Traffic Jam: A Passenger-Centred Diagnostic Framework for Congestion, Public Transport Operations and Institutional Coordination in Nairobi, Kenya

Allan Robinson

Nalla Labs, Nairobi, Kenya

Author Note

Allan Robinson https://orcid.org/0009-0002-4515-6159

This article is a preprint and has not been peer reviewed. The study was not preregistered; the field evaluation it proposes will be preregistered before data collection. Traffic Jam is an open-source research project. The paper is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0), and the model code is released under the MIT License. The paper, the model code and its outputs are archived on Zenodo and are available at nallalabs.xyz or on request from hello@nallalabs.xyz. Contributions from the public, including data, field observations, replications, corrections and critiques, are welcome. This research received no funding and is not sponsored by any institution. The author thanks an independent reviewer for comments on an earlier draft (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026), and the public agencies, researchers and data providers whose openly available work made this analysis possible; none of them reviewed or endorsed the paper.

Correspondence concerning this article should be addressed to Allan Robinson, Nalla Labs, Nairobi, Kenya. Email: allanrobinson@nallalabs.xyz. General inquiries and contributions: hello@nallalabs.xyz.

Abstract

Urban congestion in Nairobi is commonly addressed through road building, enforcement and transport technology, yet recurring delays may also arise from interactions among road capacity, public transport operations, traffic management, road-user behaviour and institutional coordination. This paper develops a passenger-centred framework for identifying where delay originates rather than only where it is observed, and sets out a research programme to test it. Using secondary data from the Japan International Cooperation Agency's 2023 household survey, TomTom speed indicators and national statistics, a wage-based model estimates the economic cost of weekday peak congestion in the Nairobi Metropolitan Area at about KSh 46.3 billion a year under central assumptions; a Monte Carlo analysis gives a 90% uncertainty interval of KSh 33.6–70.8 billion, with the value of travel time and peak delay accounting for most of the variance. An illustrative stage seat-supply model shows how delays on inbound vehicle trips could produce passenger queues at central stages; it requires field calibration. A review of congestion economics, from Pigou and Vickrey to recent field experiments, and of Nairobi's monitoring systems informs a proposed diagnosis-first pilot with a named operating owner, comparison sites and a pre-specified difference-in-differences evaluation, and a staged design for an AI-assisted coordination system in which decisions remain with responsible institutions. No primary fieldwork has yet been conducted. The contributions are a preliminary estimate, a diagnostic framework, an operational model, and a protocol for validation and replication in other cities.

Keywords: urban congestion, passenger travel time, Nairobi, paratransit, matatu, traffic diagnosis, institutional coordination, congestion pricing, AI-assisted transport management

Traffic Jam: A Passenger-Centred Diagnostic Framework for Congestion, Public Transport Operations and Institutional Coordination in Nairobi, Kenya

The Nairobi Metropolitan Area generates about 14.1 million person-trips on an average weekday. Of these, 37.2% are made on foot and 32.2% by matatu, the privately operated minibuses that form the backbone of the city's public transport (JICA Expert Team, 2025). Average road speeds in Nairobi City fell from 40.0 km/h in 2013 to 21.2 km/h in 2023, and from 27.5 km/h to 8.3 km/h at peak (JICA Expert Team, 2025). In 2025 a 10 km drive took an average of 26 minutes 47 seconds and almost 36 minutes in the evening peak (TomTom, 2026). The cost of this delay falls on commuters, operators, businesses and the public purse, and it is growing: motorcycle registrations in Kenya doubled between 2024 and 2025 ("Bikes Boom Drives Spike," 2026), and the Japan International Cooperation Agency (JICA) projects that average speeds on major roads will fall from 18.2 km/h in 2023 to 6.8 km/h by 2045 if nothing is done (JICA Expert Team, 2025).

The dominant response has been to add or upgrade infrastructure and technology. Announced commitments to a three-phase intelligent transport system covering 210 junctions, a separate concessional loan for traffic systems, camera-based enforcement and a bus rapid transit (BRT) line total between KSh 112.7 billion for clearly distinct projects and KSh 136.7 billion if a loan that may overlap with them is included ("Kenya Signs Ksh24 Billion Loan," 2025; Kenya Urban Roads Authority [KURA], 2026; "KURA Signs Deal," 2024; "NTSA to Install 1,000 Speed Cameras," 2026; "Team Europe to Build Sh47bn Green BRT," 2023; The National Treasury, 2025). These investments may well be necessary. Yet JICA's own forecasts show speeds declining even when every planned road and BRT line is built (JICA Expert Team, 2025), consistent with the international finding that new road capacity tends to fill with new traffic (Duranton & Turner, 2011). Capacity alone may not explain why some junctions fail every evening while others with similar volumes do not, or why passengers can wait a long time for a seat at a central stage while vehicles are delayed elsewhere on the route.

This study investigates peak-hour congestion in Nairobi as a potentially multi-causal problem arising from interactions among infrastructure capacity, public transport operations, traffic enforcement, road-user behaviour and institutional coordination. Rather than assuming that coordination failures are the primary cause, it develops a framework for identifying and measuring the relative contribution of these mechanisms, and a passenger-centred diagnostic approach through which recurring delays can be traced to their origins, assigned to responsible actors and evaluated using measurable outcomes. The working hypothesis is that part of recurring peak delay reflects coordination failures: problems that recur daily are seen by many actors but owned by none, and the information that could locate their cause is scattered across agencies, operators and platforms. The hypothesis would be weakened if field diagnosis showed that recurring delays are explained mainly by demand exceeding road and fleet capacity, with few recurring problems that an identifiable owner could fix.

Problem Statement

An independent advisory review of an earlier draft (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026) recommended stating the problem in passenger terms rather than in terms of vehicles or institutions. We adopt its suggested wording:

> *Passengers on the selected route experience recurring and unpredictable peak delays. The first phase will identify where their time is lost, estimate how much is avoidable and determine who has the authority and capacity to act.*

The problem is measured through passenger outcomes: waiting time at stages, end-to-end journey time, the variability of journey time (the buffer-time index) and, for institutions, the time taken to act on recurring problems. This wording separates what passengers experience, what causes it (vehicle supply, stage operations, junction capacity, incidents and behaviour) and what institutions can change, because the place where a queue is visible is not necessarily where it originates.

Research Questions

The paper addresses five questions:

1. RQ1: What is the order of magnitude of the economic burden of weekday peak congestion in the Nairobi Metropolitan Area, and which assumptions drive its uncertainty?
2. RQ2: Which mechanisms and institutional arrangements could produce recurring peak delay, and how can the location of a queue be distinguished from its origin?
3. RQ3: What do Nairobi's previous interventions and comparable cities show about which mechanisms an intervention must address, and under what conditions lessons transfer to Nairobi?
4. RQ4: How should a first intervention be diagnosed, owned and evaluated so that its effect, and separately the contribution of coordination, can be measured credibly?
5. RQ5: Do the institutions responsible for Nairobi's traffic share a coordinated system that monitors conditions, assigns recurring problems to named owners and tracks whether they act, and how could artificial intelligence (AI) help such a system detect, diagnose and resolve congestion while decisions and accountability remain with people?

These questions define a research programme rather than a single study. The present paper is its foundational framework and protocol: it answers RQ1 provisionally from secondary data, illustrates the mechanism behind RQ2 with a model, reviews the documentary evidence for RQ3 and RQ5, and designs the fieldwork that RQ2, RQ4 and RQ5 require. To keep the strength of each claim visible, the paper labels claims by type: *estimated* (computed from data under stated assumptions), *illustrated* (shown within a model under stated assumptions), *inferred* (an interpretation of documentary evidence) and *proposed* (to be tested). Table 1 summarises what the paper provides for each question.

Table 1

What This Paper Provides for Each Research Question

RQ	What this paper provides	Type of claim	What remains to be done
RQ1	A preliminary, reproducible estimate of the economic burden, with scenario and Monte Carlo uncertainty analysis	Estimated from secondary data	Mode-specific delay data; a Nairobi value-of-time survey
RQ2	A stage seat-supply model and two hypotheses (H1, H2) about where delay originates	Illustrated within models	Field calibration and competing-mechanism diagnosis
RQ3	A review of Nairobi's interventions and of the transferability of international cases	Inferred from documents	Systematic review; interviews with those involved
RQ4	A diagnosis-first pilot protocol with a pre-specified evaluation design	Proposed	Baseline data, implementation and evaluation
RQ5	An assessment of documented monitoring systems and a staged design for AI-assisted coordination	Inferred and proposed	Institutional interviews; technical testing against baselines

Contribution and Scope

The paper makes seven contributions. First, it provides a transparent, reproducible preliminary estimate of the economic cost of peak congestion, with published code, a Monte Carlo uncertainty analysis and a critique of its own assumptions. Second, it formalises the Central-Stage Seat Supply and Passenger Queue Model, which shows how inbound delays could generate passenger queues at central stages. Third, it proposes a passenger-centred diagnostic framework that weighs competing mechanisms rather than assigning each queue to a single cause. Fourth, it specifies a pilot protocol with interview and observation instruments, a pre-specified difference-in-differences design and provisional sample sizes. Fifth, it synthesises congestion economics, from Pigou to recent field experiments, and states which claims this paper borrows, finds consistent with

Nairobi evidence, qualifies or contests. Sixth, it assesses whether Nairobi's institutions share a coordinated system for monitoring traffic and holding owners to account, and specifies the long-term aim of the programme: an AI-assisted traffic coordination system for Kenya in which AI measures, detects, diagnoses, routes and evaluates while named institutions decide and answer for results. Seventh, it offers a replication framework so that other cities can apply and test the approach.

The scope is deliberately limited. The analysis covers weekday peaks in the Nairobi Metropolitan Area and relies on secondary data; no primary fieldwork has yet been conducted. The paper is therefore an evidence synthesis, a set of models and a research protocol, not an evaluation. Its estimates are intended to size the problem and to guide the design of a first test, not to predict the benefit of any specific project. The programme plans four follow-on papers, each with its own data: the cost of congestion with mode-specific measurements; where passenger queues originate, with field validation of the stage model; institutional coordination and accountability; and AI-assisted coordination tested against simpler methods.

The remainder of the paper is organised as follows. The Literature Review sets out the approach to the review and synthesises evidence on traffic capacity, congestion economics, Nairobi's travel patterns, matatu operations, governance, monitoring technology and the transferability of international interventions. The Method section maps each research question to its data and methods and describes the cost model and uncertainty analysis, the economists' models, the stage model, the diagnostic, evaluation and systems-assessment protocols, and ethics. The Results report the estimates and model outputs. The Discussion is organised by research question and closes with the paper's contribution to theory, its limitations and their mitigations, and open problems. The final section draws conclusions and staged recommendations.

Literature Review

Approach to the Review

This is a narrative, theme-based review, not a systematic review. Sources were identified between 18 and 23 November 2026 through targeted searches of the open web, scholarly repositories and indexes such as RePEc, and the websites of Kenyan agencies and development partners, using combinations of terms for urban congestion, traffic-flow models, congestion pricing, value of travel time, paratransit and matatu operations, bus bunching and headways, transport governance and institutional coordination, intelligent transport systems and AI in traffic management, and named Nairobi interventions; and by following citations in key works. Bibliographic databases such as Scopus and Web of Science were not searched, and only English-language sources were included, so relevant studies may have been missed. Sources were graded in four tiers (see Method): original datasets, peer-reviewed studies and official documents; credible institutional evidence; secondary reporting; and unverified leads. Peer-reviewed research and original institutional data were preferred, and news reports were used mainly to document policy events, investments and claims requiring verification. The search themes and example search strings are listed in Supplementary File S3, and a systematic review using bibliographic databases is listed among the open problems. The review is organised by theme, and each theme closes with a short synthesis of what the evidence does and does not show for Nairobi.

Traffic Capacity, Mixed Traffic and Nonlinear Congestion

Classical traffic-flow theory begins with the relationship between density, speed and flow. Greenshields (1935) observed an approximately linear decline in speed with density, which implies a parabolic flow–density curve: flow rises with density up to a critical point and then falls. Past that point, adding vehicles reduces the number that pass a section per hour, so a queue, once formed, can sustain itself. For planning purposes the most widely used summary is the volume–delay function of the Bureau of Public Roads (1964), in which travel time grows with a power of the volume-to-capacity ratio. Its textbook parameters produce a gentle curve until volumes approach capacity and a steep rise thereafter. For Nairobi, JICA adopted much steeper tentative parameters, reflecting mixed traffic and frequent stopping (JICA Expert Team, 2025).

At junctions, delay depends on signal timing as much as on volume. Webster (1958) derived an approximately optimal fixed cycle length from the lost time per cycle and the sum of critical flow ratios, a result that still underpins signal design. Deterministic queueing analysis adds a further insight that is central to this paper: when arrivals exceed capacity for a period, total delay grows with the square of the period's length. Measures that shorten an overload, rather than eliminate it, can therefore deliver disproportionate gains.

Taken together, traffic-flow theory indicates that delay rises nonlinearly as roads approach capacity and that short overloads can create disproportionately long queues. Most of this theory was developed for homogeneous motorised traffic, however. Its applicability to Nairobi's mix of matatus, buses, cars, motorcycles and pedestrians, with frequent stopping and loading, is uncertain, which is why JICA's tentative parameters differ sharply from textbook values and why corridor-specific calibration is listed among the open problems.

Congestion Externalities, Equilibrium, Pricing and Induced Demand

Economists have studied road congestion for more than a century, and their work is the natural benchmark for any claim about what causes a jam and what would relieve it. This section reviews the main contributions in roughly chronological order. For each, we state the claim and, in Table 2, whether this paper borrows it, finds it consistent with the available Nairobi evidence, qualifies it, or contests it for Nairobi. We also examine a popular summary of this literature that circulates widely: that nobody in a jam is doing anything wrong, that individual tactics such as leaving earlier or finding a shortcut are copied until their advantage disappears, and that relief therefore requires a change of structure rather than more individual effort.

Pigou and Knight: An Unpriced Externality or a Missing Owner?

Pigou (1920) used the example of two roads, one good but narrow and liable to congestion and the other uncongested but inferior, to argue that drivers who choose the congested road ignore the delay they impose on others, so that a charge on that road would raise total welfare. Knight (1924) challenged the argument: if the congested road had an owner, the owner would set a toll that rations its use efficiently, so the problem lies in the absence of ownership rather than in private choice as such. The two positions frame a choice that recurs throughout the literature, between a public charge and a responsible owner, and both are relevant in Nairobi, where a privately operated expressway runs beside streets and stages that nobody owns in the economic sense.

Walters: Measuring the Gap Between Private and Social Cost

Walters (1961) converted the engineer's speed–flow relationship into an average-cost curve and derived from it the marginal social cost of one more vehicle. Applying the method to highways in the United States, he concluded that efficient prices were generally much higher than prevailing charges, which implied higher fuel taxes and special tolls in congested areas. Walters' method underlies the marginal external delay and the illustrative Pigouvian toll calculated in this paper (Equation 8).

Vickrey: Time of Day, Bottlenecks and Electronic Charging

William Vickrey extended the argument from how many people use a road to when they use it. In 1951 he studied subway fares for New York City's Mayor's Committee for Management Survey, and in 1959 he studied traffic congestion in Washington, DC (Drèze, 1998). His 1959 statement to a congressional committee proposed charging for street use by time and place, recorded by electronic identifiers carried in vehicles (Vickrey, 1959/1994), decades before transponder tolling became routine. His 1963 paper opened with the proposition that pricing in no other major sector was "so irrational, so out of date, and so conducive to waste" as in urban transportation (Vickrey, 1963, p. 452). In 1969 he introduced the bottleneck model, in which commuters choose when to travel by trading the time they spend queuing against the cost of arriving earlier or later than they wish (Vickrey, 1969).

The bottleneck model shows, under its assumptions, why the popular claim that a jam cannot be beaten by effort can hold. In its equilibrium every commuter bears the same private cost whatever their departure time: those who leave early avoid part of the queue but arrive too early, and those who travel at the most convenient time queue longest (Equation 12). No individual can lower their cost by retiming, and any advantage one commuter discovers is soon shared by others. Within the model the popular claim therefore holds, with one refinement; whether Nairobi's travellers behave as the model assumes is an empirical question, and matatu passengers in particular do not choose departure times freely. The equilibrium is also not efficient, because each commuter's place in the queue delays those behind, and a charge that varies through the peak can remove that inefficiency.

Vickrey shared the 1996 Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel with James Mirrlees. The prize recognised contributions to the economic theory of incentives under asymmetric information (in Vickrey's case including his work on auctions and on income taxation) rather than congestion pricing as such (Royal Swedish Academy of Sciences, 1996). Vickrey died on 11 October 1996, three days after the prize was announced (Drèze, 1998). Popular accounts that present the prize as a reward for watching traffic therefore compress his record. His transport work is better described as the application, over five decades, of one idea, that a price can carry information and change behaviour, to a problem most of his contemporaries treated as one of construction.

Downs: Triple Convergence

Downs (2004) explained why improvements to a congested road rarely survive the peak. When a road is widened or a faster route appears, travellers converge on it from other routes, other times of day and other modes until peak congestion returns close to its former level. Triple convergence is the mechanism behind the popular observation that shortcuts and early departures are copied, and its empirical counterpart is the finding that vehicle-kilometres travelled rise almost one-for-one with road capacity (Duranton & Turner, 2011). Downs also argued that peak congestion performs a useful function, rationing scarce road space among people who all want to travel at the same time. Triple convergence is consistent with JICA's forecast that speeds keep falling even with all planned projects (JICA Expert Team, 2025), and the Results illustrate it numerically. We qualify the rationing argument: part of Nairobi's peak delay is not rationing at all but capacity lost to vehicles blocking junctions, stages spilling into live lanes and incidents left uncleared, and capacity lost in this way rations nothing.

Small, Arnott, de Palma and Lindsey: The Bottleneck Formalised

Small (1982) estimated how commuters in the San Francisco Bay Area trade travel time against arriving early or late at work, and found that lateness is penalised much more heavily than earliness. His estimates are often approximated as an early-arrival cost of about half, and a late-arrival cost of about twice, the value of travel time (Small, 2015). Arnott et al. (1993) and related work formalised Vickrey's model and derived its central results. Without a charge, total cost is split equally between queuing and schedule delay. A charge that rises and falls through the peak exactly as the queue would have done eliminates queuing, halves total social cost and transfers the other half to the road authority as revenue, but it leaves each commuter's private cost unchanged before that revenue is spent (Small, 2015). The welfare gain is real, but commuters only experience it if the revenue returns to them as better services, lower taxes or rebates. We borrow the model and its parameter ratios for an illustrative Nairobi calculation, and we argue that this revenue result, more than the efficiency result, would decide whether any charge in Nairobi was accepted.

The Vickrey Bathtub and Hypercongestion

Vickrey also described downtown traffic as a bathtub: vehicles flow in as commuters enter the central area and drain out as their trips are completed. He did not publish the model. Arnott (2013) formalised it, and Arnott (2015) summarised it for a policy audience. The rate at which trips are completed rises with traffic density up to a critical level and then falls, because

crowded streets slow every vehicle, including those trying to leave. Above the critical density the network is hypercongested: more vehicles produce fewer completed trips. In this setting a toll that keeps density below the critical level reduces the frequency and severity of jams, and commuters are better off even if they receive no benefit from the toll revenue (Arnott, 2015). Fosgerau and Small (2013) reached related conclusions for dense downtown networks. The bathtub is the economic counterpart of Greenshields' parabolic flow–density curve. We borrow it as a hypothesis for Nairobi's central business district (CBD), where evening gridlock around blocked junctions has the appearance of hypercongestion, and we argue that its lesson does not depend on charging: any instrument that holds inflow below the critical density, such as holding vehicles at termini, keeping junction boxes clear or timing signals to meter entry, could capture part of the gain.

Ostrom: Governing a Commons Without Prices

Road space at a busy stage or junction has the features of a common-pool resource: it is hard to exclude users, and each additional user reduces what is left for others. From field studies of fisheries, forests and irrigation systems, Ostrom (1990) showed that such resources are often governed successfully neither by the state nor by private owners but by institutions their users devise, with clear boundaries, accountable monitoring, graduated sanctions and low-cost ways of resolving conflict. Whether Nairobi's stages are already governed by informal rules of this kind (who loads where and when, and who enforces it) is an empirical question that the proposed diagnosis is designed to answer (Appendix A). We borrow Ostrom's design principles for the coordination forum proposed in this paper, and we argue that for Nairobi's first phase this route is more promising than either Pigou's charge or Knight's owner.

Evidence From Cities in Developing Countries

Two recent studies caution against importing conclusions from high-income cities. Using travel times for 154 Indian cities, Akbar et al. (2023) found that differences in speed across cities are driven more by uncongested speed, meaning how fast traffic moves when roads are nearly empty, than by congestion, and that raising uncongested speed would increase mobility more than optimal congestion pricing. In a large field experiment in Bangalore that introduced experimental congestion charges for commuters, Kreindler (2024) found a moderate, roughly linear relationship between traffic density and delay, implying that optimal charges would reduce travel times only slightly and yield small welfare gains. Nairobi's data point in the same direction: the free-flow speed implied by TomTom's 2025 measurements is only about 33.5 km/h, against 16.7 km/h in the evening peak (TomTom, 2026). Part of the time Nairobi commuters lose is therefore not congestion delay at all, and it lies outside the cost model in this paper.

Pricing in Practice, and Nairobi's Priced Road

Vickrey's proposals have since been implemented in several cities, including Singapore, London and Stockholm (see International Interventions and Their Transferability below). New York City began charging vehicles entering Manhattan south of 60th Street on 5 January 2025. In the first year, officials reported 27 million fewer vehicle entries into the charging zone, an 11% reduction (Office of the Governor of New York, 2026), with the revenue helping to fund the Metropolitan Transportation Authority's capital projects (Lassen, 2026). Nairobi already has a priced road. The Nairobi Expressway's operator charges by vehicle class and distance travelled but not by time of day; in 2023 a saloon car paid KSh 360 from Mlolongo to Westlands ("Nairobi Expressway Explains ETC," 2023). In Vickrey's terms the expressway sells a faster alternative to those willing to pay, as Knight's owner would; it does not price the peak on the roads beside it, where most passengers travel. It suggests that some road users will pay for time savings, but it is not evidence about congestion pricing.

Network Behaviour and Induced Demand

Wardrop (1952) formalised the user equilibrium in which no traveller can reduce their own travel time by changing route. Because each user ignores their effect on others, the equilibrium can be markedly worse than the system optimum. Roughgarden and Tardos (2002) showed that with linear delay functions selfish routing produces total travel time at most four-thirds of the optimum, and Braess (1968) showed that adding a link to a network can make every user slower. Empirically, Duranton and Turner (2011) found that vehicle-kilometres travelled in US cities rose almost one-for-one with new road capacity. Together these results caution against expecting infrastructure alone to resolve congestion and suggest that information and coordination can recover some of the gap between equilibrium and optimum without new construction.

Valuing Travel Time

Valuing congestion requires a value of travel time. Appraisal practice commonly sets the value of commuting time as a fraction of the wage rate and weights out-of-vehicle time (walking and waiting) more heavily than in-vehicle time. Wardman (2004) reviewed public transport values of time in the United Kingdom, and the convention that walking and waiting time are worth about twice in-vehicle time remains widespread, although later meta-analysis suggests somewhat lower multipliers. We found no published estimate of a value of time for Nairobi commuters across modes and income groups, a gap that matters because, as the Results show, this single parameter dominates the uncertainty in any cost estimate.

Table 2 summarises the position this paper takes on each contribution.

Table 2

How Economists Have Analysed the Traffic Jam: What This Paper Borrows, Finds Consistent, Qualifies or Contests

Economist	Central claim	Nairobi evidence or test	Our position
Pigou (1920)	Drivers ignore the delay they impose on others; a charge on the congested road raises welfare	Marginal external delay of 7.2 vehicle-minutes per trip at capacity (Results)	**Consistent**; borrowed as Equation 8
Knight (1924)	An owner charging a toll would ration road use efficiently	The expressway has an owner and a toll but no peak price; streets and stages have no owner	**Qualify:** what Nairobi lacks is an owner for each recurring problem, not only a price
Walters (1961)	Measure marginal social cost from speed–flow data; efficient prices exceed current charges	Illustrative Pigouvian toll of KSh 27 per car trip at capacity	**Borrow** the method; local speed–flow data needed (P2)
Vickrey (1963, 1969)	Price road use by time of day; in the bottleneck equilibrium every departure time costs the same	Illustrative approach: 48-minute peak queue; every arrival time costs KSh 118.6	**Consistent** within the model: “effort cannot beat a jam”; **argue** that for matatu passengers the timing decision is largely the SACCO’s
Downs (2004)	Triple convergence absorbs improvements; peak congestion rations road space	JICA: speeds keep falling with all planned projects; capacity +25% with demand +20% raises queuing loss	**Consistent** (convergence); **qualify** rationing, since much Nairobi delay is lost capacity
Small (1982); Arnott et al. (1993)	Lateness costs more than earliness; the optimal toll halves social cost but leaves private cost unchanged before revenue use	2018 CBD matatu ban suspended within 24 hours; Stockholm’s charge followed a full-scale trial	**Borrow** the model; **argue** that revenue use and visible benefits decide acceptance
Arnott (2013, 2015)	Above a critical density, more vehicles complete fewer trips; tolling helps even if revenue is wasted	Evening CBD gridlock around blocked junctions (to be tested)	**Borrow** as hypothesis H2; **argue** that non-price metering can capture part of the gain
Ostrom (1990)	Shared resources can be governed by rules their users devise, with monitoring and graduated sanctions	Stage and junction space is shared; any informal rules are undocumented	**Borrow** design principles for the coordination forum
Akbar et al. (2023)	In Indian cities, uncongested speed matters more than congestion	Implied free-flow speed 33.5 km/h against 16.7 km/h in the evening peak	**Qualify** our cost model, which omits slow free-flow travel

Economist	Central claim	Nairobi evidence or test	Our position
Kreindler (2024)	Optimal congestion charges yield small welfare gains in Bangalore	No Nairobi experiment exists	**Argue** against leading with pricing; test cheaper coordination first

Note. Consistent = the claim agrees with the available Nairobi evidence or holds within the model, but has not been tested empirically in Nairobi; borrow = the method or model is used in this paper; qualify = the claim holds only in part in Nairobi; argue = we take a different position for Nairobi. CBD = central business district; P2 refers to the open problems listed in the Discussion.

Taken together, these studies indicate that congestion is shaped by demand, capacity, traveller behaviour, scheduling preferences and pricing. They agree that unpriced road space is over-used at peak and that added capacity tends to be absorbed by new traffic. They differ on the remedy (a public charge, an owner or user institutions) and, in the developing-country evidence, on the size of the gains from pricing. Most of the models focus on vehicle traffic and on travellers who choose their own departure times and routes, so their ability to explain passenger waiting and the operational constraints of Nairobi's matatu system remains uncertain. This study therefore treats these mechanisms as competing or interacting explanations to be examined through field diagnosis, and uses the models only to illustrate what each would imply under stated assumptions.

What Is Known About Congestion in Nairobi

The JICA household interview survey of 2023 is the most recent comprehensive source on travel in the metropolitan area. It records 8.85 million motorised person-trips a day, a walking share of 37.2% and a matatu share of 32.2% of all trips, with average matatu occupancy of 11.2 passengers (JICA Expert Team, 2025). Salon and Gulyani (2019) showed that commuting patterns in Kenyan cities are shaped by income, with walking dominant among lower-income households and long matatu journeys common.

Safety and inclusion are inseparable from congestion. Women report harassment and safety concerns on matatus (Flone Initiative, 2018; UN Women et al., 2021). Pedestrians were the largest group of road deaths in Kenya in 2025, with 1,580 recorded deaths ("Road Accident-Related Deaths Rise," 2025), and the World Health Organization estimated 14,926 road deaths in Kenya in 2021, far above police-recorded totals (Bett, 2024; World Health Organization [WHO], 2023). Motorcycle taxis have grown rapidly: 241,763 motorcycles were registered in 2025, 61% of all new registrations ("Bikes Boom Drives Spike," 2026). Industry estimates of boda boda incomes exist ("Boda Boda Riders With Financed Motorcycles," 2025; "How Boda Boda Riders Make Sh1bn," 2022), but they have not been independently verified and are cited for order of magnitude only.

Nairobi has also tried several interventions whose records shed light on mechanisms. Table 3 summarises four of them, following the structure recommended by the independent review (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026): what the record supports and what a new pilot should ask. In each case the record documents implementation and disruption more clearly than measured effects on passenger journeys.

Table 3

Previous Nairobi Interventions: What the Record Supports and What a Pilot Should Ask

Intervention	What the record supports	Question for a pilot
Junction works: Globe Roundabout closure, March 2026	KeNHA closed the roundabout for 13 days to repair the Nairobi River bridge; traffic was diverted via the Globe overpass and movement into the CBD was disrupted (Etyang, 2026; Njaaga, 2026). Nearly all Thika Road public service vehicles use the roundabout to reach the CBD.	Which movement is constrained? Where does the queue go when a link closes or opens? Request turning counts and junction plans.
Corridor expansion: Thika Road	The African Development Bank reported that travel between Nairobi and Thika fell from 2–3 hours to 30–45 minutes after the upgrade (African Development Bank, 2017). The review also notes reported traffic growth and land development beyond expectations; we have not yet verified those figures.	Do corridor gains carry through to CBD approaches and to passengers' waiting and total trip times?
Market relocation: Wakulima to Kangundo Road, 2024	The County moved selected traders from a market built for 1,200 traders but used by more than 3,600; protests disrupted traffic and a court issued a temporary order (Muriuki, 2024; "Reprieve for Marikiti Traders," 2024). No estimate of the congestion effect is available.	Measure loading, pedestrian access and traffic at both sites; ask traders and customers how their trips changed.
Moving stages outside the CBD: 2018 ban and 2021 Green Park trial	A December 2018 ban on matatus in the CBD forced thousands to walk and was suspended within 24 hours (Kinyanjui, 2018; "Sonko Bows to Pressure," 2018). A 2021 three-hour test at Green Park terminus left commuters walking to the CBD; free buses were offered to people with disabilities and parents with small children (Kinyanjui, 2021; "Test Run by NMS," 2021).	Track the full journey when a stage moves: walking access, waiting, transfers, fare and vehicle travel time.

Note. Adapted from the structure of the independent advisory review (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026), with sources verified by the author. CBD = central business district.

Two lessons recur. Changes in access, such as closing a junction, relocating traders or moving a stage, shift costs onto passengers' walking, waiting and transfers, which vehicle-based measures miss. And none of these interventions published a before-and-after measurement of passenger journeys, which makes it impossible to learn from them systematically.

What is known about Nairobi therefore rests mainly on one comprehensive survey, commercial probe data that measure car travel, official statistics often reported through the press, and the records of individual interventions. These sources establish the scale of travel and the decline in speeds, but they do not measure passenger waiting at stages, the reliability of matatu journeys or the distribution of delay across modes and income groups.

Matatu Fleet Dynamics, Dispatch, Headways and Passenger Queues

Paratransit is the main motorised mode in many cities of the Global South. Cervero and Golub (2007) reviewed informal transport worldwide and found that it provides on-demand mobility, jobs and coverage where formal transit is absent, but also adds to congestion, pollution and crashes; policy responses range from suppression to regulation and integration. In Nairobi, matatus are run as private businesses whose organisation ranges from individually owned vehicles to larger, more organised fleets, and business strategy and the regulatory regime shape how vehicles are operated (McCormick et al., 2013). Williams et al. (2015) mapped Nairobi's matatu routes and published them in the General Transit Feed Specification, demonstrating that semi-formal transit can be described with open data standards.

The operational mechanism at the centre of this paper, that the regularity of vehicle arrivals governs passenger waiting, is well established for scheduled buses. Newell and Potts (1964) showed that headways on a bus line are unstable: a bus that falls behind meets more waiting passengers, takes longer to load and falls further behind, while the bus behind it catches up, so that buses bunch. Later work designed holding and dispatch strategies to restore regular headways (Daganzo, 2009). These results

were derived for scheduled services with fixed stops. In Nairobi, dispatch is decided by SACCOs and crews rather than by a published timetable, so the relevant mechanism may combine headway instability on the road with dispatch and loading practices at the stage.

We found no published study that measures how delays on vehicles' inbound journeys affect the supply of seats and passenger waiting at Nairobi's central stages. The bunching literature supplies the mechanism and the paratransit literature the operating context, but their combination in a fill-and-dispatch setting has not been modelled or measured, which motivates the stage seat-supply model in this paper.

Governance, Coordination, Enforcement and Ownership

Responsibility for Nairobi's roads is divided. The Kenya National Highways Authority (KeNHA) manages national highways that cross the city, including Thika Road; the Kenya Urban Roads Authority (KURA) manages urban roads and is implementing the intelligent transport system; the Nairobi Metropolitan Area Transport Authority (NaMATA) plans mass transit; the National Transport and Safety Authority (NTSA) regulates public service vehicles and SACCOs; the National Police Service enforces traffic law; and Nairobi City County manages county roads, parking, termini and boda boda regulation (KURA, 2026; Nairobi City County, 2026; Nairobi Metropolitan Area Transport Authority [NaMATA], n.d.). A statutory forum already exists in the County Transport and Safety Committee, which brings together the County, NTSA, police and operators (Nairobi City County, 2024). New data partnerships, such as the 2026 agreement between Safaricom and the Kenya National Bureau of Statistics on anonymised mobile data, expand what could be measured (Wolfenden, 2026). The policy review of the Africa Transport Policy Program (SSATP, 2018) identified institutional fragmentation, the absence of dedicated funding and the need to reform paratransit as priorities for Kenyan cities.

The governance literature thus documents fragmented responsibility and an existing statutory forum, and the commons literature reviewed above suggests that shared resources can be governed by user institutions with monitoring and graduated sanctions (Ostrom, 1990). What the public record does not document for Nairobi is a mechanism that assigns a recurring, junction-level or stage-level problem to a named owner and checks whether passenger journeys improve. That absence is inferred from documents and must be tested through interviews.

Monitoring Systems and Artificial Intelligence

Nairobi is building the components of a monitoring system. KURA's intelligent transport system programme is establishing a Traffic Management Centre at City Cabanas, where feeds from cameras at signalised junctions are to be monitored ("KURA Initiates Plan," 2025). The first phase, covering 25 junctions and the centre, was under construction in mid-2026 ("Ruto Unveils Ksh7.9B Smart Traffic System," 2026), and a third phase plans to use AI to count vehicles and passengers, track turning movements and violations, and send live data from 125 junctions to the centre ("Government to Roll Out AI-Powered Traffic Lights," 2026). NTSA is installing speed cameras with instant digital fines ("NTSA to Install 1,000 Speed Cameras," 2026), the County Transport and Safety Committee brings agencies and operators together (Nairobi City County, 2024), and the Safaricom mobile data partnership described above could show how people move across the city (Wolfenden, 2026). Outside government, crowd-sourced reports have already been turned into data with AI: Milusheva et al. (2021) applied a machine-learning model to 874,588 traffic-related tweets in Nairobi and geolocated 32,991 crash reports from 2012 to 2020, producing one of the first crash maps of the city.

International evidence on AI in traffic management is promising but mostly reported by the developers of the systems concerned. An adaptive signal control system piloted at nine intersections in Pittsburgh reduced average travel times through the pilot area by more than 25% in its developers' before-and-after evaluation (Smith et al., 2012), and Google reports fewer stops at junctions where cities adopted its AI-based signal timing recommendations (Google, 2023). Neither kind of result shows that AI

improves coordination among institutions, which is the problem this paper identifies. Kenya adopted a National AI Strategy for 2025–2030 in March 2025 (Ministry of Information, Communications and the Digital Economy [MICDE], 2025), but it has no dedicated AI law and relies on data protection legislation and sectoral policies (Data Protection Act, 2019; EY, 2025). The AI risk management framework of the NIST (2023) calls for clear accountability structures, with roles and responsibilities assigned to identified people and organisations, and for AI systems to be measured against their intended purpose before and during use. The literature therefore describes monitoring components and AI tools, but not a system that links what is detected to an institution that must act and to a public record of whether it did.

Much of the evidence on AI performance comes from developers, from high-income cities or from traffic conditions unlike Nairobi's, so whether these tools perform comparably in Nairobi is untested.

International Interventions and Their Transferability

International experience is useful for Nairobi only if the conditions under which an intervention worked are made explicit. Table 4 therefore records, for each case, the context that differs from Nairobi's, the outcome reported in the cited source and what may transfer. Singapore, for example, provides evidence that congestion charges can be operated through periodic adjustment against target speeds; its institutional, technological and regulatory context differs substantially from Nairobi's, however, so its experience is treated as a policy design reference rather than as evidence of expected effects in Nairobi.

Table 4

International Interventions: Context, Reported Outcome and Conditions for Transfer to Nairobi

Case	Context that differs from Nairobi	Outcome reported in cited source	What may transfer, and on what condition
Singapore: electronic road pricing (Land Transport Authority, 2026a; Land Transport Authority, 2026b)	High income; one land transport authority with charging powers; mature electronic infrastructure	Charges adjusted each quarter against target speed bands	The measure-and-adjust cycle, as a design reference; requires an institution able to act on the measurements
London and Stockholm: cordon charges ("Clearing Congestion," 2016; Eliasson et al., 2009; Resources for the Future, n.d.)	High income; strong metropolitan authorities; rail and bus alternatives	Private-vehicle traffic down by roughly a fifth; London's speeds later eroded	Trial before permanence and visible use of revenue; gains can erode without continued management
New York: congestion pricing (Lassen, 2026; Office of the Governor of New York, 2026)	High income; dense rail network; revenue earmarked for transit capital	27 million fewer vehicle entries in the first year (11%)	Earmarking revenue for public transport, once alternatives exist
Seoul: 2004 bus reform (Asian Development Bank, n.d.)	Formal bus companies; strong city government	Reform package of median lanes, route restructuring and operator contracts (components, not effects, are used here)	Contracting operators with performance terms, if regulators can monitor performance
Bogotá: TransMilenio (BRTData, n.d.)	Dedicated BRT infrastructure and fleet	Peak loads near 49,000 passengers per hour per direction	High-capacity protected lanes, if fleet and funding are sustained
Jakarta: angkot integration (Institute for Transportation and Development Policy [ITDP], 2021)	Large city with paratransit, integrated through a public network and fare system	Minibus operators brought into an integrated network and fare system	Integrating rather than displacing paratransit, through contracts
Dar es Salaam: BRT (ITDP Africa, 2022; National Audit Office of Tanzania, n.d.)	Closest regional comparator; similar income and paratransit dependence	Shorter on-road times, but 65 of 305 planned buses in service and waits up to 121 minutes	Service supply is part of the intervention; infrastructure alone does not shorten passengers' journeys

Note. Outcomes are as reported in the cited sources and are mostly vehicle-based; few report passenger waiting time or journey reliability. BRT = bus rapid transit.

Dar es Salaam is the closest regional comparison and a cautionary one. Its BRT reduced on-road travel times, but Tanzania's National Audit Office found that only 65 of the planned 305 buses were in service and that passengers waited up to 121 minutes (ITDP Africa, 2022; National Audit Office of Tanzania, n.d.). The lesson for Nairobi is that service supply is part of the intervention: protected lanes do not shorten passengers' journeys if vehicles are not there to carry them. Finally, technology-led signal optimisation, such as Google's Project Green Light, reports fewer stops in pilot cities, but its results are reported by the vendor (Google, 2023).

The cases suggest that interventions that reduced congestion combined a measurement cycle, an institution with authority to act, and attention to service supply and public acceptance, and that where operators were integrated rather than displaced, as in Seoul and Jakarta, reform proceeded through contracts. Few of the cited sources report passenger waiting or journey reliability,

so even the successful cases offer limited evidence on the outcomes this paper prioritises.

Research Gap

Existing research on urban congestion provides established approaches for modelling traffic flow, valuing travel-time losses, analysing bottlenecks, assessing congestion pricing and evaluating transport infrastructure. These approaches do not necessarily capture the operational characteristics of passenger transport systems in which vehicle dispatch, loading practices, fleet circulation and institutional responsibilities influence passenger waiting time and journey reliability.

In Nairobi, available transport-demand, traffic-speed and institutional information provides a foundation for examining congestion, but important empirical gaps remain. They include the relationship between inbound vehicle delays and passenger queues at central stages; the relative contribution of operational and infrastructure constraints to passenger delay; the distribution of congestion costs across modes and income groups; the transparency of published cost estimates, which span almost two orders of magnitude and rarely report their methods; and the effectiveness of institutional coordination, and of AI-assisted monitoring, in responding to recurring traffic problems.

The present study addresses these gaps by developing a passenger-centred diagnostic framework, estimating the economic burden of weekday peak congestion from secondary data with a published uncertainty analysis, formalising an illustrative stage seat-supply model and proposing an evaluation protocol for a future operational intervention. The empirical contribution of the proposed diagnostic and coordination framework remains to be established through primary fieldwork and comparative evaluation.

Method

Research Design

The study combines four components: a structured synthesis of secondary evidence, a set of parametric models, a protocol for field diagnosis and evaluation, and an assessment of the systems institutions use to monitor and coordinate traffic. Table 5 maps each research question to its data, method, output and current status. Figure 1 shows the conceptual model. It separates institutional and operational factors, the traffic and public transport mechanisms through which they act, and the passenger outcomes they produce; it places intervention and evaluation beneath all three, with a feedback loop from passenger outcomes to institutional decisions, and it recognises external influences that can mimic or mask the effect of any intervention.

Table 5

Methods Map: Research Questions, Data, Methods and Outputs

RQ	Data	Method	Output in this paper	Status
RQ1	JICA survey; TomTom; KNBS wages; EPRA fuel prices; published estimates	Cost model; one-at-a-time sensitivity; Monte Carlo; top-down cross-check	Preliminary estimate with uncertainty interval	Estimated; needs mode-specific data
RQ2	Model parameters (illustrative); literature	Stage seat-supply model; engineering and economists' models; diagnostic matrix	Mechanisms and hypotheses H1 and H2	Illustrated; needs field calibration
RQ3	Intervention records; international sources	Documentary synthesis; transferability assessment	Tables of interventions and transferability	Inferred
RQ4	To be collected: passenger, vehicle and junction observations; interviews	Diagnosis protocol; pre-specified difference-in-differences; business case	Protocol, instruments, sample sizes	Proposed
RQ5	Official documents and press reports; to be collected: institutional interviews	Five-function systems assessment; technology readiness review	Systems table; staged AI-assisted system design	Inferred and proposed

Note. EPRA = Energy and Petroleum Regulatory Authority; JICA = Japan International Cooperation Agency; KNBS = Kenya National Bureau of Statistics.

Figure 1

Conceptual Model: Institutional Factors, Mechanisms, Passenger Outcomes and Evaluation

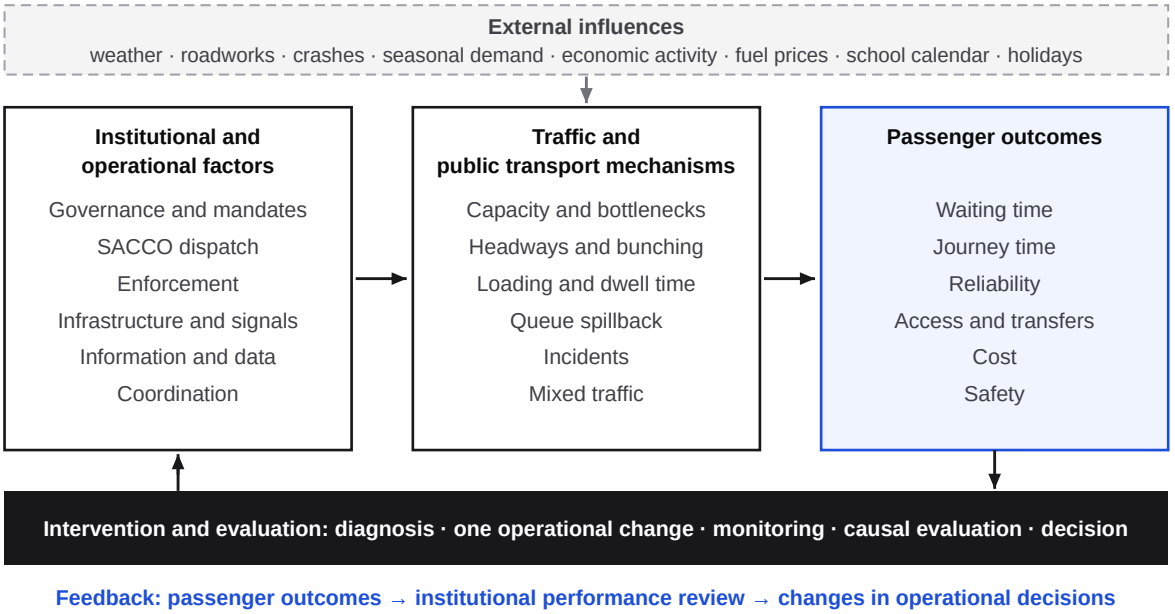

Note. Arrows indicate the hypothesised direction of influence; the model is a framework for measurement, not a tested causal structure. External influences are logged during fieldwork so that they are not mistaken for intervention effects.

Data Sources and Verification

Table 6 lists the principal data sources. Sources were accessed between 18 and 23 November 2026 and graded in four tiers: Tier 1, primary scientific and official evidence (original datasets, peer-reviewed studies, official reports and legislation); Tier 2, credible institutional evidence (development-partner and university reports); Tier 3, secondary reporting (news reports of official statistics or announcements); and Tier 4, unverified claims and research leads, which are not used as evidence. Where only a press report of an official statistic was available, the press report is cited and the claim is marked accordingly. A source-verification registry (Supplementary File S1) records, for every reference, its tier, type, link, the claims it supports, verification status, limitations and the date checked, and a data dictionary (Supplementary File S2) documents every model input. Currency conversions use Central Bank of Kenya rates of 5 November 2026: US$1 = KSh 129.96 and €1 = KSh 146.27 (“Kenyan Shilling Nears Ksh130,” 2026).

Table 6

Principal Data Sources

Source	Content used	Reference period
JICA household interview survey and demand forecast	Person-trips, modal shares, occupancy, speeds, volume-to-capacity ratios, forecasts	2013, 2023; forecasts to 2045
TomTom Traffic Index	Average and peak travel time per 10 km; congestion level	2025
Kenya National Bureau of Statistics (via press reports)	Average wage earnings; vehicle registrations	2024–2025
Energy and Petroleum Regulatory Authority (via press reports)	Pump prices; tax share of pump price	April and November 2026
NTSA and WHO	Road deaths by road-user type; modelled death estimates	2021, 2023, 2025
KURA, National Treasury, development partners	Intelligent transport system phases, loans, BRT financing	2023–2026
News reports and audit documents	Previous interventions in Nairobi; Dar es Salaam BRT performance	2018–2026

Note. Full citations appear in the reference list. JICA = Japan International Cooperation Agency; WHO = World Health Organization.

Cost-of-Congestion Model

The model estimates the annual cost of weekday peak congestion as the value of extra time spent by people travelling in matatus, buses and cars plus the cost of fuel burned during that extra time:

$$\text{Cost}_{\text{year}} = \left[\, P\, s\, d\, \text{VOT} \;+\; V\, s\, d\, \bar{f}\, p \,\right] D \tag{1}$$

where P is daily person-trips by matatu, bus and car; s is the share of those trips made in the peaks; d is the extra time per peak trip relative to free flow, in hours; VOT is the value of travel time per hour; $V\bar{f}$ is the fuel burned per hour of delay by the corresponding vehicle trips; p is the pump price per litre; and D is the number of weekdays per year. The terms are built from observed inputs as follows:

$$P = T_{\text{all}}\,(m_{\text{matatu}} + m_{\text{bus}} + m_{\text{car}}), \qquad T_{\text{all}} = \frac{T_{\text{motorised}}}{1 - w} \tag{2}$$

$$d = \frac{\delta_{10}}{60} \cdot \frac{L}{10}, \qquad \delta_{10} = t_{\text{peak},10} - \frac{t_{\text{avg},10}}{1 + c} \tag{3}$$

$$\text{VOT} = \theta\, \frac{W}{H} \tag{4}$$

$$V\bar{f} = \frac{T_{\text{all}}\, m_{\text{car}}}{o_{\text{car}}}\, f_{\text{car}} + \left(\frac{T_{\text{all}}\, m_{\text{matatu}}}{o_{\text{matatu}}} + \frac{T_{\text{all}}\, m_{\text{bus}}}{o_{\text{bus}}} \right) f_{\text{psv}} \tag{5}$$

Here $T_{\text{motorised}}$ is daily motorised person-trips and w the walking share, so T_{all} recovers total trips; m denotes modal shares of all trips; $t_{\text{avg},10}$ and $t_{\text{peak},10}$ are TomTom's average and peak travel times per 10 km and c its average congestion level, so $t_{\text{avg},10}/(1 + c)$ is the implied free-flow time; L is average trip length; θ is the value of time as a share of the hourly wage; W is average annual wage earnings and H annual paid hours; o is vehicle occupancy; and f is fuel burned per hour of delay. Boda boda passengers and pedestrians are excluded from P, so the estimate does not cover their delay. Two approximations should be noted. First, the model applies a common peak-delay estimate derived from road-speed data, which mainly reflect cars, to matatu, bus and car trips; this represents exposure to traffic delay rather than a direct measurement of mode-specific passenger travel time, and it must be replaced by mode-specific observations where available. Second, the value of time is wage-based, so the result is a wage-based economic valuation rather than a comprehensive social valuation of delay for informal workers, students, care workers and others without an hourly wage. Fuel is valued in two ways: at the pump price, which is the private

cost to operators and motorists, and at resource cost, excluding the taxes and levies that make up about 36–45% of the pump price (“Explained: Taxes and Subsidy,” 2026), which is the appropriate measure for economic appraisal. Table 7 lists every input with its low, central and high values.

Table 7

Model Inputs, Values and Sources

Input	Low	Central	High	Type and source
Motorised person-trips per day, $T_{\text{motorised}}$		8,852,297		Observed; JICA (2025)
Walking share, w		0.372		Observed; JICA (2025)
Modal shares: matatu, bus, car		0.322, 0.090, 0.132		Observed; JICA (2025)
Occupancy: matatu, bus, car		11.2, 40, 1.5		Matatu observed (JICA); others assumed
Peak delay per 10 km, δ_{10} (min)	11.4	14.7	18	Derived from TomTom (2026)
Share of trips in peaks, s	0.40	0.50	0.60	Assumed
Trip length, L (km)	8	10	12	Assumed
Value of time share, θ	0.20	0.30	0.50	Assumed (appraisal convention)
Average wage, W (KSh per year)		988,200		Observed; KNBS via press
Paid hours, H		2,000		Assumed
Fuel burn, car / matatu or bus (L/h)	0.6 / 1.2	0.8 / 1.5	1.0 / 2.0	Assumed
Pump price, p (KSh per litre)		227		Observed; EPRA via press
Weekdays, D		250		Assumed

Note. The low and high scenarios move all uncertain inputs together, so the resulting range is a plausible envelope rather than a statistical confidence interval. Blank cells indicate inputs held fixed across scenarios. EPRA = Energy and Petroleum Regulatory Authority; KNBS = Kenya National Bureau of Statistics.

Uncertainty Analysis and Cross-Checks

Uncertainty is examined in three ways. One-at-a-time sensitivity analysis swings each uncertain input between its low and high value while holding the others at their central values; it is retained as a diagnostic of which inputs matter. The low and high scenarios move all uncertain inputs together and therefore define an envelope, not a probability range. The primary uncertainty analysis is a Monte Carlo simulation of 10,000 draws in which each uncertain input is assigned a distribution (Table 8). The distributions express the same bounds as the scenarios, with the central value as the mode where one is justified; they are structured assumptions to be replaced by empirical distributions as field data become available. The two fuel-burn rates are drawn with a rank correlation of 0.5 through a Gaussian copula, because both depend on the same stop-and-go conditions. We report the median, mean, 5th and 95th percentiles, the probability that the cost exceeds policy-relevant thresholds and each input’s approximate share of output variance, computed from squared rank correlations normalised to sum to one.

Table 8

Distributions Used in the Monte Carlo Analysis

Input	Distribution	Basis	Evidence needed to replace it
Value-of-time share of wage, θ	Triangular (0.20, 0.30, 0.50)	Appraisal conventions for commuting time	Nairobi stated-preference survey by mode and income
Peak delay per 10 km, δ_{10}	Uniform (11.4, 18) minutes	A trip may fall in the morning or evening peak, whose delays bound the range	Mode-specific probe or GPS data by corridor and hour
Share of trips in peaks, s	Triangular (0.40, 0.50, 0.60)	Assumed bounds	Departure-time distribution from the JICA survey
Trip length, L	Triangular (8, 10, 12) km	Assumed bounds	Trip-length distribution from the household survey
Fuel burn, car and matatu or bus	Triangular (0.6, 0.8, 1.0) and (1.2, 1.5, 2.0) L/h; rank correlation 0.5	Assumed idle and creep consumption	Fleet or on-board fuel data
Trips, modal shares, occupancy, wage, fuel price, weekdays	Fixed	Observed (JICA, KNBS, EPRA) or conventional (250 weekdays)	Updated surveys and statistics

Note. Triangular parameters are (minimum, mode, maximum). The random seed is fixed so that results can be reproduced exactly (Appendix B).

As a top-down cross-check, we also apply a common simplified formulation in which the cost is the number of commuters materially delayed multiplied by the hours each loses, the value of time and the number of weekdays:

$$C = N\,\Delta t\,\mathrm{VOT}\,D \tag{6}$$

where N and Δt take round-number assumed values and VOT is the central value from Equation 4. We compare the physical quantity both formulations imply (person-hours lost per weekday) separately from the monetary result, and we compare both with previously published estimates converted to annual nominal values. A nominal 20% uplift for weekend and off-peak delay shows the direction of one known omission.

Congestion Engineering Relationships

Five standard relationships are used to interpret the evidence. The first is the volume–delay function of the Bureau of Public Roads (1964):

$$t = t_0\left[1 + \alpha\left(\frac{V}{C}\right)^{\beta}\right] \tag{7}$$

where t_0 is free-flow travel time and V/C the volume-to-capacity ratio. Differentiating total delay with respect to one additional vehicle gives the marginal external delay imposed on others, and multiplying by the value of time of those delayed gives the Pigouvian toll (Pigou, 1920; Vickrey, 1969):

$$V\frac{\partial t}{\partial V} = t_0\,\alpha\,\beta\left(\frac{V}{C}\right)^{\beta}, \qquad \tau^{*} = \mathrm{VOT}\cdot o\cdot V\frac{\partial t}{\partial V} \tag{8}$$

For a deterministic bottleneck with capacity μ, arrivals $\lambda > \mu$ for a period T and arrivals $\lambda_2 < \mu$ afterwards, total vehicle delay is

$$D_q = \tfrac{1}{2}\,(\lambda-\mu)\,T\left(T + \frac{(\lambda-\mu)\,T}{\mu-\lambda_2}\right) \;\propto\; T^2 \tag{9}$$

so halving the overload period reduces delay by three-quarters. Webster's approximately optimal fixed cycle length (Webster, 1958) is

$$C_0 = \frac{1.5\,L_{\mathrm{lost}} + 5}{1 - Y} \tag{10}$$

where L_{lost} is lost time per cycle in seconds and Y is the sum of critical flow ratios. Finally, two indicators are defined for publishing corridor conditions: a congestion index for segment s, and a hotspot score that weights it by the estimated number of people flowing, $\hat{P}_s(t)$:

$$\mathrm{CI}_s(t) = \min\!\left(1, \max\!\left(0,\; 1 - \frac{v_s(t)}{v_{f,s}}\right)\right), \qquad H_s = \sum_{t} \mathrm{CI}_s(t)\, \hat{P}_s(t)\, \Delta t \tag{11}$$

where $v_{f,s}$ is the 85th percentile of night-time speeds on the segment. The hotspot score is measured in person-hours, so a segment carrying full matatus ranks above one carrying the same number of private cars.

Economists' Models: The Bottleneck and the Bathtub

Two models from the economics of congestion are used to examine, under stated assumptions, claims reviewed in the Literature Review. The first is Vickrey's bottleneck in the form developed by Arnott et al. (1993). N identical commuters pass a bottleneck of capacity s vehicles per hour and wish to arrive at time t^{*}. Each values travel time at α per hour, arriving early at β per hour and arriving late at γ per hour, with $\beta < \alpha < \gamma$; these symbols follow the bottleneck literature and are unrelated to the parameters of Equation 7. In equilibrium no commuter can gain by changing departure time, so for every arrival time a in the rush period the sum of queuing cost, $\alpha T(a)$, and schedule-delay cost is the same:

$$\alpha\, T(a) + \beta\, \max(0,\, t^{*} - a) + \gamma\, \max(0,\, a - t^{*}) = \delta\, \frac{N}{s}, \qquad \delta = \frac{\beta\gamma}{\beta + \gamma} \tag{12}$$

The rush period lasts N/s hours. The longest queue, faced by the commuter arriving at t^{*}, and the total cost per morning without a charge are

$$T_{\max} = \frac{\delta N}{\alpha s}, \qquad TC_0 = \delta\, \frac{N^2}{s} \tag{13}$$

of which half is queuing and half is schedule delay. A charge that varies through the peak exactly as queuing cost would have done removes the queue without changing when people arrive:

$$\tau(a) = \alpha\, T(a), \qquad \tau_{\max} = \delta\, \frac{N}{s}, \qquad TC_1 = R = \delta\, \frac{N^2}{2s} \tag{14}$$

so total social cost TC_1 is half of TC_0, an equal amount R is collected as revenue, and each commuter's private cost is unchanged at $\delta N/s$ before the revenue is used (Small, 2015). Following Small (2015), we set $\beta = 0.5\alpha$ and $\gamma = 2\alpha$ as approximations to Small (1982), and α equal to the central value of time in Equation 4.

The second is the Vickrey bathtub as formalised by Arnott (2013). With Greenshields' linear relation between speed and density (Greenshields, 1935), the rate at which trips are completed in a network of density k is proportional to

$$g(k) = k\, v_f \left(1 - \frac{k}{k_j}\right), \qquad k_c = \frac{k_j}{2}, \qquad g_{\max} = \frac{v_f\, k_j}{4} \tag{15}$$

where v_f is free-flow speed and k_j jam density. Above the critical density k_c, additional vehicles reduce the number of trips completed per hour. Both models are applied with illustrative parameters; their purpose is to show what each argument implies under its assumptions, not to forecast or to validate the argument empirically.

The Central-Stage Seat Supply and Passenger Queue Model

Purpose. An independent review of an earlier draft (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026) drew attention to the difference between where a passenger queue is seen and where it originates. The model formalises one way that difference can arise: a delay on vehicles' inbound return trips, perhaps several kilometres from the stage, reduces the rate at which seats reach a central stage and so lengthens the queue there, even when the stage itself operates well.

Structure. With N vehicles of k seats on a route whose round-trip cycle is C_r minutes, seats arrive at the stage at rate

$$S = \frac{60\,N\,k}{C_r}, \qquad C_r = t_{\text{out}} + t_{\text{in}} + \Delta_{\text{in}} + t_{\text{term}} \tag{16}$$

where t_{out} is the loaded outbound trip, t_{in} the free-flow inbound return, Δ_{in} any extra delay on the inbound return and t_{term} terminal time at both ends. If passengers arrive at rate λ per hour and $\lambda > S$, the passenger queue grows linearly over a peak of T hours, giving a maximum queue, a maximum extra wait and a mean extra wait (in minutes) of

$$Q(T) = (\lambda - S)\,T, \qquad W_{\max} = \frac{60\,Q(T)}{S}, \qquad \bar{W} = \frac{30\,(\lambda - S)\,T}{S} \tag{17}$$

Table 9 defines the variables and the data needed to measure each.

Table 9

Variables of the Stage Seat-Supply Model and the Data Required to Measure Them

Symbol	Meaning	Measurement in the field
N	Vehicles operating on the route in the peak	SACCO fleet records; stage counts of distinct vehicles
k	Seats per vehicle	Vehicle inspection; licensing records
$t_{\text{out}}, t_{\text{in}}$	Outbound and free-flow inbound travel times (minutes)	GPS traces or timed runs; night-time runs for free flow
Δ_{in}	Extra delay on the inbound return (minutes)	GPS traces through candidate junctions on study evenings
t_{term}	Terminal time at both ends (minutes)	Arrival, loading and departure timestamps
λ	Passenger arrivals at the stage (per hour)	Counts in 15-minute intervals
S, Q, W	Seat supply, queue and waiting time	Seat counts at departure; queue counts; sampled passenger waits

Assumptions. The model assumes a fixed fleet that stays on the route, a constant cycle time within the peak, passengers who board the first vehicle with a free seat, demand that does not respond to the queue, and departures that are not held back to fill. It ignores bunching, overtaking, vehicles leaving the route and passengers who walk away or change mode.

Predicted mechanism. Seat supply falls in proportion to the increase in the cycle time; when it falls below passenger arrivals, the queue grows linearly through the peak, and waiting time rises with the length of the peak.

When the model fails. The model will mislead where operators add vehicles or shorten routes in response to delay, where crews hold vehicles at the stage until full, where bunching rather than average delay drives waits, or where demand falls because passengers switch to walking or motorcycles. In these cases the observed queue will diverge from the prediction, and the diagnostic matrix below is needed to attribute it.

Validation protocol. On each observed evening, the model is run with that evening's measured fleet, cycle times and arrivals, and its predicted maximum queue and mean wait are compared with the observed values. Calibration uses a subset of evenings and validation the remainder; fit is reported as mean absolute error and bias, and the model is judged useful for diagnosis only if it predicts the direction and approximate size of changes in waiting across evenings better than a model that ignores inbound delay.

Transferability. The model needs only fleet, seat, cycle-time and arrival data, so it can be applied to other paratransit systems with central terminals, from minibus taxis to tuk-tuk and jeepney routes, provided the dispatch practice is recorded. Its parameters in this paper are illustrative assumptions.

Field Diagnosis Protocol

As recommended by the independent review (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026), the first phase starts with one route and one stage. Selection requires a SACCO willing to take part, a central stage it serves, and returning vehicles that use the junctions under study; the Globe Roundabout and the Koja area are candidate locations for a Thika Road

route, subject to confirmation in the field. Semi-structured interviews cover drivers and conductors, stage marshals and SACCO dispatchers, passengers and people working nearby, and the road authority, County and traffic officers (Appendix A).

Observation takes place on several comparable weekday evenings using a common clock. At the stage, observers count passengers arriving and seats available in 15-minute intervals and record when each matatu arrives, begins loading, departs and clears the city centre. In parallel, observers or GPS traces follow returning vehicles through the candidate junctions. Rain, crashes, roadworks and changes in traffic control are logged. Appendix A sets out the draft instruments, sampling plan and quality-control procedures.

The diagnosis weighs competing mechanisms rather than assigning a queue to a single cause. Table 10 lists the patterns observers look for, the explanations each pattern is consistent with and the evidence that distinguishes them. Each mechanism's contribution is then estimated where possible: for example, the share of the observed seat shortfall explained by longer inbound cycles (from Equation 16), the share of waiting explained by loading and dispatch time at the stage, and the share of days on which incidents or downstream queues dominate.

Table 10

Diagnostic Matrix for a Passenger Queue at a Central Stage

Observed pattern	Possible explanations	Evidence that distinguishes them
Irregular vehicle arrivals at the stage	Inbound congestion; bunching; dispatch practice at the outer terminal; downstream incident	GPS traces; arrival headways; outer-terminal dispatch times; incident log
Long loading or dwell times	Stage layout; crowding; loading practice; holding vehicles to fill	Arrival, loading and departure timestamps; stage observation
Vehicles fail to clear the city centre	Downstream bottleneck; blocked junctions; enforcement stops	Clearance times; junction observation; vehicle tracking
Queue grows despite regular departures	Demand exceeds the seats the fleet can supply	Passenger arrival counts against seats departing
Waits vary sharply between days	Variable inbound delay; incidents; weather; unreliable supply	Distributions of headways and waits; daily event log

Note. Several mechanisms may act at once. The diagnosis reports the estimated contribution of each, with its uncertainty, rather than a single cause.

Pilot Evaluation Design

A pilot proceeds only when the diagnosis identifies a measured cause, a named operating owner agrees a costed change and comparison stages are identified that behave similarly during the baseline. A pre-analysis plan, specifying the primary and secondary outcomes, the comparison sites, the estimation model and the inference method, will be registered publicly before the change is made. The first test is limited to one operational change where feasible. Its effect is estimated by difference-in-differences:

$$\hat{\delta} = \left(\bar{Y}_{P,1} - \bar{Y}_{P,0}\right) - \left(\bar{Y}_{C,1} - \bar{Y}_{C,0}\right) \tag{18}$$

where $\bar{Y}$ is the mean outcome at the pilot (P) or comparison (C) stage before (0) or after (1) the change. Equivalently, with daily observations and controls X_{it} for weather and incidents,

$$Y_{it} = \beta_0 + \beta_1 \,\mathrm{Pilot}_i + \beta_2 \,\mathrm{After}_t + \delta \left(\mathrm{Pilot}_i \times \mathrm{After}_t\right) + \gamma^{\top} X_{it} + \varepsilon_{it} \tag{19}$$

The primary outcomes are passenger waiting time at the stage and end-to-end journey-time reliability, measured by the buffer-time index (the 95th-percentile journey time minus the mean, divided by the mean). Secondary outcomes are people moved per hour, pedestrian access and spillover to adjacent streets and stages. Figure 2 illustrates the design.

A single treated stage and a single comparison stage make it hard to separate the effect of the change from site-specific shocks, so the design is strengthened in five ways. First, at least two comparison stages are used where feasible. Second, pre-intervention trends are checked with an event-study specification,

$$Y_{it} = \alpha_i + \lambda_t + \sum_{k \neq -1} \delta_k \, \mathbf{1}[t - t_0 = k] \, \text{Pilot}_i + \gamma^\top X_{it} + \varepsilon_{it} \tag{20}$$

where α_i and λ_t are site and day effects, t_0 is the week of the change and the coefficients δ_k for weeks before it should be close to zero. Third, placebo tests assign false intervention dates in the baseline and false treated sites among the comparison stages. Fourth, because daily traffic outcomes are serially correlated, conventional standard errors overstate precision (Bertrand et al., 2004); inference therefore uses outcomes collapsed to pre- and post-period means, standard errors robust to autocorrelation, or permutation tests across comparison sites, as specified in the pre-analysis plan. Fifth, spillovers to the comparison stages and all concurrent changes (roadworks, enforcement campaigns, fare changes) are logged and reported.

Figure 2

Difference-in-Differences Design for a One-Change Pilot

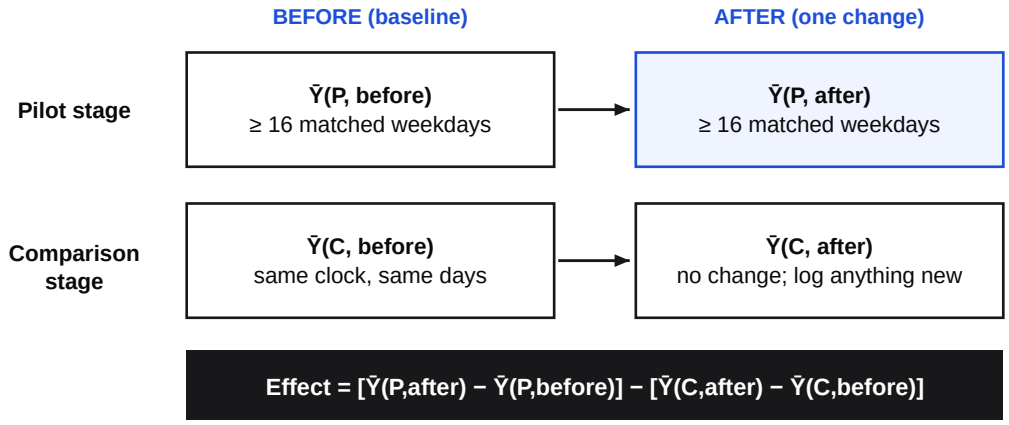

Note. Ȳ = mean daily outcome. Observations at both stages use the same clock and the same days; concurrent changes at either site are logged.

The number of weekdays needed depends on the day-to-day variability of the outcome. With independent days, a two-sided test at $\alpha = 0.05$ with 80% power requires, per cell,

$$n = \frac{4 \, (z_{1-\alpha/2} + z_{1-\beta})^2 \, \text{CV}^2}{\Delta^2} \tag{21}$$

where CV is the coefficient of variation of the daily outcome and Δ the proportional effect to be detected. If daily outcomes at the pilot and comparison stages are differenced first, removing shocks common to both such as rain, the requirement per period becomes

$$n = \frac{2 \, (z_{1-\alpha/2} + z_{1-\beta})^2 \, \text{CV}_d^2}{\Delta^2} \tag{22}$$

with CV_d the coefficient of variation of the daily difference. If daily outcomes follow a first-order autoregressive process with correlation ρ, the effective sample is smaller and the requirement becomes approximately

$$n_\rho = n \, \frac{1 + \rho}{1 - \rho} \tag{23}$$

We assume $\mathrm{CV} = 0.15$, $\mathrm{CV}_d = 0.10$ and ρ between 0 and 0.4 for illustration. These sample sizes are provisional: the final observation period will be set from baseline variance, serial correlation, the characteristics of the comparison sites, anticipated spillovers, missing observations and the inference method, and the assumptions and sensitivity of the calculation will be reported before implementation.

A separate question is whether coordination itself contributed anything. A gain from a package of marshalling, enforcement and layout changes does not isolate the contribution of a coordination forum (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026). We therefore measure coordination through process outcomes recorded in a joint event log: the share of recurring problems with a named owner, the time from a problem being logged to an owner acting, and which joint decisions changed operating practice. These are compared with usual practice in the baseline period. Travel-time effects are attributed to the operational change, not to the forum.

Pilot Business Case Method

The citywide model sizes the burden; it is not a business case for a pilot (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026). A pilot's benefit is computed bottom-up from observed passenger minutes, weighting walking and waiting time relative to in-vehicle time (Wardman, 2004), and netting off any extra walking, transfer or waiting time the change creates:

$$B = D \cdot \frac{\mathrm{VOT}}{60} \sum_{k} w_k \, \Delta t_k \, n_k \tag{24}$$

where n_k passengers each gain (positive) or lose (negative) Δt_k minutes of time type k with weight w_k. Costs are recorded in two groups: operating costs C_{op} (observers, marshals, data and agency staff time) and capital costs C_{cap} (software and equipment). The cost per generalised passenger-hour saved and the benefit–cost ratio are

$$c_h = \frac{C_{\mathrm{op}} + C_{\mathrm{cap}}}{G}, \qquad \mathrm{BCR} = \frac{B}{C_{\mathrm{op}} + C_{\mathrm{cap}}} \tag{25}$$

where G is generalised passenger-hours saved per year. Corridor percentages are never scaled to the whole city.

Assessment of Coordination Systems and the Role of AI

RQ5 is answered in two steps. First, we list the systems that Nairobi's institutions use, or plan to use, to monitor traffic and coordinate their responses, drawing on official documents and press reports. Each system is assessed against five functions that a coordination system must perform: measuring conditions; detecting problems; diagnosing where delay originates; assigning each recurring problem to a named owner and tracking whether it is acted on; and evaluating and reporting outcomes for passengers. A function is recorded as present only where a public source describes it, and otherwise as planned or not documented. Second, for each function we specify how AI could support it, the evidence so far, the simple baseline it must outperform before it is adopted and who retains the decision, drawing on the literature reviewed above and on the NIST (2023) framework. Each function is also classified by readiness (available technology, adaptation needed, research and development) and by whether it needs AI at all, because functions such as routing a problem to its owner, escalation deadlines and audit trails may be better served by transparent rule-based software.

Ethics, Approvals and Data Protection

No primary data have been collected for this paper, and no ethics approval has yet been sought. Before any fieldwork, the protocol will be submitted to an Institutional Scientific and Ethics Review Committee accredited by the national research regulator, and a research licence will be obtained from that regulator, as Kenyan law requires for research in the country (National Commission for Science, Technology and Innovation [NACOSTI], n.d.). The approving committee, the approval reference and the licence number will be reported in subsequent papers.

Participants will receive information in English or Kiswahili, give informed consent before interviews, may decline any question and may withdraw at any time without consequence. Interviews will be audio-recorded only with consent; recordings will be transcribed, de-identified and deleted after verification. Passenger and vehicle counts are aggregate, and no faces, number plates or individual trajectories will be recorded or published; published cells with fewer than about 15 people will be suppressed. Field staff will work in pairs at stages and junctions, avoid live traffic lanes and stop observation if conditions become unsafe.

The project will apply privacy-by-design and data-minimisation principles. Before any data collection or system deployment, it will establish the lawful basis for processing, the purpose of each data stream, access controls, encryption, retention and deletion schedules, and procedures for assessing re-identification risks. Aggregation and anonymisation will be applied where appropriate, but the project will not assume that aggregated or pseudonymised data, such as mobile network or GPS records, are free of privacy risk. Where processing is likely to result in a high risk to the rights and freedoms of data subjects, a data protection impact assessment will be completed as section 31 of the Data Protection Act requires (Data Protection Act, 2019), and institutional data-sharing agreements will be signed before operational deployment. Results will be published whether or not they favour the intervention or its partners.

Results

Mobility Patterns

Walking and matatus together account for 69.4% of the 14.1 million daily person-trips in the metropolitan area (Figure 3). Cars, taxis and trucks account for 13.2%, buses for 9.0% and motorcycles for 7.6%. Because a matatu carries on average 11.2 passengers and a car about 1.5, a vehicle-based view of the road systematically under-weights the people most affected by delay.

Figure 3

Share of Person-Trips by Main Mode, Nairobi Metropolitan Area, 2023

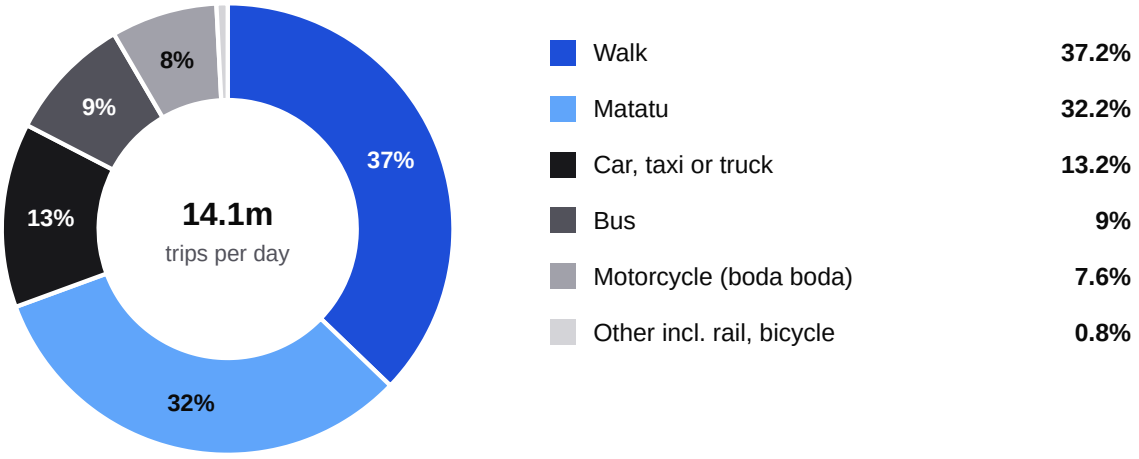

Note. Data from JICA Expert Team (2025), household interview survey. “Other” combines commuter rail (0.2%), bicycle (0.2%) and other modes (0.4%).

Average speeds roughly halved between 2013 and 2023 (Figure 4). In Nairobi City the all-day average fell from 40.0 to 21.2 km/h and the peak average from 27.5 to 8.3 km/h. JICA’s forecasts (Figure 5) show speeds on major roads falling to 6.8 km/h by 2045 under a do-nothing scenario and to 11.2 km/h even with the planned road and BRT programme (JICA Expert Team, 2025).

Figure 4

Average Road Speeds in Nairobi, 2013 and 2023

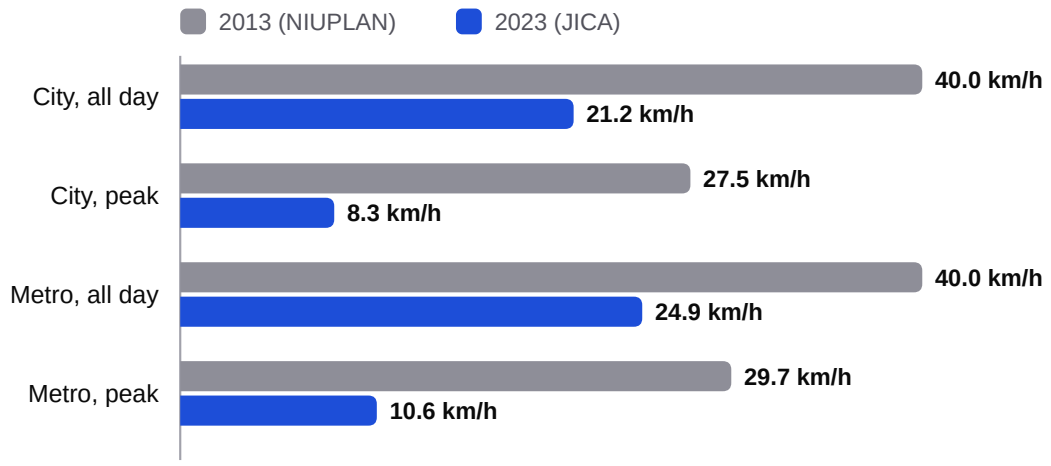

Note. Data from JICA Expert Team (2025), comparing the 2013 NIUPLAN survey with the 2023 survey. Metro = Nairobi Metropolitan Area.

Figure 5

Forecast Average Speed on Major Roads, 2023–2045 (km/h)

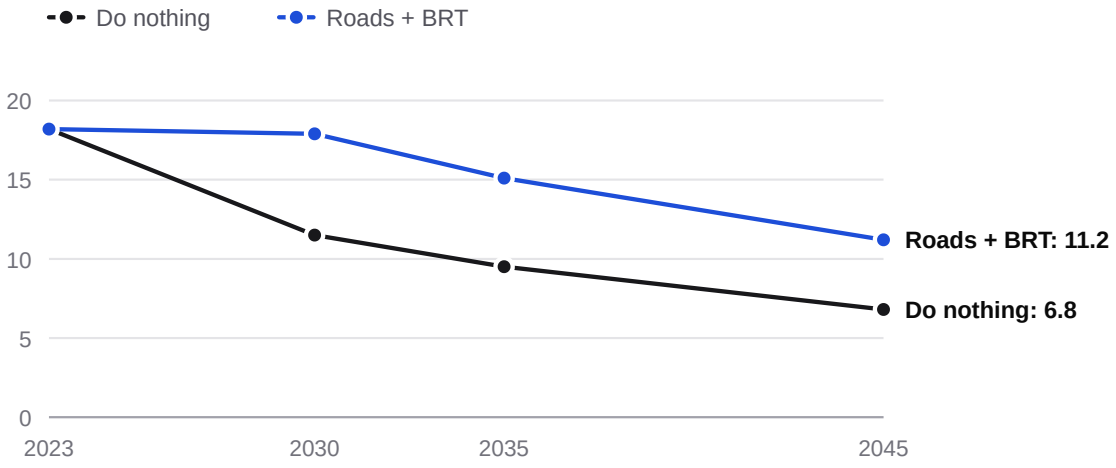

Note. Model scenarios from JICA Expert Team (2025, Table 7.3). The “Roads + BRT” scenario includes the NIUPLAN road programme and bus rapid transit.

Fleet Growth and Road Safety

New motorcycle registrations doubled from 118,308 in 2024 to 241,763 in 2025, while registrations of cars, vans, lorries and buses were almost flat (Figure 6). Vulnerable road users dominate road deaths: pedestrians, motorcyclists, pillion passengers and pedal cyclists made up 3,138 of the 4,195 deaths recorded by road-user type in 2025, or 75% (Figure 7). The WHO’s modelled estimate of 14,926 deaths in 2021 (95% CI [13,002, 16,851]) is roughly three times the totals recorded by police in recent years (Bett, 2024; WHO, 2023), which suggests that recorded crash data substantially under-count harm.

Figure 6

New Vehicle Registrations in Kenya, 2024 and 2025

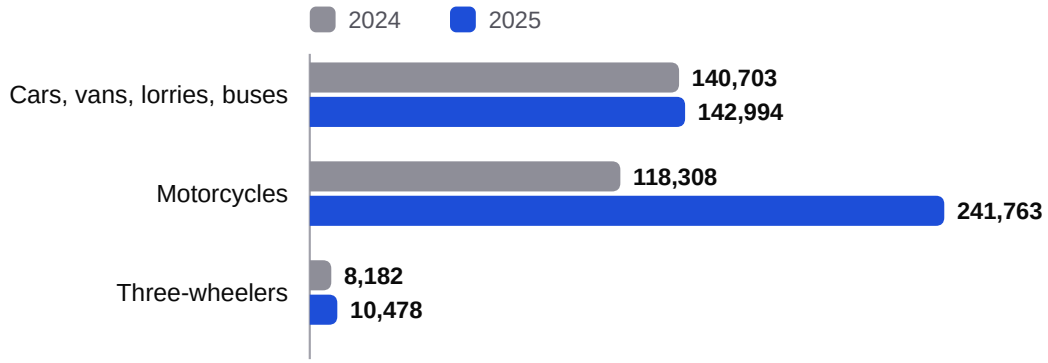

Note. Data from the Kenya National Bureau of Statistics *Economic Survey 2026*, as reported in "Bikes Boom Drives Spike" (2026).

Figure 7

Road Deaths by Road-User Type, Kenya, 2025

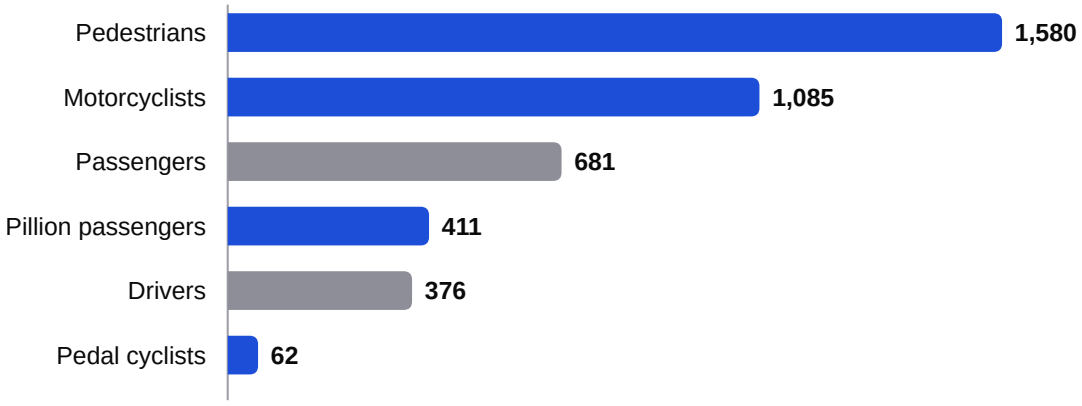

Note. Data from NTSA as reported in "Road Accident-Related Deaths Rise" (2025). Blue bars denote vulnerable road users.

The Cost of Peak Congestion

Applying Equations 1–5 with central inputs gives the following worked calculation:

Total trips: 8,852,297 ÷ (1 − 0.372) = 14,096,014 per day
P = 14,096,014 × (0.322 + 0.090 + 0.132) = 7.7 million trips; peak trips Ps = 3.8 million
Free-flow time = 26.78 ÷ 1.496 = 17.9 min per 10 km; mean peak delay δ_{10} = 14.7 min, so d = 0.245 h
VOT = 0.30 × 988,200 ÷ 2,000 = KSh 148.2 per hour
Person-hours lost = 3.83 million × 0.245 = 0.94 million per weekday
Time cost = 0.94 million × 148.23 × 250 = KSh 34.8 billion
Fuel = about 50 million litres × KSh 227 = KSh 11.4 billion
Total = KSh 46.3 billion per year (≈ US$ 356 million)

Table 11 reports the three scenarios, and Figure 8 shows the composition of the central estimate. Time accounts for 75% of the total. For an individual making two peak trips a day, the delay amounts to about 122 hours a year, or 15.3 eight-hour working days, valued at about KSh 18,154 at the central value of time. These figures are estimated from secondary data under stated assumptions. They do not include pedestrian delay or the full costs experienced by motorcycle users, so the model cannot by itself show how the burden is distributed across modes and income groups.

Table 11

Estimated Annual Cost of Weekday Peak Congestion, Nairobi Metropolitan Area (KSh Billion)

Component	Low	Central	High
Value of time lost	11.5	34.8	102.5
Fuel burned in delay	4.4	11.4	26
Total	**15.8**	**46.3**	**128.4**
Fuel at resource cost (excluding taxes and levies)	2.7	7.1	16.1
Total with fuel at resource cost	14.2	41.9	118.6
Person-hours lost per weekday (millions)	0.46	0.94	1.66

Note. Scenario inputs are listed in Table 7. Values are in 2026 shillings and cover weekday peaks only. The headline total values fuel at the pump price, the private cost; the resource-cost rows exclude taxes and levies, which are transfers within the economy, and are the appropriate basis for economic appraisal. The low and high columns are scenario envelopes, not confidence limits.

Figure 8

Composition of the Central Cost Estimate

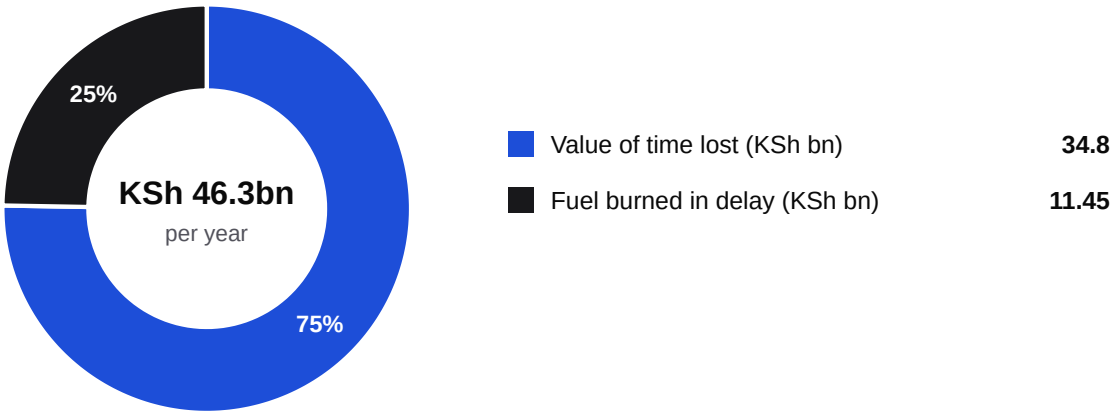

Note. Fuel is valued at the pump price, including taxes and levies. At resource cost the fuel component is KSh 7.1 billion and the total KSh 41.9 billion.

Uncertainty Analysis and Cross-Checks

The value-of-time share has by far the largest influence on the estimate: moving it from 20% to 50% of the wage moves the total from KSh 34.7 billion to KSh 69.5 billion (Figure 9). Peak delay, the share of trips in the peaks and trip length each move the total by roughly ±20%. Fuel-burn assumptions move it by less than ±4%.

Figure 9

One-at-a-Time Sensitivity of the Central Estimate (KSh Billion per Year)

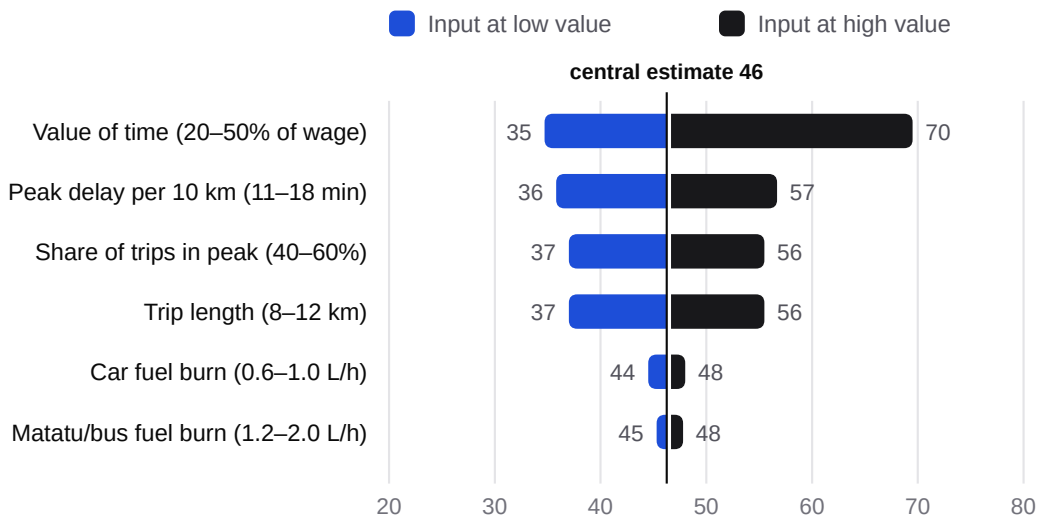

Note. Each input is moved to its low and high value while the others are held at central values. The vertical line marks the central estimate of KSh 46.3 billion.

The Monte Carlo analysis (Figure 10) gives a median annual cost of KSh 49.2 billion (mean 50.2) and a 90% uncertainty interval of KSh 33.6–70.8 billion. The central scenario, KSh 46.3 billion, lies at about the 40th percentile, because the value-of-time distribution is skewed towards higher values. Under these assumptions the probability that the cost exceeds KSh 40 billion is about 80.4%, exceeds KSh 50 billion about 47% and exceeds KSh 75 billion about 2.8%. The extremes of the scenario envelope, which require every input to sit at its bound at once, lie outside the simulated range. The value of time accounts for about 39% of the output variance and peak delay for 35%; the share of trips in the peaks (13%) and trip length (12%) account for most of the rest, and fuel burn for almost none. Person-hours lost per weekday have a 90% interval of 0.69–1.2 million. These intervals describe uncertainty in the stated inputs only; they do not capture uncertainty in the model’s structure, such as the transfer of car-based delay to matatu passengers.

Figure 10

Distribution of the Annual Cost of Weekday Peak Congestion in 10,000 Monte Carlo Simulations

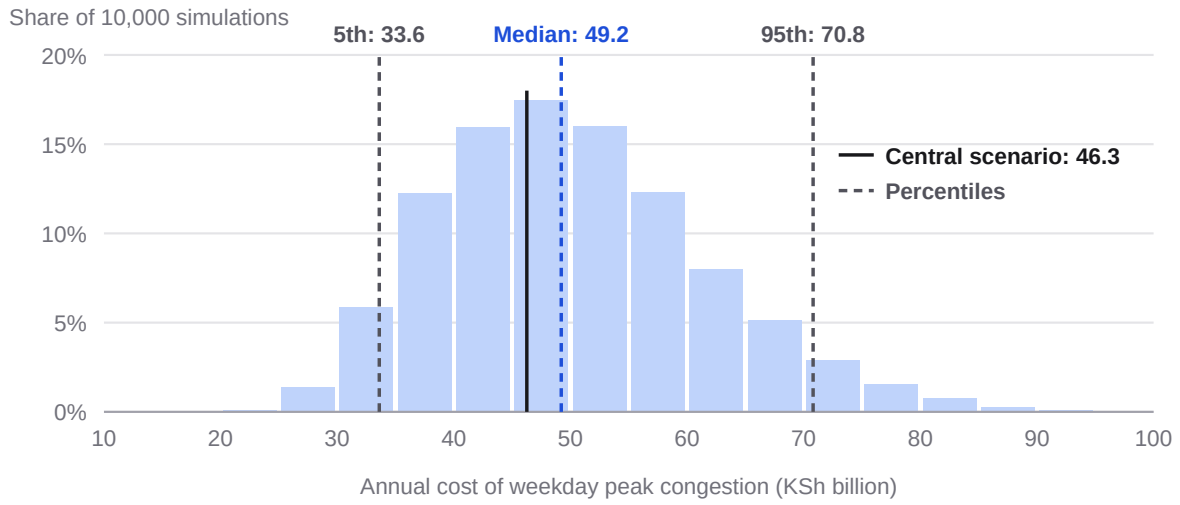

Note. Input distributions are listed in Table 8. Bins are KSh 5 billion wide. Dashed lines mark the 5th percentile, median and 95th percentile; the solid line marks the central scenario.

The top-down cross-check (Equation 6), with 1.0 million commuters materially delayed, 1.0 hour lost per commuter per weekday and the central value of time, gives KSh 37.1 billion a year. Its implied 1.0 million person-hours per weekday is close to this study's 0.94 million. Varying the number of commuters between 0.7 and 1.3 million and the hours lost between 0.75 and 1.5 gives KSh 19.5–72.3 billion. The two formulations therefore agree on the physical quantity, and both are sensitive to assumptions that field data must replace. Figure 11 places both alongside published estimates, which range from KSh 1.9 billion (2008 prices) to KSh 120 billion a year ("Effect of Motor Vehicle Congestion," n.d.; Kenya Institute for Public Policy Research and Analysis [KIPPRA], n.d.; Mwita, 2025; "Nairobi Traffic a Health, Economic Burden," 2021; "Traffic Jams Cost Nairobi," 2012; UN-Habitat, n.d.). Valuing fuel at resource cost lowers the central estimate to KSh 41.9 billion; adding 20% for weekend and off-peak delay raises it to KSh 55.5 billion.

Figure 11

Published and Modelled Estimates of the Annual Cost of Congestion in Nairobi (KSh Billion)

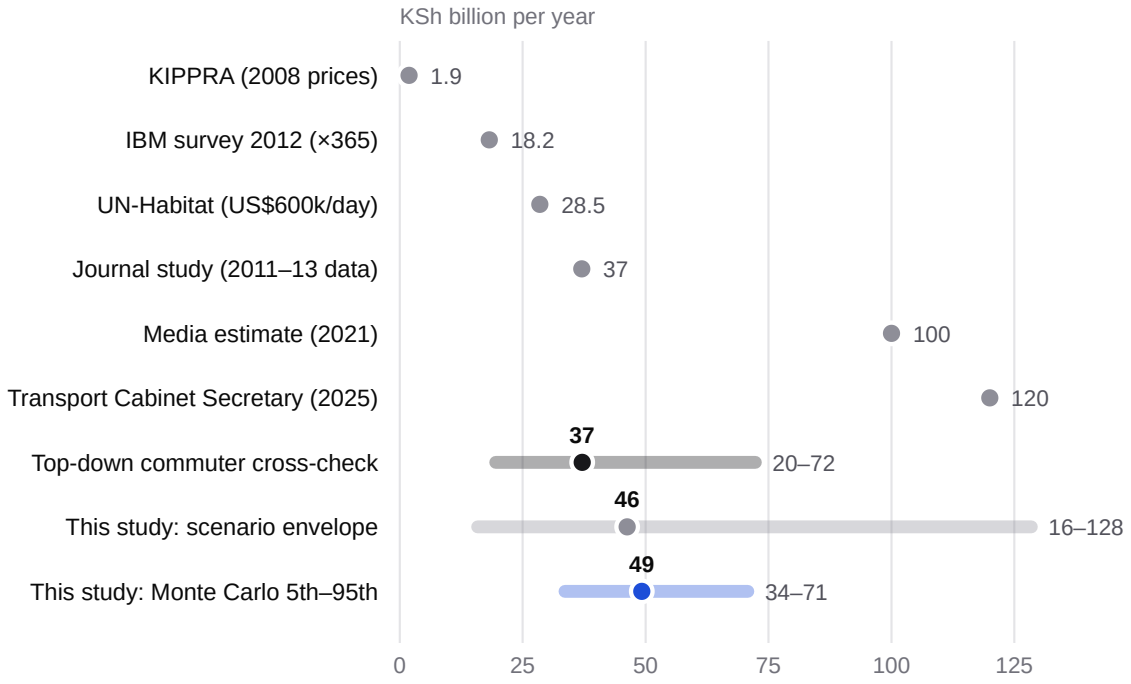

Note. Dots are point estimates; bars show modelled ranges. Daily figures are multiplied by 365 and are nominal, not inflation-adjusted. The UN-Habitat figure is converted at the November 2026 exchange rate.

Congestion Engineering Results

With JICA's tentative Nairobi parameters (α = 2.0, β = 3.0), Equation 7 implies that a trip already takes 1.8 times its free-flow time at the 2023 volume-to-capacity ratio of 0.72, rising to 4.5 times by 2045 (Figure 12). Using textbook parameters on a 10 km link with a 12-minute free-flow time, Equation 8 gives a marginal external delay of 7.2 vehicle-minutes at V/C = 1.0 and 14.9 at 1.2, which corresponds to a Pigouvian toll of about KSh 27 per car trip at capacity. Equation 9 with μ = 2,000, λ = 2,200 and λ_2 = 1,600 vehicles per hour gives 150 vehicle-hours of delay for a 60-minute overload and 37.5 for a 30-minute overload, a 75% reduction (Figure 13). Equation 10 with 12 seconds of lost time and Y = 0.75 gives an optimal cycle of 92 seconds.

Figure 12

Travel Time Relative to Free Flow as Roads Fill: Textbook and Nairobi Parameters

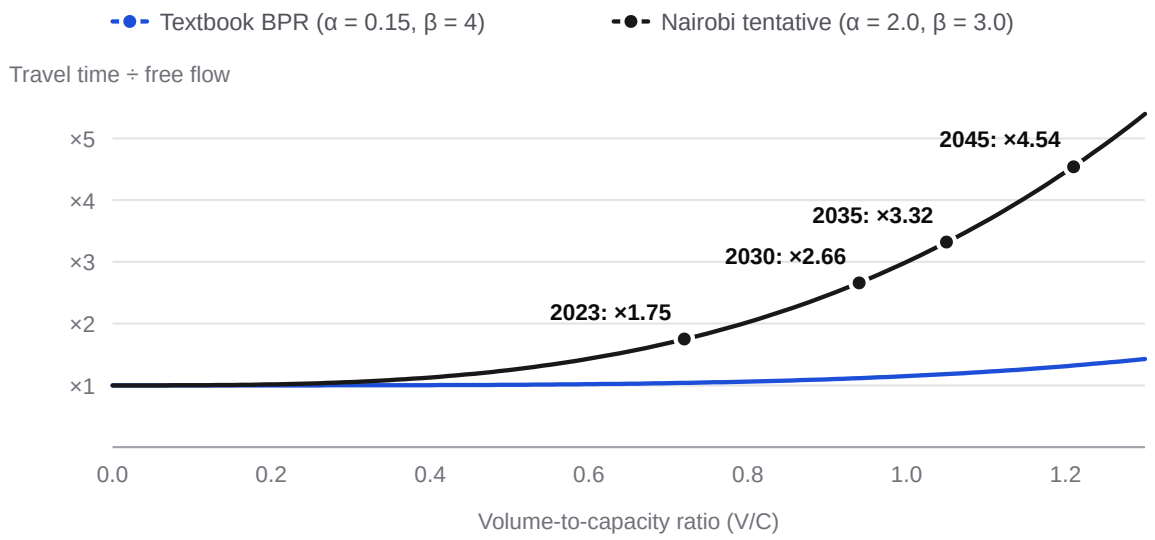

Note. Points mark JICA's do-nothing volume-to-capacity ratios for 2023, 2030, 2035 and 2045 under the Nairobi parameters.

Figure 13

Vehicles Queued at a Bottleneck When the Overload Lasts 60 or 30 Minutes

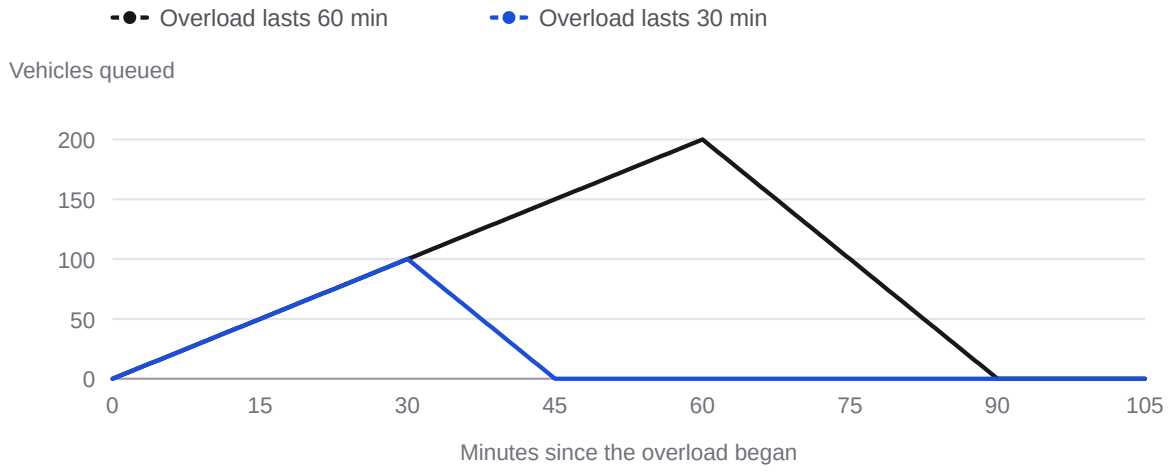

Note. Deterministic queueing model with capacity 2,000 vehicles per hour, arrivals of 2,200 per hour during the overload and 1,600 per hour afterwards. Illustrative values.

Figure 14 illustrates passenger throughput across selected modes and systems. The first three values are calculated scenarios for a single lane: a protected bus lane carrying 120 buses of 60 passengers would move 7,200 people an hour, 2.7 times a car lane, and mixed-traffic matatus about 3,360. The commuter-rail value is a system target and the TransMilenio value a published peak load, so the bars are not measured on the same basis and indicate orders of magnitude only.

Figure 14

Illustrative Passenger Throughput Across Selected Transport Modes and Systems (People per Hour)

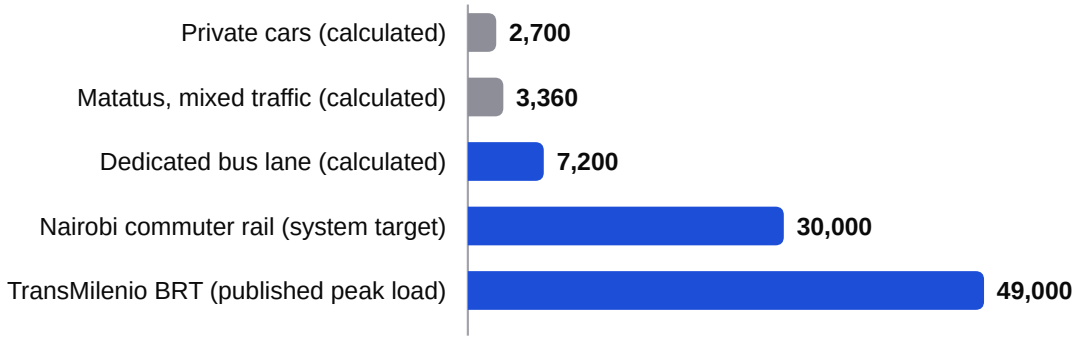

Note. Measurement type is shown in each label. Calculated scenarios use 1,800 cars per lane at 1.5 occupants, 300 matatus at 11.2 passengers and 120 buses at 60 passengers per hour; commuter rail target from "UK and G7 Back Nairobi Railway City" (2025); TransMilenio peak load from BRTData (n.d.).

Applying the Economists' Models to Nairobi

A Vickrey Bottleneck on a Nairobi Approach

To show what the bottleneck model implies at Nairobi's scale, we apply Equations 12–14 to an illustrative two-lane approach to the CBD with capacity s = 3,600 vehicles per hour and N = 7,200 vehicles whose users wish to arrive at 08:00. The value of travel time is the central value of KSh 148.23 per hour, and the early and late penalties follow Small's approximate ratios (Small, 2015). Capacity, demand and the desired arrival time are assumptions chosen for illustration.

β = 0.5 × 148.23 = KSh 74.11 per hour; γ = 2 × 148.23 = KSh 296.46 per hour
δ = (74.11 × 296.46) ÷ (74.11 + 296.46) = KSh 59.29 per hour
Rush period N/s = 7,200 ÷ 3,600 = 2 hours, with arrivals from 06:24 to 08:24
Private cost per trip = 59.29 × 2 = KSh 118.6 at every arrival time
Longest queue T_{max} = 118.6 ÷ 148.23 = 0.8 hours = 48 minutes, or 2,880 vehicles
Total cost without a charge TC_0 = 59.29 × 7,200² ÷ 3,600 = KSh 0.85 million per morning, half of it queuing
Optimal charge: peaks at KSh 118.6 at 08:00 and removes KSh 0.43 million of queuing per morning, about KSh 106.7 million a year over 250 weekdays

Figure 15 shows the result. Whatever time a commuter arrives, the sum of queuing cost and the cost of arriving early or late is the same KSh 118.6. A commuter who leaves earlier to beat the queue exchanges queuing for early arrival and gains nothing, which, within the model, is what the claim that individual effort cannot beat a jam means. The optimal charge has exactly the shape of the queuing cost: it replaces the queue with a price, so the 48 minutes of queuing at 08:00 become a charge of KSh 118.6, and the approach carries the same number of vehicles without a queue. Each commuter still pays KSh 118.6 in total. The gain of KSh 0.43 million per morning arrives as revenue, and whether commuters feel it depends on how that revenue is used (Arnott et al., 1993; Small, 2015).

Figure 15

Cost per Vehicle Trip by Arrival Time at an Illustrative Nairobi Bottleneck

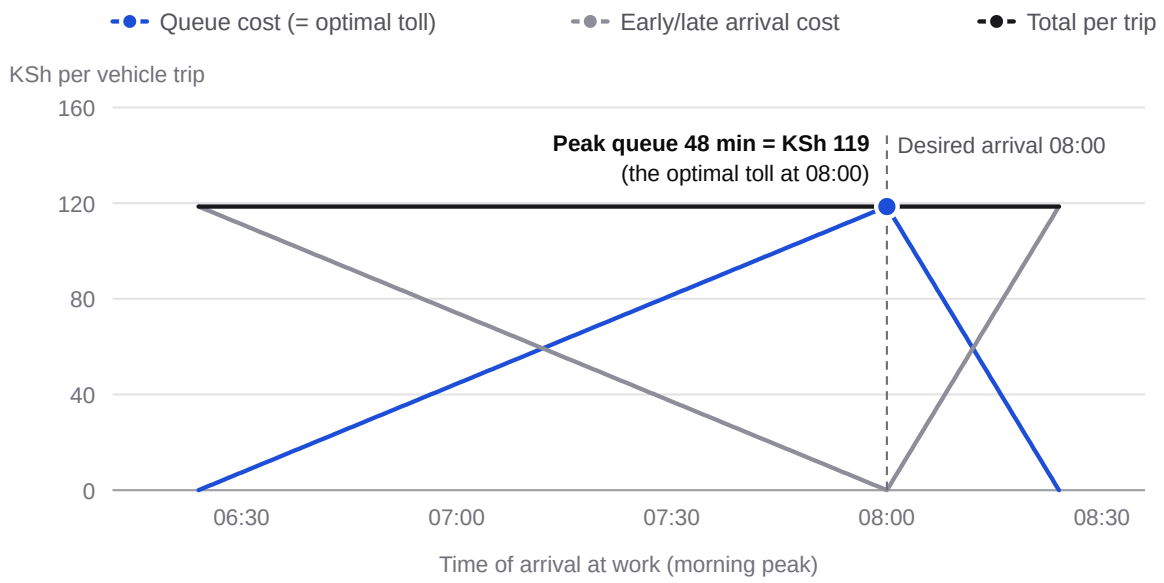

Note. Equations 12–14 with s = 3,600 vehicles per hour, N = 7,200, desired arrival 08:00, α = KSh 148.23 per hour, $\beta = 0.5\alpha$ and $\gamma = 2\alpha$. Without a charge, commuters pay the blue line in queuing time; with the optimal charge they pay the same amount as a toll and do not queue. Illustrative values.

The same model illustrates Downs' triple convergence. Because total queuing cost rises with N^2/s, a 25% increase in capacity alone reduces the annual queuing loss on the approach from KSh 106.7 million to KSh 85.4 million. If the added capacity draws 20% more vehicles into the peak, from other routes, other times and other modes, the loss rises to KSh 122.9 million, above where it started. The gain from added capacity therefore lasts only if demand does not converge, which the evidence on induced demand suggests is unlikely in the long run (Duranton & Turner, 2011). A late penalty of four times the value of time, perhaps more realistic for workers who lose pay when late, raises the loss to KSh 118.6 million.

The Bathtub: Where Extra Vehicles Reduce Flow

Equation 15, with an illustrative free-flow speed of 30 km/h on central streets and a jam density of 150 vehicles per kilometre per lane, gives a critical density of 75 vehicles per kilometre and a maximum completion rate of about 1,125 vehicles per hour per lane (Figure 16). At 90 vehicles per kilometre, 20% above the critical density, completions fall to 1,080, the same as at 60 vehicles per kilometre: the extra vehicles add delay and nothing else. If Nairobi's CBD operates on this hypercongested branch in the evening peak, holding back even a modest share of inbound vehicles would raise the number of completed trips. We record this as hypothesis H2, testable with speed and density data from probe vehicles and counts.

Figure 16

The Vickrey Bathtub: Trips Completed per Hour Rise With Density to a Critical Level, Then Fall

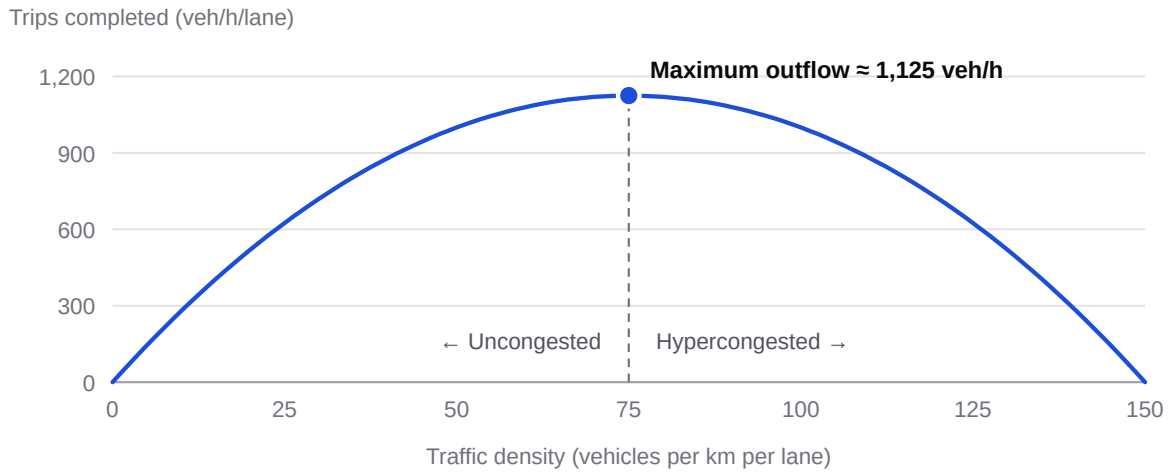

Note. Equation 15 with Greenshields' speed–density relation, a free-flow speed of 30 km/h and a jam density of 150 vehicles per kilometre per lane. Illustrative values.

Stage Supply Scenarios

Table 12 and Figure 17 apply Equations 16 and 17 to an illustrative route: 30 matatus of 14 seats, a loaded outbound trip of 32.6 minutes, a free-flow inbound return of 17.9 minutes, 30 minutes of terminal time, and 300 passengers arriving per hour over a three-hour evening peak. With no extra inbound delay the route supplies 313 seats an hour and no queue forms. An extra 20 minutes on the inbound return cuts supply to 251 seats an hour; the queue then grows by 49 passengers an hour and reaches 148, the last arrivals wait about 35 minutes and the average extra wait is about 18 minutes. Nothing at the stage has changed.

Table 12

Seat Supply and Passenger Queues at a Central Stage Under Different Inbound Delays

Extra inbound delay	Cycle (min)	Seats per hour	Queue growth per hour	Queue after 3 h	Maximum extra wait (min)	Mean extra wait (min)
None	80.5	313	0	0	0	0
10 min	90.5	278	22	65	13.9	7
20 min	100.5	251	49	148	35.4	17.7
30 min	110.5	228	72	216	56.8	28.4

Note. Illustrative assumptions: 30 vehicles, 14 seats, 300 passengers per hour, three-hour peak. Waiting times exclude the normal wait between departures.

Figure 17

Passengers Queuing at a Central Stage by Extra Delay on the Inbound Return Trip

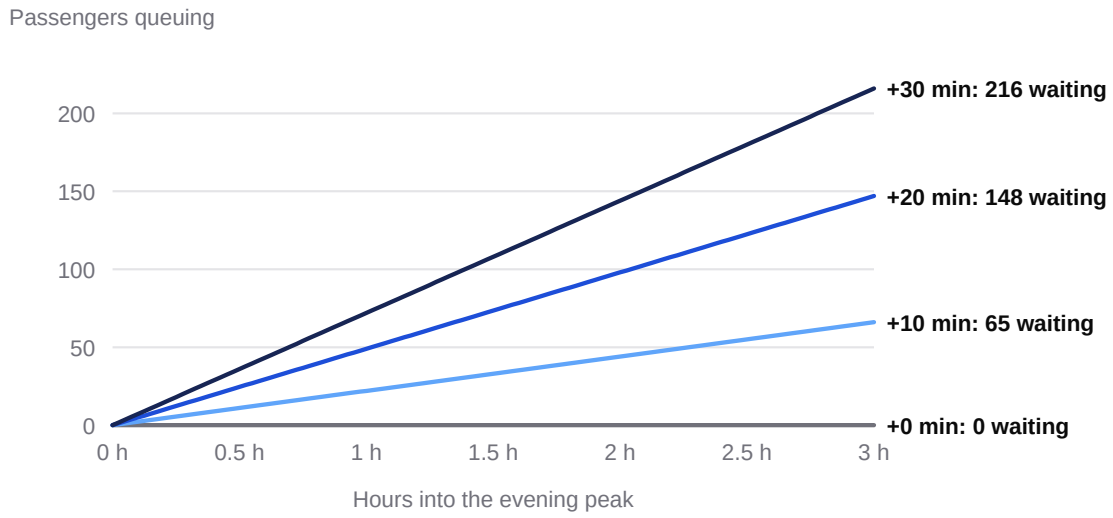

Note. Deterministic model (Equations 16–17) with the assumptions in Table 12.

Committed Investment

Announced commitments to traffic technology, enforcement and bus rapid transit total KSh 112.7 billion for clearly distinct projects and KSh 136.7 billion if the China Exim loan, whose relationship with the KURA phases is not publicly documented, is counted separately (Table 13; Figure 18) (“Eyes on Cabinet,” 2026; “Government to Roll Out AI-Powered Traffic Lights,” 2026; “Kenya Signs Ksh24 Billion Loan,” 2025; Kipkemoi, 2026; “KURA Signs Deal,” 2024; KURA, 2026; “NTSA to Install 1,000 Speed Cameras,” 2026; “Team Europe to Build Sh47bn Green BRT,” 2023; The National Treasury, 2025). These are announcements, loans, grants and partnership commitments, not disbursements, and their status ranges from contracts under construction to budget allocations. None of the announcements specifies an independent before-and-after measurement of passenger journeys.

Table 13

Project-Level Registry of Announced Commitments

Project	KSh bn	As announced	Funding type	Status	Overlap risk
BRT Line 3	50.8	€347.6m	EU grant; EIB and AFD loans; Government of Kenya	Financing announced; target 2030	Separate project
NTSA speed cameras and instant fines	42	KSh 42bn	Public–private partnership	Announced, February 2026	Needs verification
Nairobi ITS loan, China Exim Bank	24	US$185m	Concessional loan	Agreement signed, November 2025	Possible overlap with ITS phases
ITS Phase II (60 junctions)	10.8	US$83.8m	Loan (Korea EDCF)	Tender issued, June 2026	Unknown
ITS Phase I (25 junctions and Traffic Management Centre)	7.9	KSh 7.9bn	Contract (Samsung C&T)	Under construction, 2026	Low
ITS Phase III (125 junctions)	1.2	KSh 1.18bn	Budget allocation	Planned	Unknown
Distinct projects	**112.7**	Excludes the China Exim loan			
Headline maximum	**136.7**	Includes items whose overlap cannot yet be resolved			

Note. Values as announced, converted at Central Bank of Kenya rates of 5 November 2026. AFD = Agence française de développement; EDCF = Economic Development Cooperation Fund; EIB = European Investment Bank; ITS = intelligent transport system.

Figure 18

Announced Commitments to Traffic Technology, Enforcement and Bus Rapid Transit (KSh Billion)

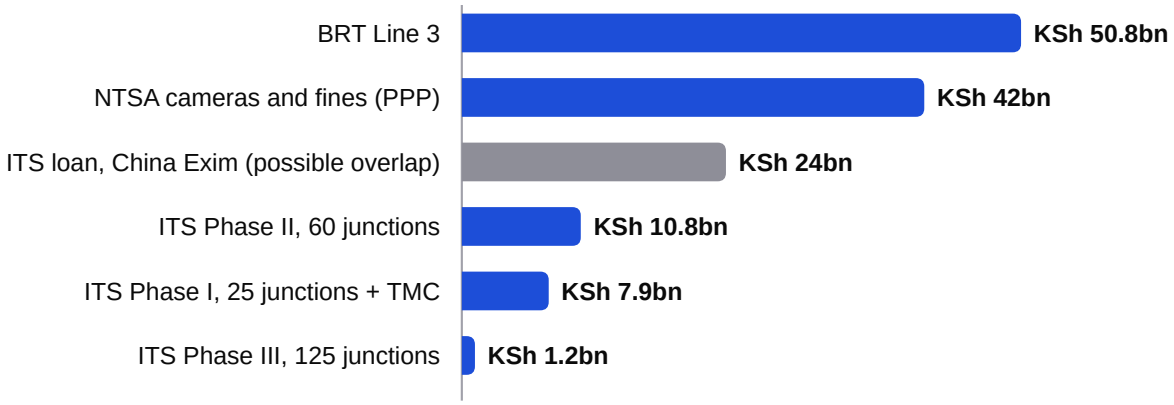

Note. Commitments as announced, not disbursements. The grey bar marks the loan whose relationship with the KURA phases is not publicly documented.

Existing Monitoring and Coordination Systems

Table 14 answers the first part of RQ5. Nairobi has, or is building, several systems that measure traffic and detect problems: junction cameras and detectors feeding the Traffic Management Centre, speed cameras, mobile network data, commercial probe speeds, crowd-sourced reports and open route data. None of the documented systems, however, diagnoses where delay originates, assigns recurring problems across institutions to a named owner, tracks whether that owner acted or reports what changed for passengers. The Traffic Management Centre is the nearest to a coordination hub, but its documented scope is signalised junctions under KURA, while stages and termini, SACCO dispatch, junctions on national highways and enforcement sit with other institutions. Nairobi is acquiring the means to see its traffic faster than the means to turn what it sees into accountable action.

Table 14

Monitoring and Coordination Systems in Nairobi Assessed Against Five Coordination Functions

System (owner)	Measure	Detect	Diagnose origin	Assign and track	Evaluate and report
Intelligent transport system and Traffic Management Centre (KURA)	Yes, at signalised junctions	Planned: AI cameras in Phase III	Not documented	Not documented	Not documented
Speed cameras and instant fines (NTSA)	Partly: speeds at camera sites	Yes: traffic violations	No	Individual drivers only	Not documented
County Transport and Safety Committee (County, NTSA, police, operators)	No	No	Not documented	Forum exists; no public tracking	Not documented
Mobile network data (Safaricom and KNBS)	Planned: population movement	No	No	No	No
Commercial probe data (TomTom, Google)	Yes: road speeds	Partly: incidents	No	No	No
Crowd-sourced reports (Ma3Route, social media)	Partly	Yes: incidents and crashes	No	No	No
Open matatu route data (Digital Matatus)	Routes and stops only	No	No	No	No

Note. “Not documented” means that no public source describing the function was found; it does not mean that the function is absent. Sources: (“Government to Roll Out AI-Powered Traffic Lights,” 2026; “KURA Initiates Plan,” 2025; Milusheva et al., 2021; Nairobi City County, 2024; “NTSA to Install 1,000 Speed Cameras,” 2026; “Ruto Unveils Ksh7.9B Smart Traffic System,” 2026; TomTom, 2026; Williams et al., 2015; Wolfenden, 2026). KNBS = Kenya National Bureau of Statistics; KURA = Kenya Urban Roads Authority; NTSA = National Transport and Safety Authority.

Evaluation Requirements and the Pilot Business Case

Equations 21 and 22 show how strongly the design affects the observation effort (Figure 19). To detect a 10% change in mean waiting time with independent days and a coefficient of variation of 0.15 requires 71 weekdays in each of the four cells. Differencing each day’s outcome against a well-matched comparison stage, with an assumed coefficient of variation of 0.10, reduces this to 16 weekdays before and 16 after. A 5% effect would need 63 paired weekdays per period. Several evenings are therefore enough for diagnosis, but evaluation needs a longer baseline, consistent with the review’s advice (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026). If daily outcomes are serially correlated, the requirement rises (Equation 23): with a day-to-day correlation of 0.2, detecting a 10% change in the paired design needs about 24 weekdays per period, and with 0.4 about 38. These figures are provisional and will be recalculated from the baseline before any intervention.

Figure 19

Weekdays of Observation Needed to Detect a Given Proportional Effect

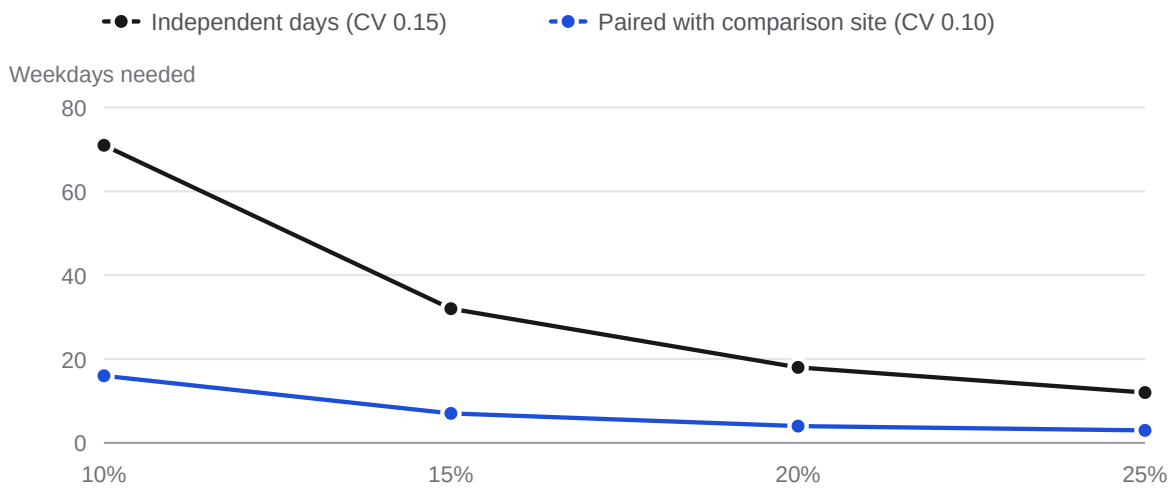

Note. Two-sided α = .05, power = .80. Coefficients of variation are assumptions to be replaced by baseline estimates.

Table 15 applies Equations 24 and 25 to a hypothetical stage serving 1,500 passengers in the evening peak, where a change saves 5 minutes of waiting but adds 3 minutes of walking for 300 passengers. With a weight of 2.0 on waiting and walking time the change saves 55,000 generalised passenger-hours a year, worth KSh 8.2 million against first-year costs of KSh 3.3 million. If only 2 minutes of waiting are saved, the benefit–cost ratio falls to 0.79 (Figure 20). The business case rests on measured minutes.

Table 15

Illustrative Pilot Business Case at One Stage

Item	Wait cut 5 min, weight 2.0	Wait cut 5 min, weight 1.5	Wait cut 2 min, weight 2.0
Generalised minutes saved per day	13,200	9,900	4,200
Generalised passenger-hours saved per year	55,000	41,250	17,500
Annual benefit (KSh million)	8.2	6.1	2.6
Operating costs: marshals (KSh million)	0.75	0.75	0.75
Operating costs: observers and analysis	1.20	1.20	1.20
Operating costs: agency staff time (in kind)	0.40	0.40	0.40
Operating costs: data and communications	0.15	0.15	0.15
Capital costs: software and equipment	0.80	0.80	0.80
Cost per generalised passenger-hour (KSh)	60	80	189
Benefit–cost ratio, first year	2.5	1.9	0.79

Note. All quantities are hypothetical and illustrate the method; they are not forecasts. Value of time KSh 148.23 per hour; 250 weekdays. Weights follow the convention reviewed by Wardman (2004).

Figure 20

Illustrative Annual Benefit and First-Year Cost of a One-Stage Pilot (KSh Million)

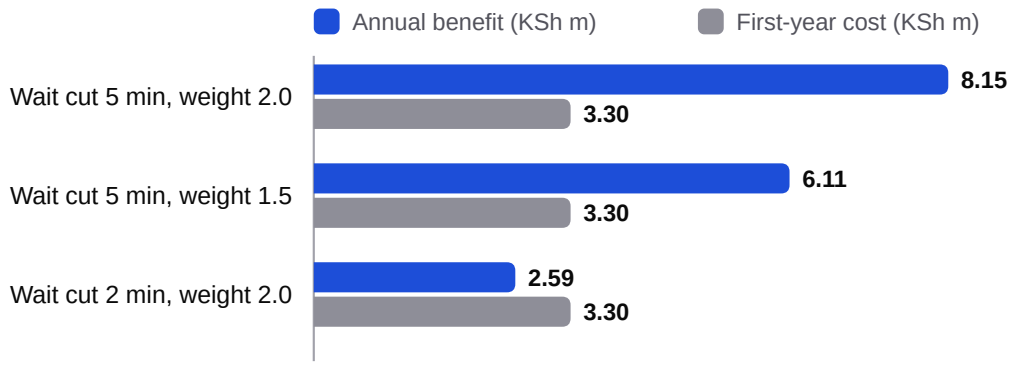

Note. Hypothetical values from Table 15.

Discussion

The Discussion is organised by research question. For each, it states what the paper's evidence shows, what it does not show and what follows, before turning to the paper's contribution to theory, its limitations and the open problems.

RQ1: The Economic Burden

Under central assumptions, weekday peak congestion costs the Nairobi Metropolitan Area about KSh 46.3 billion a year, with a Monte Carlo median of KSh 49.2 billion and a 90% interval of KSh 33.6–70.8 billion. The physical quantity behind it, roughly one million person-hours each weekday, is better established than its monetary value: two independent formulations agree on it, whereas the value placed on an hour drives most of the uncertainty. What policymakers can responsibly infer is that the burden is in the tens of billions of shillings a year and that a Nairobi value-of-time survey and mode-specific delay data would do most to narrow the range. What they cannot infer is the distribution of the burden across modes and income groups, which the model does not measure, or the value of any specific project.

The estimate lies within the range of published figures but well below the highest (Figure 11). The lower historical estimates partly reflect a smaller fleet and a stronger shilling; the highest recent figure, KSh 120 billion, was cited without a published method (Mwita, 2025). The almost 60-fold spread among published estimates is itself a finding: without transparent methods, estimates cannot be compared or improved, which is why the model, its inputs and its random seed are published.

The citywide estimate sizes the burden; it should not be used to value a pilot. Scaling a corridor percentage to the whole city would assume that conditions, mechanisms and responses are uniform across the network, an assumption the stage-supply analysis gives reason to doubt. A pilot is judged instead by observed passenger minutes saved net of any extra walking, transfers or waiting, against its operating and capital costs recorded separately (Table 15). The citywide value of a 10% cut in peak delay, about KSh 4.6 billion a year, is relevant only to programmes that operate across the network, such as the intelligent transport system.

RQ2: Mechanisms and Where Queues Originate

The evidence reviewed is consistent with treating Nairobi's peak congestion as more than a capacity shortfall, but it does not establish how much of the delay any one mechanism causes. Figure 21 summarises five systems that act on the road at peak. Infrastructure agencies may optimise for works delivered, operators for daily cash targets, enforcement for visible presence, road users for their own fastest route and planning institutions for annual budgets and multi-year projects; each optimisation can be reasonable on its own terms while the combined result is a jam. Three reinforcing loops are plausible and testable: longer rounds may lead crews to wait longer at stages, which blocks lanes; congestion may push commuters onto motorcycles, increasing filtering and crashes; and new capacity attracts new trips (Duranton & Turner, 2011). Peak spreading, as people leave home earlier, is a balancing loop that reduces congestion at a cost in sleep and family time.

Figure 21

Five Interacting Systems That Shape Peak Road Space

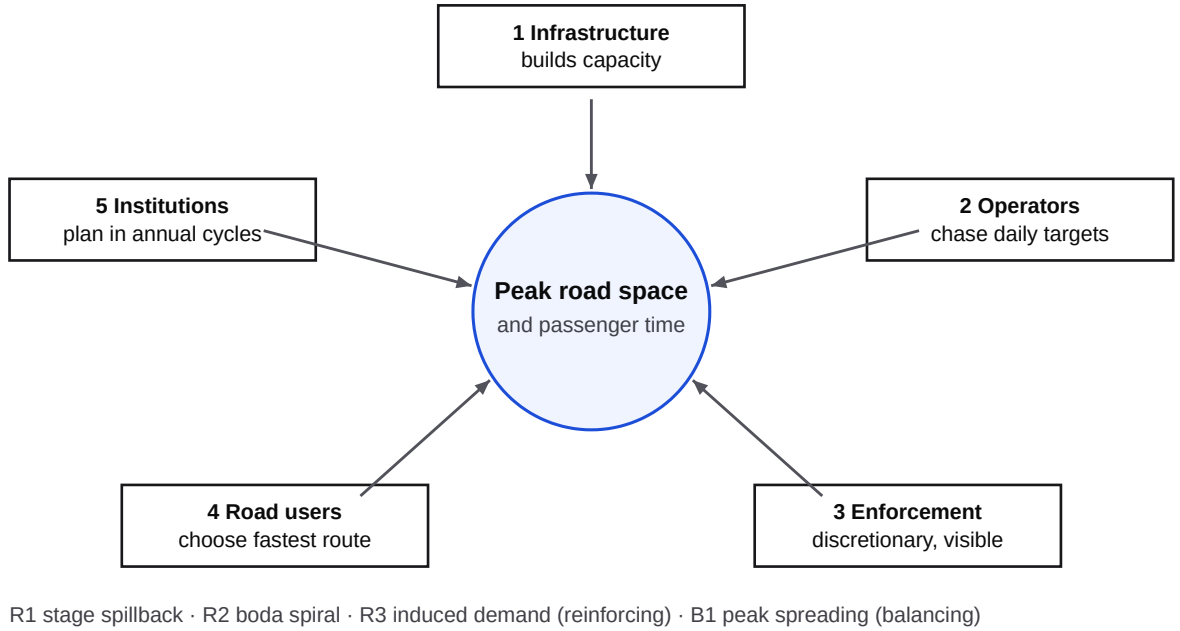

Note. Arrows show each system's effect on peak road space and passenger time. R = reinforcing loop; B = balancing loop.

A time-scale mismatch is a further plausible mechanism. Queues form within minutes and incidents within seconds, but the institutions able to change signal plans, stage layouts or enforcement practices typically review performance far less often. The records of previous interventions (Table 3) are consistent with this: decisions were implemented, disrupted and sometimes reversed without any shared measurement of what passengers experienced.

The stage model shows that, under its assumptions, an inbound delay of 20 minutes could produce a maximum extra wait of about 35 minutes at a central stage (Table 12). It does not show that this happens in Nairobi. A crowd at a central stage is visible and attracts attention, but its origin may lie in delays on vehicles' return trips, in stage operations, in a downstream bottleneck or in demand that exceeds the fleet, and several of these could act together (Table 10). An independent review of an earlier draft raised one hypothesis from field discussion (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026): that delays on the inbound return trip through the Globe Roundabout and the Koja area reduce the supply of seats for passengers leaving the city centre in the evening. We record this as hypothesis H1. It is plausible because nearly all Thika Road public service vehicles use the Globe Roundabout to reach the CBD (Njaaga, 2026), but it has not been observed. Hypothesis H2, that the CBD becomes hypercongested in the evening peak, follows from the bathtub model and is equally untested.

The distinction matters for the choice of intervention. Moving a stage, marshalling boarding or enforcing against loading in live lanes would not shorten waits if the binding constraint were the number of seats arriving; conversely, changes at upstream junctions could reduce stage waits even though the stage itself is untouched. Dar es Salaam's experience, where protected lanes shortened on-road travel but a shortage of buses left passengers waiting up to 121 minutes (National Audit Office of Tanzania, n.d.), illustrates the same logic at the scale of a whole system.

RQ3: What Previous Interventions and Other Cities Suggest

Nairobi's recent interventions (Table 3) document implementation and disruption more clearly than effects on passenger journeys. Two inferences follow, both to be tested in interviews with those involved. First, changes in access, such as closing a junction, relocating traders or moving a stage, can shift costs onto passengers' walking, waiting and transfers, which vehicle-based measures miss; the reversal of the 2018 CBD matatu ban within 24 hours (Kinyanjui, 2018) suggests how quickly such costs can make a measure untenable. Second, none of the interventions published a before-and-after measurement of passenger journeys, so the city has had little opportunity to learn which mechanisms each one addressed. The missing evidence is therefore not only about outcomes but about mechanisms.

The international cases (Table 4) are most useful as design references rather than as predictions. Four lessons may transfer, each on a condition: a fixed measure-and-adjust cycle, if an institution has the authority to act on the measurements; contracting with operators rather than displacing them, if regulators can monitor performance; trials before permanence and visible use of revenue, if any charge is ever proposed; and attention to service supply, because, as Dar es Salaam shows, infrastructure alone does not shorten passengers' journeys. Nairobi's income levels, institutional powers, enforcement capacity and dependence on paratransit differ from most of the cases, so the size of any effect in Nairobi cannot be inferred from them.

RQ4: Designing and Evaluating a First Intervention

A one-route, one-stage pilot is appropriate at this stage because the dominant mechanism is unknown, the cost of a wrong citywide intervention is high and the institutions involved have not yet worked together on a measured problem. Its weakness is statistical: a single treated site is exposed to site-specific shocks, daily outcomes are serially correlated, and a comparison stage may change during the pilot or receive spillovers. The protocol addresses these risks with a pre-analysis plan, more than one comparison stage where feasible, event-study and placebo checks, inference that allows for serial correlation and a log of concurrent changes. The provisional sample sizes in Figure 19 suggest that a paired design could make a test feasible within weeks rather than months, but they will be recalculated from the baseline. Continuation depends on the gates in Table 20: an effect on passenger waiting or reliability estimated with uncertainty, and a credible cost per passenger-hour saved.

Calls for agencies to collaborate are common and rarely testable. Table 16 makes the cooperation concrete by stating, for each candidate fix, the action, who can authorise it, who must carry it out and what each partner contributes. Which row applies is decided by the diagnosis.

Table 16

Candidate Fixes by Mechanism: Who Authorises, Who Acts and What Each Partner Contributes

Mechanism found	Candidate action	Who can authorise	Who carries it out	Partner contributions
Vehicle supply (late returns)	Signal-timing or turning change on the inbound route at peak; dispatch rule that holds headways	Road authority for the junction (KeNHA or KURA); SACCO management	Signal engineers; traffic officers; SACCO dispatchers	Counts and timing data (study team); GPS traces (SACCO); officer deployment (police)
Stage operations	Boarding layout, queue marshalling or dispatch interval	County (stage); SACCO	Stage marshals; crews	Marshal time (SACCO); space and signage (County); observation (study team)
Downstream bottleneck	Box-junction enforcement or minor geometry change	Road authority; police command	Traffic officers; contractor	Enforcement schedule (police); design (road authority)
Demand above planned capacity	Added peak vehicles or staggered dispatch	SACCO; NTSA for licensing	SACCO	Fleet (SACCO); licensing (NTSA); evaluation (study team)

Note. The table is a template; the actual owner of each junction or stage is confirmed during discovery. KeNHA = Kenya National Highways Authority; KURA = Kenya Urban Roads Authority; NTSA = National Transport and Safety Authority; SACCO = savings and credit co-operative.

Attribution needs care. A common event log can improve diagnosis, but a gain from a package of marshalling, enforcement and layout changes does not show that the coordination forum caused it (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026). We therefore separate two claims. The effect of an operational change on passengers is estimated by the difference-in-differences design. The contribution of coordination is assessed through process outcomes, compared with usual practice: whether recurring problems acquire named owners, how quickly owners act and which joint decisions altered practice. Where a minimal package is inseparable, for example when marshalling cannot work without an officer present, the package is the treatment and no claim is made about its components.

Figure 22 shows the operating cycle the study proposes. It begins with measurement on a common clock, locates the origin of delay, assigns a named owner and a participating SACCO, changes one thing, evaluates against a comparison site and decides whether to keep, adjust, replicate or stop. A joint event log records who decided what and when. Only after results hold at a second site would the approach expand into a standing Corridor Cell, a public atlas and an index. Tools that rely on artificial intelligence, such as incident detection or travel-time forecasting, are introduced only when they outperform simple baselines on held-out data, and decisions remain with the responsible agencies.

Figure 22

The Proposed Measurement and Coordination Cycle

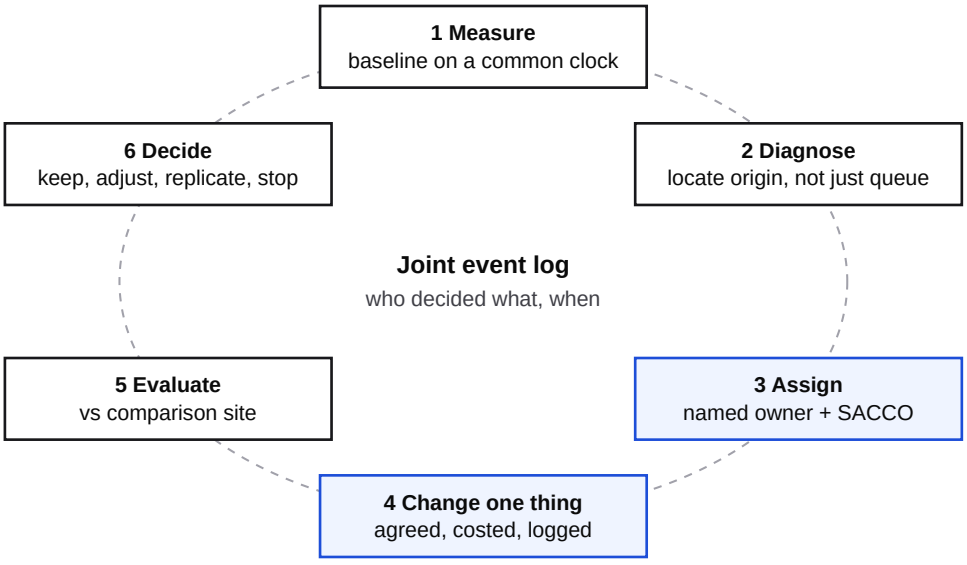

Note. Steps 3 and 4, highlighted, are where a named owner and a participating SACCO commit to act; no expansion is considered until both have been completed and evaluated.

RQ5: Institutional Coordination and the Role of AI

RQ5 asks whether Nairobi's institutions share a coordinated system and how AI could help. On the documentary evidence, the answer to the first part is no, although this inference must be tested in interviews, because coordination may occur informally without being documented. The components exist or are being built, but they are not connected to one another, to named owners or to passengers' outcomes (Table 14). The answer to the second part is that AI is most useful where that gap is widest: turning large volumes of camera, probe, mobile and crowd-sourced data into specific problems, each routed to the

institution that can fix it, with a record of what happened next. Figure 23 sets out the system this research programme aims to build: a national, AI-assisted traffic coordination system in which AI measures, detects, diagnoses, routes and evaluates, while named institutions decide and answer for the results.

Figure 23

Proposed AI-Assisted Traffic Coordination System

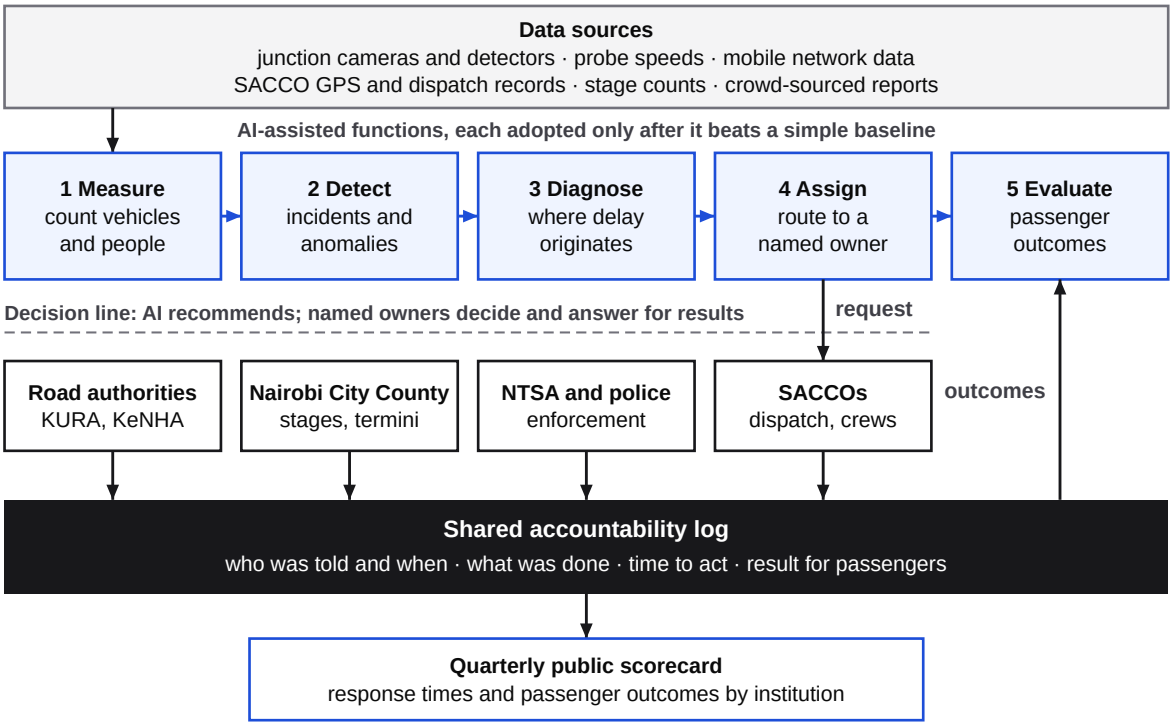

Note. AI functions sit above the decision line; institutions below it decide and act. NaMATA and the National Treasury use the scorecard for planning and funding. KeNHA = Kenya National Highways Authority; KURA = Kenya Urban Roads Authority; NTSA = National Transport and Safety Authority; SACCO = savings and credit co-operative.

Table 17 describes each function, the evidence so far, its readiness, whether it needs AI at all, the baseline it must beat and who decides. Two functions already have Nairobi precedents: KURA plans AI counting at 125 junctions (“Government to Roll Out AI-Powered Traffic Lights,” 2026), and machine learning has turned crowd-sourced reports into a crash map of the city (Milusheva et al., 2021). Diagnosis and assignment, the two functions that matter most for coordination, have none. Adaptive signal control, which changes signal timing automatically, belongs among the actions that a road authority may choose once a junction has been diagnosed; the Pittsburgh results (Smith et al., 2012) suggest it can help at junctions where the junction itself is the binding constraint.

Table 17

Functions of a Coordination System: Evidence, Readiness, Need for AI and Decision Rights

Function	What the system would do	Evidence so far	Readiness	AI or rules?	Baseline to beat	Who decides
Count vehicles	Count and classify vehicles from junction and stage cameras	Established computer-vision methods; AI counting planned at 125 junctions	Available; Nairobi accuracy untested	AI	Manual counts on sample days	Data owner (KURA, County)
Count passengers	Count passengers waiting and boarding at stages	Harder in dense crowds	Adaptation and validation needed	AI	Manual counts by crowd density	Data owner with study team
Detect incidents and anomalies	Flag crashes, blocked junctions and unusual queues from cameras, probe speeds and crowd reports	Crowd reports geolocated with machine learning in Nairobi	Adaptation needed	AI or statistics	Fixed speed-drop thresholds	Control-room operator
Diagnose origin of delay	Estimate the contribution of supply, loading, junction capacity and incidents to observed waits	No evaluated application	Research and development	Models with human review	An engineer's diagnosis from the same data	Study team with site owner
Assign and track	Route each recurring problem to its owner, set deadlines, escalate and keep an audit trail	No documented system in Nairobi	Software workflow	Rules	A manual log kept by a secretariat	Named owner; forum for escalation
Evaluate and report	Estimate effects against comparison sites; draft quarterly scorecards in English and Kiswahili	Standard statistical methods	Can be automated	Statistics; AI for drafting text only	An analyst's manual evaluation	Independent evaluator; results published

Note. A function is adopted only when it outperforms its baseline on held-out Nairobi data, and its performance (for example, counting accuracy, false-alarm rate and detection delay) is published. KURA = Kenya Urban Roads Authority.

Three design principles follow from the evidence. First, AI is not used simply because the programme is about AI, and it must earn each role: functions such as routing a problem to its owner, deadlines, escalation and audit trails are better served by transparent rule-based software, and an AI function is adopted only when it offers a measurable advantage over a simpler method: a function is adopted only when it beats a simple baseline on held-out Nairobi data, and its performance is published, because most published results for AI in traffic management come from the developers of the systems concerned. Second, accountability stays with people. The system recommends and records, but a named institution decides, and the shared log records who was told, when, what they did and what happened to passengers; response times and outcomes are reported by institution each quarter. This is the institutional piece missing from the documented systems, and it is consistent with guidance that assigns accountability for AI systems to identified organisations (National Institute of Standards and Technology [NIST], 2023). Third, the system must protect people, through the safeguards set out in Table 18. Alignment with the National AI Strategy (MICDE, 2025) would help, but the absence of a dedicated AI law (EY, 2025) makes these safeguards a design responsibility rather than a legal default.

The system should be built in the same sequence as the pilot (Table 20). The diagnosis and the first pilot need only simple measurement and a shared log kept by hand. Detection and diagnosis tools are trained and tested on the data the pilot produces; assignment and tracking become software once the coordination forum has used the manual log; and a national system linking the Traffic Management Centre, NTSA, the County, KeNHA, NaMATA and SACCOs is proposed only after results hold at a

second site. AI is not a substitute for the coordination this paper describes. Without named owners and a measured cause, AI-powered signals optimise individual junctions while the origin of delay may lie elsewhere, as the stage-supply analysis shows. With them, AI can make coordination faster, cheaper and more transparent.

The systems assessment has limits. It relies on public documents and press reports, so it may miss coordination that happens informally or in internal systems, and it cannot assess how well documented systems perform. Interviews with the agencies and operators named in Table 14 are needed to test it.

AI Governance, Data Protection, Safety and Public Accountability

A system that combines cameras, mobile network data, GPS traces, passenger counts, crowd-sourced reports, SACCO records and AI models raises governance questions that must be settled before deployment, not after. Table 18 lists the main risks and the safeguards this programme commits to.

Table 18

Governance Risks of an AI-Assisted Coordination System and Planned Safeguards

Risk area	Safeguard
Purpose and minimisation	Each data stream has a written purpose; only data needed for that purpose are collected; no data are reused for unrelated purposes without a new assessment
Lawful basis and impact assessment	Lawful basis established for each stream; data protection impact assessment completed where processing is high-risk (Data Protection Act, 2019); data-sharing agreements signed between institutions
Re-identification	Aggregated, pseudonymised mobile and GPS data are not assumed to be anonymous; re-identification risk is tested before release; small cells are suppressed
Security	Encryption in transit and at rest; role-based access; access logs reviewed; incident-response plan
Retention and ownership	Retention and deletion schedules for each stream; data ownership and access rights stated in agreements
Errors, bias and drift	Published accuracy and false-alarm rates by site and condition; checks for under-representation of passengers, pedestrians and low-income areas; periodic re-testing to detect model drift
Human oversight	AI recommends and records; named institutions decide; every automated alert can be overridden and every decision is logged
Enforcement uses	Coordination data are not used to penalise individuals; enforcement remains governed by existing law and procedures
Procurement and vendor independence	Open standards and data formats; evaluation results owned by the public partners; no exclusive vendor lock-in
Public accountability and redress	Quarterly public scorecards; a channel for operators and the public to report errors and request correction

Contribution to Theory and Practice

The paper contributes five things whose value does not depend on agreeing with its hypotheses: a passenger-centred conceptual model that treats congestion as the product of several interacting mechanisms; a new operational model, the stage seat-supply model, that links inbound delay to passenger queues; a measurement approach that weights delay by the people affected (the hotspot score) and attributes it across competing mechanisms (the diagnostic matrix); a pre-specified evaluation design for a single-site operational change; and a staged design for coordination in which AI is adopted function by function against baselines. None of these has yet been established empirically; each is offered in a form that others can test.

Engaging the Economists: What Is Borrowed, What Is Consistent and What Is Contested

The economics literature summarised in Table 2 is more consistent with this paper's framing than opposed to it, although none of its results has yet been tested in Nairobi. Three of them hold within the models and agree with the available Nairobi evidence. First, within Vickrey's model individual effort cannot beat a jam, because every arrival time costs the same (Figure 15), and Wardrop's user equilibrium (Wardrop, 1952) and Downs' triple convergence (Downs, 2004) explain why every

individual advantage, whether an earlier start, a back road or a navigation app, is copied until it disappears. Second, added capacity may be a fragile remedy: in the illustrative bottleneck, once converging demand is allowed for, a 25% capacity increase leaves the queuing loss higher than before. Third, in these models the inefficiency is structural: it arises from the way reasonable choices interact rather than from anyone's misbehaviour, which is why the working hypothesis concerns coordination rather than discipline.

We borrow two tools. The bottleneck model supplies the vocabulary of queuing cost, schedule delay and peak spreading, and a first estimate of what a time-varying instrument could achieve. The bathtub supplies a testable hypothesis about the CBD (H2) and a lesson that does not depend on prices: holding inflow below the critical density is valuable however it is achieved. Singapore's quarterly adjustment of charges against target speed bands (Land Transport Authority, 2026a) is, in effect, Vickrey's idea operated as a measure-and-adjust cycle. It is the same cycle this paper proposes (Figure 22), but with coordination rather than charging as the first instrument.

We contest the literature's direct application to Nairobi on four points where Nairobi differs from the settings in which it was developed. First, the decision maker is different. The bottleneck model assumes that each traveller chooses when to depart. For the 32.2% of trips made by matatu (JICA Expert Team, 2025), the moment a vehicle enters the peak is decided largely by SACCO dispatch and by crews waiting to fill seats, so any time-of-day instrument would have to work through operators' dispatch rules as much as through passengers' choices. Second, the use of revenue decides legitimacy. Because an optimal charge leaves each commuter's private cost unchanged until the revenue is spent (Small, 2015), a charge in Nairobi would be judged by where the money goes. The suspension of the 2018 CBD matatu ban within 24 hours (Kinyanjui, 2018) shows how quickly a measure that visibly burdens passengers can be withdrawn, and Stockholm introduced its charge through a full-scale trial (Eliasson et al., 2009). Third, the prize may be smaller than the theory suggests. Evidence from India (Akbar et al., 2023; Kreindler, 2024) indicates that where uncongested speeds are low and delay rises only moderately with density, optimal charges yield modest gains; in Nairobi, with an implied free-flow speed of about 33.5 km/h, improving junction operations and the network may matter more. Fourth, Nairobi's one priced road prices distance, not time ("Nairobi Expressway Explains ETC," 2023): it offers a way out of the jam to those who pay, rather than a smaller jam for everyone.

These arguments do not reject congestion pricing; they place it later in the sequence. A charge requires what Nairobi does not yet have: measured speeds and densities on named corridors, a credible estimate of what commuters will pay to save time and to avoid arriving late (P1, P11), an institution trusted to collect and spend the revenue, and evidence of which mechanism binds where. The diagnosis-first pilot proposed here generates the first two and tests the third on a small scale. If it shows that the CBD is hypercongested and that coordinated, non-price measures cannot hold inflow below the critical level, the case for a Vickrey-style charge, with transparent use of revenue, becomes stronger and can be made with Nairobi data rather than borrowed results.

A Framework Other Cities Can Test

The contribution intended for international use is not Nairobi's cost estimate or any single model, but a research and diagnostic framework for studying passenger-centred congestion and institutional coordination in cities with complex transport systems. Figure 24 sets out its seven phases, and Appendix C specifies the minimum dataset each city should report so that results can be compared and pooled. The framework makes no claim that Nairobi's mechanisms apply elsewhere; it asks each city to measure its own.

Figure 24

Passenger-Centred Congestion Diagnostic Framework for Replication in Other Cities

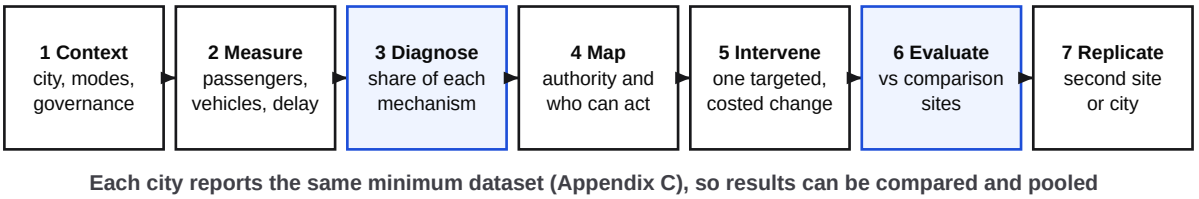

Note. Phases 3 and 6, highlighted, are those this paper specifies in most detail. Appendix C lists the reporting template.

Limitations and Mitigations

Table 19 lists the main limitations of the study, their likely consequences and how the programme will address each.

Table 19

Limitations, Consequences and Mitigations

Limitation	Likely consequence	Mitigation
No primary fieldwork	Mechanisms and hypotheses are illustrated, not observed	Diagnosis on one route and stage before any intervention (P9, P12)
Wage-based value of time	Valuation may not represent informal workers, students or care workers; largest source of variance	Nairobi stated-preference survey by mode, income and purpose (P1)
Car-based probe delay applied to all motorised modes	Matatu waiting and loading time not captured	Mode-specific GPS and stage observations; structural uncertainty stated alongside the Monte Carlo results
Exclusions: pedestrians, boda boda users, crashes, emissions, reliability, off-peak delay	Estimate is incomplete in scope; distribution across groups unknown	Distributional analysis (P8); reliability valuation (P5)
Fuel at pump prices in the headline figure	Overstates resource cost	Resource-cost estimate reported alongside (Table 11)
Network averages	Corridor variation hidden	Corridor calibration of volume–delay functions (P2)
Illustrative parameters in the stage, bottleneck and business-case models	Outputs show mechanisms, not Nairobi values	Validation protocol for the stage model; field calibration
Statistics accessed through press reports	Possible transcription errors	Source registry flags each case; replace with original documents
Narrative, English-only review without database searches	Relevant studies may be missed	Systematic review with registered protocol (P15)
Documentary assessment of coordination systems	Informal coordination may be missed	Institutional interviews (Appendix A)
Hypothesis H1 arose from discussion	May not hold	Tested in diagnosis; reported whether or not it holds

Future Research

Fifteen open problems would most improve the evidence base:

1. P1: Estimate a value of travel time for Nairobi across modes and income groups through a stated-preference survey.
2. P2: Calibrate the volume–delay function for named corridors using probe speeds and counts.
3. P3: Define an appropriate “uncongested” benchmark and test how much the choice changes the cost estimate.
4. P4: Model mixed traffic with a fast-growing share of motorcycles, distinguishing filtering from blocking.
5. P5: Estimate a buffer-time index and the value of reliability for Nairobi commuters.

6. P6: Measure saturation flow at junctions with and without loading in lanes, blocked junctions and manual signal override.
7. P7: Test whether the near one-to-one relationship between road capacity and driving holds in a city where most people do not own cars.
8. P8: Analyse how the time cost is distributed by income, mode and home location.
9. P9: Calibrate the stage seat-supply model with synchronised stage and inbound-trip observations, including bunching.
10. P10: Develop and validate process indicators for the contribution of institutional coordination.
11. P11: Estimate the costs that Nairobi commuters and matatu operators attach to arriving early and late, which any time-of-day instrument would require.
12. P12: Test whether the CBD becomes hypercongested in the evening peak by estimating a network-level flow–density relationship from probe speeds and counts (hypothesis H2).
13. P13: Test whether AI-based counting, incident detection and delay attribution outperform simple baselines on Nairobi data, and audit them for bias, including the under-representation of passengers and pedestrians in probe and crowd-sourced data.
14. P14: Evaluate whether a shared accountability log, with response times reported by institution, changes how quickly institutions act and whether passenger journeys improve, as a first test of an AI-assisted coordination system for Kenya.
15. P15: Conduct a systematic review of passenger-centred congestion diagnosis and paratransit operations using bibliographic databases, with a registered protocol, to update and test the narrative review in this paper.

The programme will report this work as a series of four papers, each with its own data: the cost of congestion with mode-specific measurements (P1–P3, P5, P8); where passenger queues originate, with field validation of the stage model (P4, P6, P9, P12); institutional coordination and accountability (P10, P14); and AI-assisted coordination tested against simpler methods (P13). Each will be preregistered where it involves an evaluation, and its data and code will be released with it.

Conclusion and Recommendations

This study examined weekday peak congestion in the Nairobi Metropolitan Area through a passenger-centred framework that combines secondary-data analysis, congestion economics, transport operations, institutional assessment and a proposed field evaluation protocol.

Under the central assumptions of the economic model, weekday peak congestion is estimated to impose an annual cost of about KSh 46.3 billion, with a Monte Carlo 90% interval of KSh 33.6–70.8 billion. The range reflects uncertainty in the value of travel time, peak delay, peak exposure and trip length, and the model's structure adds further uncertainty that the interval does not capture. The estimate should therefore be read as an indication of the scale of the burden rather than a precise measurement of its total social cost. The available travel-demand data show that walking and matatu travel account for a large share of daily trips, but the model does not estimate pedestrian delay or the full costs borne by motorcycle users, so the distribution of the burden across modes and income groups remains an empirical question.

The analysis also develops an illustrative stage seat-supply model to show how delays on vehicles' return journeys may contribute to passenger queues at central stages. The mechanism is plausible and provides a basis for field investigation, but its contribution to actual passenger waiting in Nairobi remains to be measured. The review of interventions, of congestion economics and of institutional systems suggests that infrastructure, monitoring technology and enforcement may be insufficient when operational problems cross institutional boundaries, and that Nairobi is acquiring the means to observe its traffic faster than the means to turn observation into accountable action. The extent to which institutional coordination independently improves passenger outcomes, however, remains an empirical question.

The study recommends a staged research and policy programme, set out below by time horizon.

Immediate Actions (0 to 6 Months)

1. **Fund a short diagnosis first.** Government and development partners should fund an eight- to ten-week diagnosis on one route and one stage, using synchronised passenger, vehicle and junction observations, with the responsible road authority and a willing SACCO involved from the start (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026). Ethics approval and a research licence should be obtained before fieldwork.
2. **Make the decision to proceed depend on an owner and a measured cause.** A pilot should proceed only when there is a baseline, a named operating owner, an agreed and costed change, comparison stages and a registered pre-analysis plan. Survey responses and crowd-sourced reports should inform discovery but not decide whether to proceed.
3. **Measure existing investments.** KURA and its partners should take independent before-and-after passenger measurements at intelligent transport system junctions before they are switched on, and the relationship between the separately financed traffic-system loans should be documented publicly.

Medium-Term Actions (6 to 18 Months)

4. **Test one operational change and measure passengers.** The first test should change one thing where feasible, with passenger waiting time and end-to-end journey reliability as primary outcomes, and should track people moved per hour, pedestrian access and spillover.
5. **Start a shared accountability log.** The agencies and operators involved in the pilot should keep a shared, time-stamped log that records each recurring problem, its named owner, the action taken and the result for passengers, initially by hand and later in rule-based software, and should publish response times and outcomes.
6. **Expand only after results hold at a second site.** Results and their limits should be published; if gains persist and the cost per passenger-hour saved is credible, the test should be repeated at a second site before any wider rollout.

Long-Term Actions (18 Months and Beyond)

7. **Build a shared, AI-assisted traffic coordination system in stages.** The ministry responsible for transport, with KURA, KeNHA, NTSA, the National Police Service, Nairobi City County, NaMATA and SACCO representatives, should extend the accountability log into a shared coordination system. AI functions for counting, incident detection and diagnosis should be added one at a time, each only after it outperforms a simple baseline on Nairobi data and its governance safeguards are in place, with decisions remaining with the named institutions. Once results hold at a second site, the Traffic Management Centre at City Cabanas is a natural host for a national system.
8. **Treat congestion pricing as a later option, built on Nairobi evidence.** A time-varying charge should be considered only once corridor speeds and densities, commuters' valuation of time and lateness, and a trusted use of revenue have been established. If the CBD proves hypercongested and coordinated measures cannot keep inflow below the critical density, the case for a charge should then be made with Nairobi data (Arnott, 2015; Vickrey, 1969).
9. **Replicate the framework beyond Nairobi.** Universities, city agencies and development partners in other cities with paratransit-based transport should apply the diagnostic framework and report the minimum dataset in Appendix C, so that what transfers, and what does not, can be learned across cities.

Figure 25 and Table 20 set out the sequence that follows from these recommendations.

Figure 25

Programme Sequence, December 2026 to June 2028

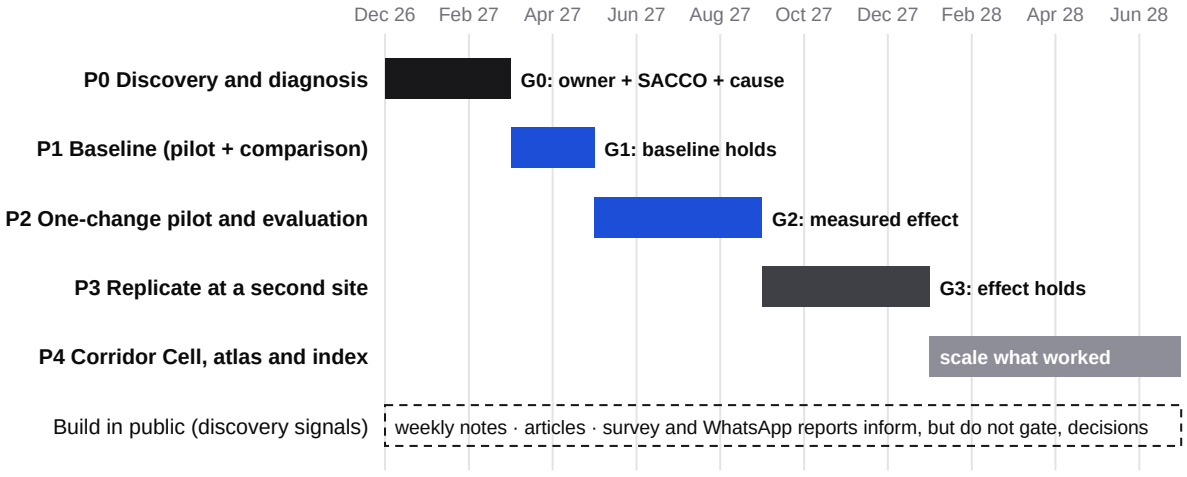

Note. Dates are targets; each phase begins only when the preceding gate is met. G = gate.

Table 20

Phases and Decision Gates

Phase	Main activities	Gate to proceed (all required)
P0 Discovery and diagnosis (8–10 weeks)	Interviews; synchronised stage and return-trip observation on several evenings; review of previous interventions; data agreements	Named site owner; SACCO willing to work through an operational change; measured cause with each mechanism's contribution estimated; candidate comparison stage
P1 Baseline	Matched weekdays at pilot and comparison stages, with the length set from baseline variance and serial correlation (provisionally at least 16); estimate variability; register the pre-analysis plan	Comparison stages behave similarly before the change; agreed and costed single change; pre-analysis plan registered; ethics approval and research licence in place
P2 One-change pilot	Implement change; matched weekdays as set in the pre-analysis plan; log concurrent changes and decisions	Effect on passenger waiting and reliability estimated with uncertainty; results published
P3 Replication	Repeat at a second site with its own comparison	Effect holds; cost per passenger-hour saved is credible
P4 Corridor Cell, atlas and index	Standing coordination forum; shared accountability log in software; AI functions adopted after beating baselines; public atlas; quarterly open data and scorecards	Institutional commitment and sustainable funding

Note. Sample sizes follow Equations 22 and 23 with the assumptions in the Method section and are provisional until the baseline is observed.

The principal contribution of this paper is a preliminary framework for investigating congestion through passenger outcomes, operational mechanisms and institutional accountability. Its wider applicability will depend on empirical validation, transparent publication of data and methods, and replication across routes and cities. Future research should prioritise mode-specific travel-time valuation, calibrated traffic-flow models, validation of the stage seat-supply model, stronger comparative evaluation designs and measurement of the contribution of institutional coordination. The framework is intended to support evidence-based transport policy, open research and accountable institutional learning, and its effectiveness will be judged not by the sophistication of its models or technology but by whether it produces measurable and equitable improvements in passenger journeys.

An Invitation to Contribute

Traffic Jam is an open-source project. It is not sponsored or funded by any institution, and its evidence will improve fastest if others test, correct and extend it. Contributions of every kind are welcome: counts and observations from stages and junctions; replications of the cost, bottleneck and stage-supply models under different assumptions; work on any of the fifteen open problems listed in the Discussion, especially those that bring an AI-assisted coordination system for Kenya closer; applications of the replication framework in other cities; corrections to sources; and accounts from passengers, crews, marshals and officials of how the city's peaks actually work. The paper, the model code and its outputs are archived on Zenodo and available at nallalabs.xyz or on request from hello@nallalabs.xyz, which also receives general inquiries and contributions. Correspondence about the paper should be addressed to allanrobinson@nallalabs.xyz.

References

Africa Transport Policy Program. (2018). *Policies for sustainable accessibility and mobility in cities of Kenya*. World Bank. https://ssatp.org/sites/default/files/publication/SSATP_UTM_FinalReport_KENYA.pdf

African Development Bank. (2017, November 9). *Thika Highway continues to impact Kenyan lives five years on* [Press release]. AllAfrica. https://allafrica.com/stories/201711091310.html

Akbar, P. A., Couture, V., Duranton, G., & Storeygard, A. (2023). Mobility and congestion in urban India. *American Economic Review, 113*(4), 1083–1111. https://doi.org/10.1257/aer.20181662

Arnott, R. (2013). A bathtub model of downtown traffic congestion. *Journal of Urban Economics, 76*, 110–121.

Arnott, R. (2015). A bathtub model of downtown traffic congestion. *ACCESS Magazine, 1*(46). https://escholarship.org/uc/item/24h06883

Arnott, R., de Palma, A., & Lindsey, R. (1993). A structural model of peak-period congestion: A traffic bottleneck with elastic demand. *American Economic Review, 83*(1), 161–179.

Asian Development Bank. (n.d.). *Fixing a city's broken public bus service* [Case study]. Development Asia. https://development.asia/case-study/fixing-citys-broken-public-bus-service

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics, 119*(1), 249–275. https://doi.org/10.1162/003355304772839588

Bett, H. (2024, March 14). Road traffic fatalities in Kenya, 2023. *The Star*. https://www.the-star.co.ke/news/2024-03-14-road-traffic-fatalities-in-kenya-2023

Bikes boom drives spike in vehicle registrations. (2026, April 30). *The Star*. https://www.the-star.co.ke/news/2026-04-30-bikes-boom-drives-spike-in-vehicle-registrations

Boda boda riders with financed motorcycles earn more, new report. (2025). *Capital FM*. https://capitalfm.africa/boda-boda-riders-with-financed-motorcycles-earn-more-new-report/

Braess, D. (1968). Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung, 12*, 258–268. https://doi.org/10.1007/BF01918335

BRTData. (n.d.). *Bogotá*. Global BRTData. https://brtdata.org/location/latin_america/colombia/bogota

Bureau of Public Roads. (1964). *Traffic assignment manual for application with a large, high speed computer*. U.S. Department of Commerce. https://catalog.hathitrust.org/Record/000968330

Cervero, R., & Golub, A. (2007). Informal transport: A global perspective. *Transport Policy, 14*(6), 445–457. https://doi.org/10.1016/j.tranpol.2007.04.011

Clearing congestion. (2016). *The Technologist*. https://technologist.eu/?p=1032

Daganzo, C. F. (2009). A headway-based approach to eliminate bus bunching: Systematic analysis and comparisons. *Transportation Research Part B: Methodological, 43*(10), 913–921.

Data Protection Act, No. 24 of 2019, Laws of Kenya (2019). https://www.kenyalaw.org

Downs, A. (2004). *Still stuck in traffic: Coping with peak-hour traffic congestion* (Rev. ed.). Brookings Institution Press.

Drèze, J. H. (1998). William S. Vickrey, June 21, 1914–October 11, 1996. In *Biographical memoirs* (Vol. 75). National Academies Press. https://doi.org/10.17226/9649

Duranton, G., & Turner, M. A. (2011). The fundamental law of road congestion: Evidence from US cities. *American Economic Review, 101*(6), 2616–2652. https://doi.org/10.1257/aer.101.6.2616

Effect of motor vehicle congestion on the economic performance of Kenya: A case of Nairobi City County. (n.d.). *Strategic Journal of Business & Change Management*. https://www.strategicjournals.com/index.php/journal/article/download/96/96

Eliasson, J., Hultkrantz, L., Nerhagen, L., & Smidfelt Rosqvist, L. (2009). The Stockholm congestion-charging trial 2006: Overview of effects. *Transportation Research Part A: Policy and Practice, 43*(3), 240–250. https://doi.org/10.1016/j.tra.2008.09.007

Etyang, P. (2026, March 26). Traffic restored at Globe Roundabout after repairs. *The Star*. https://www.the-star.co.ke/news/2026-03-26-traffic-restored-at-globe-roundabout-after-repairs

Explained: Taxes and subsidy in April–July fuel pricing cycle. (2026, April 15). *Capital FM*. https://capitalfm.africa/news/2026/04/explained-taxes-and-subsidy-in-april-july-fuel-pricing-cycle/

EY. (2025, October 30). *Kenya releases the National AI Strategy 2025–2030*. https://taxnews.ey.com/news/2025-2193-kenya-releases-the-national-ai-strategy-2025-2030

Eyes on Cabinet as France says ready to fund Nairobi BRT. (2026). *Nation*. https://beta.nation.co.ke/kenya/business/eyes-on-cabinet-as-france-says-ready-to-fund-nairobi-brt-5191002

Flone Initiative. (2018). *Report on gender equity assessment of Nairobi's public minibus transport services*. UN-Habitat. https://unhabitat.org/report-on-gender-equity-assessment-of-nairobis-public-minibus-transport-services-0

Fosgerau, M., & Small, K. A. (2013). Hypercongestion in downtown metropolis. *Journal of Urban Economics, 76*, 122–134.

Google. (2023). *Project Green Light's work to reduce urban emissions using AI*. https://blog.google/outreach-initiatives/sustainability/google-ai-reduce-greenhouse-emissions-project-greenlight/

Government to roll out AI-powered traffic lights at 125 Nairobi junctions. (2026). *Eastleigh Voice*. https://eastleighvoice.co.ke/city-affairs/348540/government-to-roll-out-ai-powered-traffic-lights-at-125-nairobi-junctions

Greenshields, B. D. (1935). A study of traffic capacity. *Highway Research Board Proceedings, 14*, 448–477. https://trid.trb.org/view/120649

How boda boda riders make Sh1bn every day. (2022). *Business Daily*. https://www.businessdailyafrica.com/bd/economy/how-boda-boda-riders-make-sh1bn-everyday-3743962

Institute for Transportation and Development Policy. (2021). *Jakarta is what resiliency looks like*. https://framework.itdp.org/?p=17710

ITDP Africa. (2022). *Optimising BRT operations in Dar es Salaam*. https://africa.itdp.org/?p=17983

JICA Expert Team. (2025). *The project for capacity building for bus operation policy and management in Nairobi Metropolitan Area: Technical report on traffic demand forecast and public transport strategy*. Japan International Cooperation Agency. https://openjicareport.jica.go.jp/pdf/12391272_03.pdf

Kenya Institute for Public Policy Research and Analysis. (n.d.). *Evaluating the impact of road traffic congestion mitigation measures in Nairobi Metropolitan Region* (Discussion paper). https://repository.kippra.or.ke/server/api/core/bitstreams/e50b4af4-14cb-4c48-a326-4048b65e5d9b/content

Kenya signs Ksh24 billion loan with China EximBank to fund Nairobi traffic management plan. (2025, November 26). *Kenyans.co.ke*. https://www.kenyans.co.ke/news/118363-kenya-signs-ksh-24-billion-loan-china-eximbank-fund-nairobi-traffic-management-plan

Kenya Urban Roads Authority. (2026). *Specific procurement notice: Nairobi intelligent transport system (ITS) establishment and junctions improvement project (Phase II)*. https://gaa.go.ke/sites/default/files/2026-06/KURA%20Specific%20Procurement%20Notice%20Establishment%20And%20Junctions%20Improvement%20Proje

ct%20%5BPHASE%20II%29.pdf

Kenyan shilling nears Ksh130 against U.S. dollar as pound hits Ksh175. (2026, November). *Kenyans.co.ke*. https://www.kenyans.co.ke/news/129201-kenyan-shilling-nears-ksh130-against-us-dollar-pound-hits-ksh175

Kinyanjui, M. (2018, December 13). Roller coaster ride for matatus, commuters in year of troubles. *The Star*. https://www.the-star.co.ke/news/2018-12-13-roller-coaster-ride-for-matatus-commuters-in-year-of-troubles

Kinyanjui, M. (2021, April 27). Test run at Green Park Terminus confused but okay. *The Star*. https://www.the-star.co.ke/news/realtime/2021-04-27-test-run-at-green-park-terminus-confused-but-okay

Kipkemoi, F. (2026, June 16). KURA invites bids for Sh10.8bn Nairobi smart traffic project. *The Star*. https://www.the-star.co.ke/news/2026-06-16-kura-invites-bids-for-sh108bn-smart-traffic-project

Knight, F. H. (1924). Some fallacies in the interpretation of social cost. *The Quarterly Journal of Economics, 38*(4), 582–606. https://doi.org/10.2307/1884592

Kreindler, G. (2024). Peak-hour road congestion pricing: Experimental evidence and equilibrium implications. *Econometrica, 92*(4), 1233–1268. https://doi.org/10.3982/ECTA18422

KURA initiates plan to monitor and manage Nairobi traffic using smart cameras. (2025, October 27). *People Daily*. https://peopledaily.digital/news/kura-initiates-plan-to-monitor-and-manage-nairobi-traffic-using-smart-cameras

KURA signs deal for building Nairobi's intelligent transport system and improve 25 junctions. (2024, November 27). *People Daily*. https://peopledaily.digital/news/kura-signs-deal-for-building-nairobis-intelligent-transport-system-and-improve-25-junctions

Land Transport Authority. (2026a, March). *Revised ERP rates from 23 March 2026*. https://www.lta.gov.sg/content/ltagov/en/newsroom/2026/3/news-releases/revised-erp-rates-23mar26.html

Land Transport Authority. (2026b, November 16). *Revised ERP rates from 23 November 2026*. https://www.lta.gov.sg/content/ltagov/en/newsroom/2026/11/news-releases/revised-erp-rates-from-23-nov-26.html

Lassen, D. (2026, January 6). New York, MTA say first year of congestion pricing is a success. *Trains*. https://www.trains.com/pro/passenger/commuter-regional/new-york-mta-say-first-year-of-congestion-pricing-is-a-success/

McCormick, D., Mitullah, W., Chitere, P., Orero, R., & Ommeh, M. (2013). Paratransit business strategies: A bird's-eye view of matatus in Nairobi. *Journal of Public Transportation, 16*(2), 135–152. https://digitalcommons.usf.edu/jpt/vol16/iss2/7

Milusheva, S., Marty, R., Bedoya, G., Williams, S., Resor, E., & Legovini, A. (2021). Applying machine learning and geolocation techniques to social media data (Twitter) to develop a resource for urban planning. *PLOS ONE, 16*(2), Article e0244317. https://doi.org/10.1371/journal.pone.0244317

Ministry of Information, Communications and the Digital Economy. (2025, March 27). *Kenya's Artificial Intelligence (AI) Strategy 2025–2030 launched at the KICC, Nairobi*. https://ict.go.ke/kenyas-artificial-intelligence-ai-strategy-2025-2030-launched-kicc-nairobi

Muriuki, B. (2024, September 12). Nairobi County defends relocation of traders from Wakulima market to Kangundo road. *Citizen Digital*. https://citizen.digital/article/nairobi-county-defends-relocation-of-traders-from-wakulima-market-to-kangundo-road-cites-health-waste-concerns-n349397

Mwita, M. (2025, March 20). Samsung to build Nairobi's Sh7.9bn intelligent traffic system. *The Star*. https://www.the-star.co.ke/business/2025-03-20-samsung-to-build-nairobis-sh79bn-intelligent-traffic-system

Nairobi City County. (2024, November 30). *Finding sustainable solutions to the traffic challenges facing Nairobi central business district*. https://nairobi.go.ke/finding-sustainable-solutions-to-the-traffic-challenges-facing-nairobi-central-business-district

Nairobi City County. (2026, April 23). *Nairobi City County moves to streamline boda boda sector with new regulations*. https://nairobi.go.ke/nairobi-city-county-moves-to-streamline-boda-boda-sector-with-new-regulations

Nairobi Expressway explains ETC Ksh3,000 initial charge & its link to China. (2023, March 16). *Kenyans.co.ke*. https://www.kenyans.co.ke/news/86961-nairobi-expressway-explains-etc-ksh3000-initial-charge-its-link-china

Nairobi Metropolitan Area Transport Authority. (n.d.). *Bus rapid transit corridors*. https://namata.go.ke/brt_corridor

Nairobi traffic a health, economic burden to city residents. (2021, August 16). *The Star*. https://www.the-star.co.ke/business/kenya/2021-08-16-nairobi-traffic-a-health-economic-burden-to-city-residents/

Nairobi traffic thesis: Independent advisory review [Unpublished manuscript]. (2026, November).

National Audit Office of Tanzania. (n.d.). *Performance audit report on the management of urban transport services in major cities*. https://nao.go.tz/reports/view/management-of-urban-transport-services-in-major-cities

National Commission for Science, Technology and Innovation. (n.d.). *Frequently asked questions*. Retrieved November 23, 2026, from https://www.nacosti.go.ke/faqs/

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1

Newell, G. F., & Potts, R. B. (1964). Maintaining a bus schedule. In *Proceedings of the 2nd Australian Road Research Board Conference* (Vol. 2, Part 1, pp. 388–393). Australian Road Research Board.

Njaaga, D. (2026, March 25). KeNHA reopens Globe Roundabout after 13-day closure. *The Standard*. https://standardmedia.co.ke/nairobi/article/2001543844/kenha-reopens-globe-roundabout-after-13-day-closure

NTSA to install 1,000 speed cameras, roll out instant digital fines. (2026, February 25). *The Star*. https://www.the-star.co.ke/news/2026-02-25-ntsa-to-install-1000-speed-cameras-roll-out-instant-digital-fines

Office of the Governor of New York. (2026, January 5). *Less traffic, better transit: On its first anniversary, Governor Hochul celebrates transformational success of congestion pricing* [Press release]. https://www.governor.ny.gov/news/less-traffic-better-transit-its-first-anniversary-governor-hochul-celebrates-transformational

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press. https://doi.org/10.1017/CBO9780511807763

Pigou, A. C. (1920). *The economics of welfare*. Macmillan. https://www.econlib.org/library/NPDBooks/Pigou/pgEW.html

Reprieve for Marikiti traders as court stops County from relocating them to Kangundo. (2024, September 13). *Nation*. https://nation.africa/kenya/counties/nairobi/reprieve-for-marikiti-traders-as-court-stops-nairobi-county--4760722

Resources for the Future. (n.d.). *The experience of area licenses and cordon tolling to date*. https://resources.org/common-resources/the-experience-of-area-licenses-and-cordon-tolling-to-date

Road accident-related deaths rise by 3.4pc in 2025. (2025, December 26). *The Standard*. https://standardmedia.co.ke/national/article/2001537222/road-accident-related-deaths-rise-by-34pc-in-2025

Roughgarden, T., & Tardos, É. (2002). How bad is selfish routing? *Journal of the ACM, 49*(2), 236–259. https://doi.org/10.1145/506147.506153

Royal Swedish Academy of Sciences. (1996, October 8). *The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 1996* [Press release]. NobelPrize.org. https://www.nobelprize.org/prizes/economic-sciences/1996/press-release/

Ruto unveils Ksh7.9B smart traffic system to tackle Nairobi congestion crisis. (2026, June 29). *People Daily*. https://peopledaily.digital/news/ruto-unveils-ksh7-9b-smart-traffic-system-to-tackle-nairobi-congestion-crisis

Salon, D., & Gulyani, S. (2019). Commuting in urban Kenya: Unpacking travel demand in large and small Kenyan cities. *Sustainability, 11*(14), Article 3823. https://doi.org/10.3390/su11143823

Small, K. A. (1982). The scheduling of consumer activities: Work trips. *American Economic Review, 72*(3), 467–479.

Small, K. A. (2015). The bottleneck model: An assessment and interpretation. *Economics of Transportation, 4*(1–2), 110–117. https://doi.org/10.1016/j.ecotra.2015.01.001

Smith, S. F., Barlow, G. J., Xie, X.-F., & Rubinstein, Z. B. (2012). *Real-time adaptive traffic signal control for urban road networks: The East Liberty pilot test* (Report No. CMU-RI-TR-12-20). Robotics Institute, Carnegie Mellon University.

Sonko bows to pressure, suspends matatu CBD ban. (2018, December 4). *The Star*. https://www.the-star.co.ke/news/2018/12/04/sonko-bows-to-pressure-suspends-matatu-cbd-ban_c1860423

Team Europe to build Sh47bn green BRT in Nairobi. (2023). *Construction Kenya*. https://www.constructionkenya.com/11264/team-europe-brt-kenya/

Test run by NMS at Green Park Terminus signals actualisation of city matatu ban. (2021, April 28). *People Daily*. https://peopledaily.digital/news/test-run-by-nms-at-green-park-terminus-signals-actualisation-of-city-matatu-ban/amp

The National Treasury. (2025, November 26). *Signing of a USD 185 million concessional loan agreement to finance the Intelligent Transport System (ITS) project*. https://www.treasury.go.ke/node/1343

TomTom. (2026). *TomTom Traffic Index 2025: Nairobi*. https://www.tomtom.com/traffic-index/nairobi-traffic/

Traffic jams cost Nairobi Sh50m a day. (2012, May 11). *Capital FM*. https://www.capitalfm.co.ke/business/2012/05/traffic-jams-cost-nairobi-sh50m-a-day/

UN Women, Kenyatta University, & Stockholm Environment Institute. (2021). *Women and public transport in Nairobi*. https://data.unwomen.org/sites/default/files/documents/Publications/Nairobi_public-transport_report.pdf

UN-Habitat. (n.d.). *Promoting sustainable transport solutions for East African cities (SUSTRAN)*. https://unhabitat.org/programme/promoting-sustainable-transport-solutions-for-east-african-cities-sustran

Vickrey, W. S. (1963). Pricing in urban and suburban transport. *American Economic Review, 53*(2), 452–465.

Vickrey, W. S. (1969). Congestion theory and transport investment. *American Economic Review, 59*(2), 251–260.

Vickrey, W. S. (1994). Statement to the Joint Committee on Washington, DC, Metropolitan Problems (with a foreword by Richard Arnott and Marvin Kraus). *Journal of Urban Economics, 36*(1), 42–65. (Original work published 1959)

Walters, A. A. (1961). The theory and measurement of private and social cost of highway congestion. *Econometrica, 29*(4), 676–699. https://doi.org/10.2307/1911814

Wardman, M. (2004). Public transport values of time. *Transport Policy, 11*(4), 363–377. https://doi.org/10.1016/j.tranpol.2004.05.001

Wardrop, J. G. (1952). Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers, 1*(3), 325–362. https://doi.org/10.1680/ipeds.1952.11259

Webster, F. V. (1958). *Traffic signal settings* (Road Research Technical Paper No. 39). Her Majesty's Stationery Office.

Williams, S., White, A., Waiganjo, P., Orwa, D., & Klopp, J. (2015). The digital matatu project: Using cell phones to create an open source data for Nairobi's semi-formal bus system. *Journal of Transport Geography, 49*, 39–51. https://doi.org/10.1016/j.jtrangeo.2015.10.005

Wolfenden, K.-L. (2026, June 8). Safaricom and KNBS partner to harness mobile data for national planning. *TechAfrica News*. https://techafricanews.com/2026/06/08/safaricom-and-knbs-partner-to-harness-mobile-data-for-national-planning/

World Health Organization. (2023). *Global status report on road safety 2023: Kenya country profile*. https://cdn.who.int/media/docs/default-source/country-profiles/road-safety/road-safety-2023-ken.pdf

Appendix A

Draft Field Instruments and Fieldwork Plan

The instruments below are drafts. They will be translated into Kiswahili, piloted on two evenings, revised and submitted for ethics approval before use. Table 21 sets out the proposed sample; Instruments A to D follow in Tables 22, 23 and 24; and Table 25 sets out quality-control, analysis and safety procedures.

Table 21

Proposed Sampling and Fieldwork Plan

Instrument	Participants or sites	Proposed sample	Selection and duration
A: Passenger questionnaire	Passengers queuing at the pilot and comparison stages	About 400 per stage across the baseline	Systematic selection of every nth arriving passenger; about 8 minutes; adults only
B: Stage and matatu operations observation	Pilot and comparison stages	Every vehicle and 15-minute interval on each study evening	Two observers per stage on a common clock; evening peak
C: Junction observation	Candidate inbound junctions	Each study evening	Fixed observation points; video only where permitted and with faces and plates excluded
D: Interviews	Drivers, conductors, marshals, dispatchers, passengers, nearby workers, officials	About 30 to 40 interviews, until no new themes emerge	Purposive selection across roles; 30 to 60 minutes; recorded only with consent

Note. Sample sizes are planning figures and will be revised after the pilot evenings. Participants must be 18 or older and able to give informed consent; people who decline are not approached again.

Table 22

Instrument A: Passenger Questionnaire Domains

Domain	Items
Trip	Origin and destination area (not address); purpose; usual departure time
Time components	Access time to the stage; time waited today; expected boarding time; in-vehicle and transfer time on a usual day
Reliability	How often the trip takes much longer than usual; extra time allowed to arrive on time
Cost and choices	Fare paid; whether the passenger has changed, delayed or abandoned the trip because of waiting; alternative modes considered
Access and safety	Walking distance; perceived safety while waiting and travelling; disability or accessibility needs; experiences of harassment, asked privately and only with consent
Respondent	Age band; gender; occupation type; income band (optional)

Note. No names, telephone numbers or exact addresses are collected.

Table 23

Instruments B and C: Stage, Matatu Operations and Junction Observation (Common Clock, 15-Minute Intervals)

Location	Fields recorded
Stage (Instrument B)	Interval start; passengers arriving; passengers waiting at interval end; vehicles arriving; seats available; for each vehicle: sequence number, seat capacity, occupancy on departure, arrival, start of loading, departure, time clearing the CBD, queue position and dispatch practice (timed or fill-and-go); operational interruptions
Junction (Instrument C)	Traffic volume by direction and vehicle type; turning movements; signal cycle and green time or manual control; queue length and discharge; blocked-junction occurrences; loading and unloading in live lanes; pedestrian crossings; incidents; enforcement activity; passage times of the study route's vehicles
Context log	Rain; crashes and breakdowns; roadworks; enforcement presence; any change in traffic control or stage rules

Note. Vehicle identifiers are hashed or replaced with sequence numbers; number plates are not stored.

Table 24

Instrument D: Stakeholder and Institutional Interview Guide

Who to speak with	What to learn
Drivers and conductors	Where time is lost on a full round; what determines when they leave, where they pick up and how long they wait; what changed after previous interventions
Stage marshals and SACCO dispatchers	How many vehicles and seats are expected each peak; how vehicles are dispatched; when headways break down; what the SACCO can change
Passengers and people working nearby	When waits are longest; what extra walking, transfer, fare or safety concern an alternative stage would create; where pedestrians and loading compete for space
Road authority, County and traffic officers	Which institution controls the junction and the stage; what signal plans, restrictions, past trials or roadworks affect the route; who can authorise a pilot change
Institutions (KURA, KeNHA, County, NTSA, police, NaMATA)	Mandate and authority; data and systems held; history of interventions; response times; dependencies on other agencies; decision bottlenecks; budget constraints; accountability mechanisms

Note. Adapted from the independent advisory review (*Nairobi Traffic Thesis: Independent Advisory Review*, 2026). Interviews require informed consent; no personal identifiers are recorded in analysis files.

Table 25

Quality Control, Analysis and Field Safety Procedures

Area	Procedure
Training	Two days of observer and interviewer training, including practice counts against video and a written protocol
Inter-observer agreement	Ten percent of intervals double-counted; agreement reported as the intraclass correlation for counts and Cohen's kappa for categorical codes; retraining below agreed thresholds
Translation	English and Kiswahili versions, with back-translation of the questionnaire and consent forms
Recording and transcription	Audio recorded only with consent; transcribed and de-identified; recordings deleted after verification
Qualitative analysis	Thematic analysis with a codebook developed from the first interviews; two coders on a subset, with disagreements resolved by discussion
Missing data and unusual days	Missing intervals flagged, not imputed, in primary analysis; days with major incidents, strikes or extreme weather analysed separately and reported
Data management	Encrypted devices; daily upload to access-controlled storage; field sheets checked each evening
Field safety	Observers work in pairs, away from live lanes, with a supervisor on call; observation stops if conditions become unsafe

Appendix B

Reproducibility and Supplementary Files

The models are implemented in a single Python file, *trafficjam_cost_model.py*, which uses only the standard library and writes all results to a JSON file. It contains the cost model (Equations 1–5), the uncertainty analysis, including the Monte Carlo simulation with a fixed random seed, the top-down cross-check (Equation 6), historic estimate conversions, the engineering relationships (Equations 7–11), the economists' bottleneck and bathtub models (Equations 12–15), the stage seat-supply model (Equations 16–17), the sample-size calculations (Equations 21–23), the pilot business case (Equations 24–25) and the registry of announced commitments. Every input is labelled in the code as observed, derived or assumed. The file is released under the MIT License and, with its outputs, is archived on Zenodo and available at nallalabs.xyz or on request from hello@nallalabs.xyz.

Four supplementary files accompany the paper: S1, a source-verification registry listing every reference with its tier, type, link, the claims it supports, verification status, limitations and the date checked; S2, a data dictionary of every model input with its value, unit, range, distribution, source and type; S3, a log of the literature search themes and example search strings; and S4, the city replication template described in Appendix C. Contributions are welcome: readers can change any assumption, rerun the file and send corrections, extensions or field data to hello@nallalabs.xyz.

Appendix C

Replication Framework and City Reporting Template

The framework in Figure 24 is intended for any city where congestion involves paratransit or mixed traffic. Each application proceeds through seven phases: describe the context; measure passenger, vehicle and delay outcomes; diagnose the contribution of each mechanism; map who has authority and who can act; make one targeted, costed change; evaluate it against comparison sites; and replicate at a second site or city. To allow results to be compared and pooled, each application should report the minimum dataset in Table 26 (Supplementary File S4).

Table 26

Minimum Reporting Dataset for Each City Application

Group	Fields
Setting	Country; city; population; urban density; corridor or route studied
Travel and transport	Modal shares; informal or paratransit share; fleet and dispatch model on the route; road network characteristics
Governance	Institutions with authority over the junctions, stages, operators and enforcement involved; existing coordination forums
Data	Data sources, periods and access conditions; source tiers
Outcomes	Passenger waiting time; journey time and buffer-time index; people moved per hour; institutional response times
Diagnosis	Mechanisms assessed and estimated contribution of each, with uncertainty
Intervention and evaluation	Change made; owner; cost; comparison sites; pre-analysis plan reference; effect estimates with confidence intervals
Limitations and transfer	Limitations; what was expected to transfer from other cities and what did

Appendix D

Glossary

Table 27

Terms Used in This Paper

Term	Meaning
Accountability log	A shared, time-stamped record of each recurring problem, its named owner, the action taken and the result for passengers
Adaptive signal control	Signal timing that changes automatically in response to measured traffic
Bathtub model	Vickrey's picture of downtown traffic as water flowing into and draining out of a tub; outflow falls once density passes a critical level
Boda boda	Motorcycle taxi
Bottleneck model	Vickrey's model in which commuters choose when to travel, trading queuing time against arriving early or late
BRT	Bus rapid transit: high-capacity buses in protected lanes
Buffer-time index	The extra time, as a share of the mean, needed to arrive on time on 95% of days
Comparison site	A similar stage or corridor that does not receive the intervention
Corridor Cell	A standing working group of the agencies and operators responsible for one corridor
Difference-in-differences	An estimate of an effect as the change at the treated site minus the change at a comparison site
Event study	A regression that estimates the treated–comparison difference week by week, used to check trends before an intervention
Generalised time	Travel time weighted by type, with walking and waiting usually weighted above in-vehicle time
Headway	Time between successive vehicles on a route
Hypercongestion	A state in which adding vehicles reduces the number of vehicles or trips completed per hour
ITS	Intelligent transport system: connected signals, detectors, cameras and a control centre
Matatu	Privately operated minibus
Monte Carlo simulation	Repeating a calculation many times with inputs drawn from probability distributions, to describe the uncertainty of the result
Paratransit	Flexible, privately run public transport such as matatus and boda bodas
Pigouvian toll	A charge equal to the delay a trip imposes on others
Pre-analysis plan	A public, time-stamped statement of outcomes, comparisons and statistical methods, registered before an intervention is evaluated
Rule-based software	Software that follows explicit, auditable rules (for example, routing and deadlines) rather than learned models
SACCO	Savings and credit co-operative; the legal form of matatu operators
Schedule delay	The difference between when a traveller wants to arrive and when they actually arrive
Traffic Management Centre	KURA's control centre at City Cabanas, where junction camera and signal data are monitored
Triple convergence	Downs' term for travellers shifting from other routes, times and modes onto an improved road until peak congestion returns
V/C ratio	Traffic volume divided by road capacity
Value of time	What a traveller would pay to save one hour of travel