and Differences-In-Differences with Heterogeneous Treatment Effects: A Survey,” *The Econometrics Journal*.
- FAJGELBAUM, P. D. AND C. GAUBERT (2020): “Optimal Spatial Policies, Geography, and Sorting,” *The Quarterly Journal of Economics*, 135, 959–1036.
- FAJGELBAUM, P. D. AND E. SCHAAL (2020): “Optimal Transport Networks in Spatial Equilibrium,” *Econometrica*, 88, 1411–1452.
- FAN, Y., A. GUTHRIE, AND D. 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.
- FIELBAUM, A., S. JARA-DIAZ, AND A. GSCHWENDER (2016): “Optimal Public Transport Networks for a General Urban Structure,” *Transportation Research Part B: Methodological*, 94, 298–313.
- GABAIX, X. (2019): “Behavioral inattention,” in *Handbook of Behavioral Economics: Applications and Foundations 1*, Elsevier, vol. 2, 261–343.
- GADUH, A., T. GRACNER, AND A. ROTHENBERG (2022): “Life in the Slow Lane: Unintended Consequences of Public Transit in Jakarta,” *Journal of Urban Economics*, 128.
- GOLDSZMIDT, A., J. A. LIST, R. D. METCALFE, I. MUIR, V. K. SMITH, AND J. WANG (2020): “The Value of Time in the United States: Estimates from Nationwide Natural Field Experiments,” *NBER Working Paper 28208*.
- HÖRCHER, D., D. J. GRAHAM, AND R. J. ANDERSON (2017): “Crowding Cost Estimation with Large Scale Smart Card and Vehicle Location Data,” *Transportation Research Part B: Methodological*, 95, 105–125.
- ILIOPOULOU, C., K. KEPAPTSOGLU, AND E. VLAHOGIANNI (2019): “Metaheuristics for the Transit Route Network Design Problem: A Review and Comparative Analysis,” *Public Transport*, 11, 487–521.
- INSTITUTE FOR TRANSPORTATION AND DEVELOPMENT POLICY (2019): “Transjakarta: A Study in Success,” <https://www.itdp.org/2019/07/15/transjakarta-study-success/>, [Online; accessed 9-June-2023].
- JIA, P. (2008): “What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry,” *Econometrica*, 76, 1263–1316.


- JUTPI2 (2019): “Jabodetabek Urban Transportation Policy Integration Project Phase 2 – Annex 2: JABODETABEK Urban Transportation Master Plan,” <https://openjicareport.jica.go.jp/pdf/12356366.pdf>, [Online; accessed 9-June-2023].
- KREINDLER, G. E. (2023): “Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications,” *NBER Working Paper 30903*.
- LARCOM, S., F. RAUCH, AND T. WILLEMS (2017): “The Benefits of Forced Experimentation: Striking Evidence from the London Underground Network*,” *The Quarterly Journal of Economics*, 132, 2019–2055.
- MAJID, H., A. MALIK, AND K. VYBORNY (2018): “Infrastructure Investments, Public Transport Use and Sustainability: Evidence from Lahore, Pakistan,” *Working Paper*.
- MCFADDEN, D. (1974): “The Measurement of Urban Travel Demand,” *Journal of Public Economics*, 3, 303–328.
- MOHRING, H. (1972): “Optimization and Scale Economies in Urban Bus Transportation,” *The American Economic Review*, 62, 591–604.
- NIKOLAEV, A. G. AND S. H. JACOBSON (2010): “Simulated Annealing,” in *Handbook of Metaheuristics*, Springer, 1–39.
- ROTHLAUF, F. (2011): *Design of Modern Heuristics: Principles and Application*, vol. 8, Springer.
- SANTAMARIA, M. (2022): “Reshaping Infrastructure: Evidence from the Division of Germany,” *Working Paper*.
- SMALL, K. A., E. T. VERHOEF, AND R. LINDSEY (2007): *The Economics of Urban Transportation*, Routledge.
- SMERU (2014): “Poverty and Livelihood Map of Indonesia 2015,” <https://www.indonesiapovertymap.org/>, [Online; accessed 9-June-2023].
- THE WORLD BANK (2022): “With Bus Rapid Transit, African Cities Are Riding Toward a Better Future,” *Feature Story, Online [9-June-2023]*, [Online; accessed 9-June-2023].
- TSIVANIDIS, N. (2022): “The Aggregate and Distributional Effects of Urban Transit Infrastructure: Evidence from Bogotá’s TransMilenio,” *Working Paper*.
- TYAZHELNIKOV, V. AND X. ZHOU (2021): “PPML, Gravity, and Heterogeneous Trade Elasticities,” *Working Paper*.
- WARDMAN, M. (2004): “Public Transport Values of Time,” *Transport Policy*, 11, 363–377.
- WEI, F. AND R. B. MACHEMEHL (2006): “Using a Simulated Annealing Algorithm to Solve the Transit Route Network Design Problem.” *Journal of Transportation Engineering*, 132, 122 – 132.
- ZÁRATE, R. D. (2023): “Spatial Misallocation, Informality and Transit Improvements: Evidence from Mexico City,” *Working Paper*.

Figure 1: TransJakarta Network Expansion Since 2016

(a) Expansion of TransJakarta Route Network




(b) Number of BRT and Non-BRT Routes Over Time




Notes: The top panel shows two snapshots of the TransJakarta network in 2016 and 2019. The bottom panel shows the number of active routes over time.

Figure 2: **Bus Ridership** Impacts of BRT and Non-BRT Network Expansions

(a) Event Types 1 and 2: New Direct Route




(b) Event Type 3: Additional Busses (Direct)




Notes: BRT graphs on the left, and non-BRT graph on the right. In the graphs on the first row, the blue, solid lines report monthly coefficients for the first event type (direct route with no change in travel time on the bus). The red, dashed lines report monthly coefficients for the second event type (quicker direct route). The graphs in the second row show the impact of the third event type. All specifications are the event study version of equation (1). 95% confidence intervals are based on standard errors clustered two-way at the origin and destination grid level.

Figure 3: **Aggregate Trip Volume (Lack of) Impacts of BRT and non-BRT Network Expansions**

(a) Event Types 1 and 2: New Direct Route




(b) Event Type 3: Additional Busses (direct)



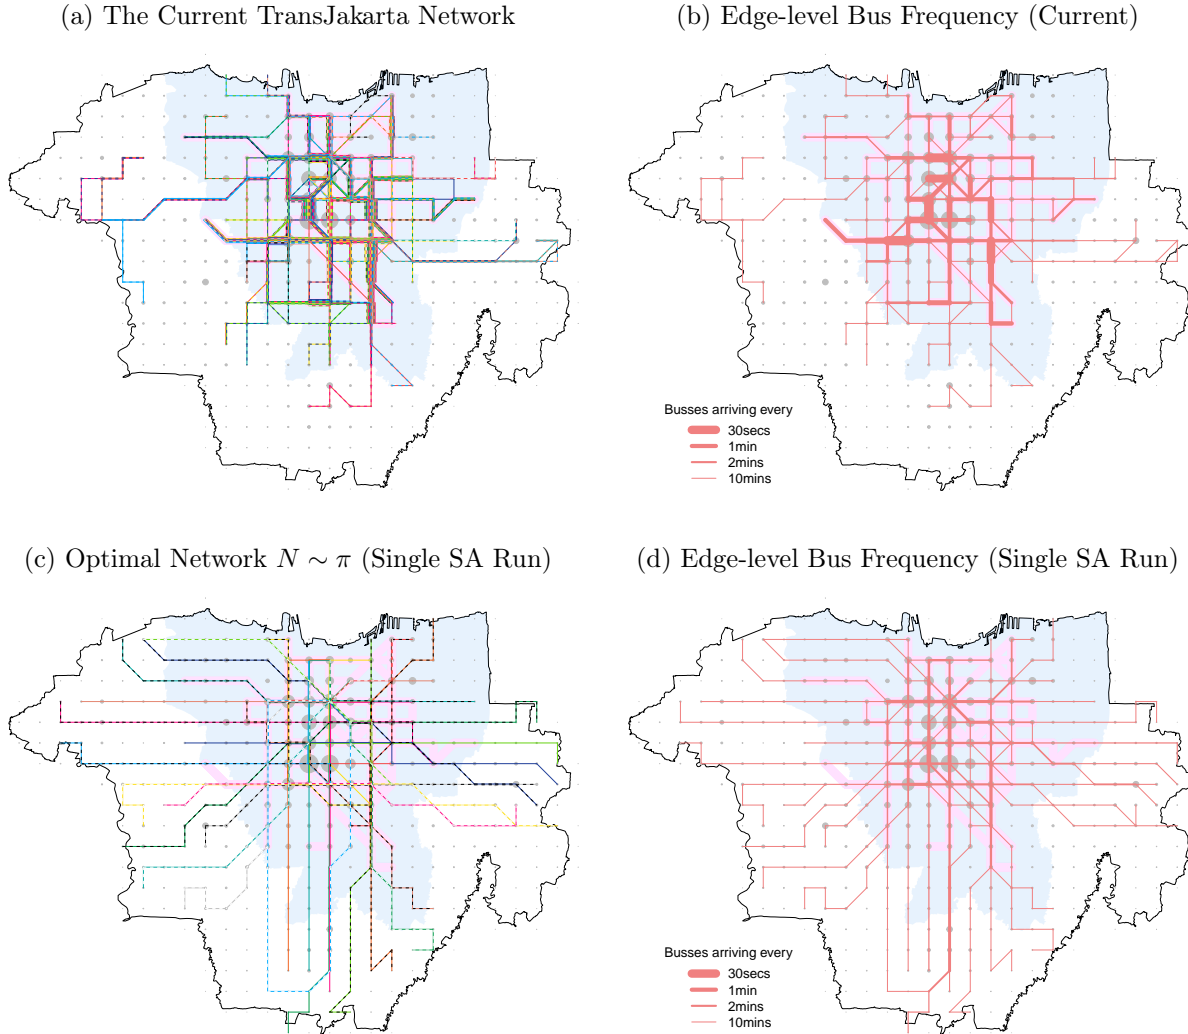
Notes: The outcome is aggregate trip flows from smartphone data. The time sample is restricted to after March 2018. The specifications are otherwise the same as in Figure 2.

Figure 4: Trip Duration Impact of BRT Event 2 (New and Quicker Direct Route)



Notes: The outcome is the log of trip duration (wait time plus time on the bus), measured from linked tap-in and tap-out observations for the same smart card. The sample is restricted to BRT destination stations that are tap-out compliant, defined as stations where tap-out transactions account for at least 30% of all transactions. The specification is the same as for Event 2 in Figure 2. 95% confidence intervals are based on standard errors clustered two-way at the origin and destination grid level.

Figure 5: Current TransJakarta Network and Optimal Network Example



Notes: This figure depicts the current TransJakarta network and one of the networks N obtained using the simulated annealing (SA) algorithm, asymptotically sampled from the planner's distribution π . The study area (black outline) includes the Jakarta DKI area (light blue) as well as the urban districts of Tangerang, South Tangerang, Depok, and Bekasi. Different bus lines are printed in different colors. Edges connected through BRT dedicated bus lanes are denoted by bright pink underlay. Panel (a) plots the current TransJakarta network (107 bus lines). Panel (b) shows bus frequency on each edge in the current TransJakarta network. To construct this graph, for each directed edge we compute the total bus arrival over all routes that share that edge. For each edge, we then take the average over the two directions. Panels (c) and (d) repeat the exercise for one of the 200 SA networks, chosen randomly.

Table 1: Impact of Network Expansion on Travel Time, Wait Times, and **Bus Ridership**

(a) BRT

	log Min Travel Time			log Bus/hr (origin)			Bus Ridership		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
E1: New Direct Line	0.036*** (0.007)			-0.051* (0.026)			0.158*** (0.021)		
E2: New Direct Line (quicker)		-0.294*** (0.027)			0.070 (0.036)			0.268*** (0.037)	
E3: Additional Busses			-0.027* (0.011)			0.320*** (0.035)			0.092*** (0.023)
Orig × Dest FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Orig × Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Dest × Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	OLS	OLS	OLS	OLS	OLS	OLS	PPML	PPML	PPML
Mean outcome	33.2	44.2	12.9	35.2	27.0	20.0	111.3	76.9	210.4
Median outcome	31.0	42.4	10.6	28.7	20.7	16.9	55.9	39.1	94.6
Unique origin x destination pairs	18,433	18,449	5,319	18,433	18,449	5,319	18,433	18,449	5,319
Unique origins	148	148	148	148	148	148	148	148	148
<i>N</i>	3,154,672	3,143,019	793,965	3,154,672	3,143,019	793,965	3,154,672	3,143,019	793,965
<i>R</i> ²	0.981	0.980	0.998	0.923	0.921	0.997			

(b) Non-BRT

	log Min Travel Time			log Bus/hr (origin)			Bus Ridership		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
E1: New Direct Line	0.030** (0.009)			0.292*** (0.044)			1.016*** (0.153)		
E2: New Direct Line (quicker)		-0.716*** (0.072)			0.205*** (0.030)			1.320*** (0.267)	
E3: Additional Busses			-0.089*** (0.021)			0.425*** (0.015)			0.450** (0.141)
Orig × Dest FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Orig × Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Dest × Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	OLS	OLS	OLS	OLS	OLS	OLS	PPML	PPML	PPML
Mean outcome	34.8	49.3	13.0	11.6	11.2	7.1	19.2	7.5	26.0
Median outcome	24.8	47.2	11.0	7.8	8.5	6.3	0.0	0.0	0.0
Unique origin x destination pairs	2,588	2,652	1,263	2,588	2,652	1,263	2,588	2,652	1,263
Unique origins	57	57	64	57	57	64	57	57	64
<i>N</i>	306,722	313,994	144,763	306,722	313,994	144,763	306,722	313,994	144,763
<i>R</i> ²	0.992	0.987	0.998	0.967	0.967	0.992			

Notes: The two panels report results for BRT and non-BRT, respectively. Columns 1-6 capture the “first stage” impacts of the relevant events and are estimated using OLS. The bus ridership effects (columns 7-9) are estimated by PPML. “Min Travel Time” captures the time on the bus (excluding wait time) of the quickest route between an origin and a destination, given the routes available in the TransJakarta network at that time. “Buses per Hour” measures the total number of buses arriving at the origin grid over all the routes connecting the origin and destination. (The set of routes is restricted to direct routes for event 3.) The coefficients capture the average effect in the first 10 months after the respective event. The origin and destination data is aggregated at the level of 1km hexagonal grids. Standard errors clustered two-way at the origin and destination grid level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Impact of BRT and Non-BRT Network Expansions on **Aggregate Trip Volume**

(a) BRT

	log Min Travel Time			log Bus/hr (origin)			Bus Ridership			All trips		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
E1: New Direct Line	0.057*** (0.008)			-0.158*** (0.031)			0.108*** (0.016)			-0.008 (0.051)		
E2: New Direct Line (quicker)		-0.214*** (0.026)			-0.108*** (0.030)			0.208*** (0.032)				-0.073 (0.049)
E3: Additional Busses			-0.006 (0.036)			0.221*** (0.058)				0.046 (0.041)		0.166 (0.117)
Orig × Dest FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Orig × Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Dest × Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	OLS	OLS	OLS	OLS	OLS	OLS	PPML	PPML	PPML	PPML	PPML	PPML
Mean outcome	32.6	42.0	13.4	37.1	32.9	13.6	122.9	93.7	195.5	1895.5	1729.5	10117.3
Median outcome	30.6	42.2	10.9	32.5	27.7	11.6	61.9	54.8	153.7	0.0	0.0	6412.4
Unique origin x destination pairs	13,802	13,675	2,497	13,802	13,675	2,497	13,802	13,675	2,497	13,802	13,675	2,497
Unique origins	124	124	89	124	124	89	124	124	89	124	124	89
N	1,184,781	1,172,746	197,039	1,184,781	1,172,746	197,039	1,184,781	1,172,746	197,039	1,184,781	1,172,746	197,039
R ²	0.986	0.985	0.999	0.948	0.948	0.999						

(b) Non-BRT

	log Min Travel Time			log Bus/hr (origin)			Bus Ridership			All trips		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
E1: New Direct Line	0.027* (0.011)			0.253*** (0.048)			0.999*** (0.199)			-0.042 (0.060)		
E2: New Direct Line (quicker)		-0.736*** (0.094)			0.200*** (0.036)			1.337*** (0.322)				0.031 (0.060)
E3: Additional Busses			-0.097** (0.031)			0.365*** (0.023)				0.425* (0.176)		-0.177* (0.073)
Orig × Dest FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Orig × Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Dest × Week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Estimator	OLS	OLS	OLS	OLS	OLS	OLS	PPML	PPML	PPML	PPML	PPML	PPML
Mean outcome	35.9	44.3	12.4	12.0	12.0	7.4	13.0	7.6	40.7	3948.6	6060.7	6880.5
Median outcome	24.5	43.6	10.8	7.8	8.5	6.9	0.0	0.0	0.0	635.5	956.9	1017.5
Unique origin x destination pairs	1,906	1,950	860	1,906	1,950	860	1,906	1,950	860	1,906	1,950	860
Unique origins	49	49	53	49	49	53	49	49	53	49	49	53
N	156,656	160,409	61,377	156,656	160,409	61,377	156,656	160,409	61,377	156,656	160,409	61,377
R ²	0.995	0.989	0.999	0.976	0.976	0.997						

Notes: This table reports the impact on the aggregate trip volume as measured from smartphone data (columns 10-12). The outcome is aggregate trip volume from smartphone data. The time sample is restricted to after March 2018. For comparison, columns 7-9 report the impact on bus ridership in the same sample. Standard errors clustered two-way at origin and destination grid level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Impact of BRT Event Type 2 on Total Trip Duration

	<u>log Min Travel Time</u>	<u>log Bus/hr (origin)</u>	<u>Bus Ridership</u>	<u>log Trip Duration</u>
	(1)	(2)	(3)	(4)
E1: New Direct Line	0.055*** (0.012)	-0.212*** (0.038)	0.110*** (0.025)	-0.003 (0.005)
E2: New Direct Line (quicker)	-0.284*** (0.038)	0.165*** (0.045)	0.071* (0.036)	-0.057*** (0.012)
Orig × Dest FE	Yes	Yes	Yes	Yes
Orig × Week FE	Yes	Yes	Yes	Yes
Dest × Week FE	Yes	Yes	Yes	Yes
Estimator	OLS	OLS	PPML	OLS
Median outcome (E1)	30.1	33.7	68.4	54.2
Median outcome (E2)	40.1	27.8	61.2	57.6
Unique origin x destination pairs	13,455	13,455	13,455	13,455
Unique origins	148	148	148	148
N	595,685	595,685	595,685	595,685
R^2	0.991	0.970		0.881

Notes: Column 4 in this table reports the impact on log trip duration (wait time plus time on the bus) of BRT Event 2. The outcome is trip duration measured from linked tap-in and tap-out observations for the same smart card. The sample is restricted to BRT destination stations that are tap-out compliant, defined as stations here tap-out transactions account for at least 30% of all transactions. We jointly estimate Events 1 and 2, so the “E2: New Direct Line (quicker)” coefficient captures the *additional* effect for Event 2. First-stage and bus ridership impacts are reported in columns 1-3. Standard errors clustered two-way at the origin and destination grid level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4: Estimated Demand Model Parameters

	(1) BRT	BRT	(2) non-BRT
Wait time $\alpha_{\text{wait}}/\alpha_{\text{time}}$	2.4 [1.8, 4.3]	2.4 [1.9, 4.3]	4.2 [3.2, 6.6]
Travel time α_{time}	1	1	1
Transfer shifter $\mu_{\text{transfer}}/\alpha_{\text{wait}}$ (minutes)	2.3 [-2.0, 4.5]	2.3 [-2.0, 4.4]	-1.1 [-5.3, 0.9]
Attention Cutoff η	1.34 [1.25, 1.45]	1.34 [1.25, 1.45]	1.44 [1.43, 1.52]
Logit parameter σ	0.060 [0.023, 0.113]		0.060 [0.023, 0.105]
<i>Moments:</i>			
BRT events (1-3)	Yes		Yes
BRT event 2 trip duration	Yes		Yes
non-BRT events (1-3)			Yes

Notes: We use a classical minimum distance with the optimal weighting matrix, and 20 random initial conditions. To construct the 95% confidence intervals, we re-estimate the model 100 times. Each time, we target a data moment vector $\widehat{m}^k = \widehat{m} + \varepsilon^k$, where ε^k is randomly drawn from the multivariate normal distribution $\mathcal{N}(\mathbf{0}, \widehat{\Omega})$. $\widehat{\Omega}$ is the estimated variance-covariance matrix of the reduced form analysis, jointly estimated in a seemingly unrelated regression framework. During this procedure, we use 4 random initial conditions for each estimation.

Table 5: Properties of Current and Optimal Networks

Statistic	Current Network	Baseline Optimal	Difference Baseline – Current
Panel A: Network and access properties			
Locations with a bus station (share)	0.42	0.66	0.24 (0.0015)
Location pairs connected by bus (share)	0.12	0.39	0.27 (0.0018)
Trips with access to bus (share)	0.73	0.91	0.18 (0.00094)
Total network mileage (km)	543.9	713.9	169.9 (2.0)
Number of bus lines	107	27.76	-79.24 (0.2)
Panel B: Ridership equilibrium properties			
Welfare	-0.0031	-0.0027	0.00036 (0.00000065)
Ridership (%)	1.57	2.57	1.0063 (0.0028)
Average bus occupancy (pax/bus)	13.62	33.00	19.39 (0.034)
Average experienced bus occupancy (pax/bus)	23.82	46.1	22.27 (0.056)
Average bus trip duration (hrs)	0.59	0.81	0.23 (0.00073)
Transfer trips (share of all bus trips)	0.76	0.71	-0.049 (0.00054)

Notes: This table reports characteristics of the TransJakarta network (column 1) and of optimal networks, on average over 200 independent simulated annealing (SA) runs (column 2). Panel A reports network properties, while Panel B reports properties of the ridership equilibrium. “Trips with access to bus” denotes the share of all trips (measured using smartphone data and assumed fixed in counterfactuals) that can be completed using (have access to) the bus network. “Total network mileage” computes the total length of all edges traversed by the network. The difference in welfare in Panel B in equivalent variation terms is a reduction of 23.1 minutes of bus travel time for each bus trip in the current network. “Ridership” corresponds to the share of total trips across Jakarta using the bus network. “Average bus occupancy” is the occupancy level of the average bus, assuming busses are active for 17 hours each day (5AM-10PM) and using the model to predict total bus ridership for each route and pair of consecutive stations. “Average experienced occupancy” weights this number by ridership itself to get at the average amount of other passengers on a bus for a typical passenger. In column 3, in parentheses, bootstrapped standard errors are obtained using the 200 independent runs.

Table 6: Optimal Networks Comparative Statics

Statistic	Current Network	Baseline Optimal	Comparative Statics (Large Changes)		
			Wait time $2 \times \alpha_{\text{wait}}$	Time on Bus $2 \times \alpha_{\text{time}}$	Transfer $\mu_{\text{transfer}} - 15\text{mins}$
Panel A: Coverage measures					
Locations with a bus station (share)	0.42	0.66 (0.0015)	0.5 (0.0032)	0.59 (0.0024)	0.65 (0.0031)
Location pairs connected by bus (share)	0.12	0.39 (0.0018)	0.23 (0.0028)	0.3 (0.0025)	0.41 (0.004)
Total network mileage (in km)	543.95	713.87 (2.032)	498.9 (3.15)	664.89 (3.24)	744.068 (3.73)
Panel B: Speed measures					
Bus time relative to quickest possible bus route	1.36	1.35 (0.01)	1.3 (0.024)	1.26 (0.021)	1.48 (0.029)
Panel C: Directness measures					
Connected directly (share of all connected pairs)	0.21	0.12 (0.0006)	0.14 (0.0016)	0.13 (0.0011)	0.16 (0.0018)

Notes: This table reports how network shape measures (rows) change as a function of parameter changes (columns). Columns 3-5 report characteristics of optimal networks achieved with different sets of preference parameters, each obtained from 40 independent simulated annealing runs with altered parameter vectors: double value of wait time ($\alpha_{\text{wait}}^{\text{BRT}}$ and $\alpha_{\text{wait}}^{\text{non-BRT}}$) relative to our estimate, double value of time on bus (α_{time}) relative to our estimate, and an additional penalty of taking a transfer ($\mu_{\text{transfer}}^{\text{BRT}}$ and $\mu_{\text{transfer}}^{\text{non-BRT}}$) equivalent to 15 minutes of wait time. See Table 5 for variable definitions. To construct “bus time relative to quickest possible bus route,” for each origin-destination pair that is directly connected, we consider the quickest direct bus connection in the network, and divide its duration by the duration of the quickest possible bus route between the two locations in the entire geographic environment. We then average this ratio over all directly connected origin-destination pairs. “Connected directly (share of all connected pairs)” computes the share of all location pairs that are connected by bus, for which a direct connection also exists. Bootstrapped standard errors are in parentheses. Table A.9 reports the *local* comparative statics version of this table.