Public Roads on Private Lands: Land Costs and Optimal Road Improvements in Urban Uganda

Jeanne Sorin*

July 26, 2026

Abstract

Cities in low-income countries need large investments in road infrastructure. Yet evidence on whether they pay off is limited, in part because the cost of the land they require is hard to measure and routinely ignored. I exploit the timing of the upgrades to estimate the local benefits of 114 km of donor-funded road improvements in Kampala, Uganda (2018–2026), which I embed in a general-equilibrium quantitative spatial model. I collect two new surveys of real-estate brokers and landowners and show that the land-acquisition costs could have amounted to 40% of the project budget under compensation at the legal market rate. As payments for land must be raised from scarce domestic public funds featuring a fiscal wedge, they carry a deadweight loss and paying all owners at market rate would have sharply reduced the otherwise substantial net welfare gains from the improvements. The wedge is highest where roads help most and therefore distorts where it is optimal to build: under market-value compensation the welfare-maximizing program captures only 70% of the first-best gains attainable without paying owners. The misallocation can be mitigated by donor financing of land or alternative compensation like the claim-based rule adopted for most of the studied improvements, under which the probability of compensation varies across the city's three coexisting land-tenure regimes, recovers 89% of those gains. The broader lesson is one of the second best: requiring full compensation to protect owners need not raise welfare when public funds are scarce, with direct implications for how road investments in fast-growing cities are designed, funded, and evaluated.

*University of California Berkeley; email: jsorin@uchicago.edu. I gratefully acknowledge financial support from IGC, STEG-CEPR and the University of Chicago. I am especially thankful to Milena Almagro, Fiona Burlig, Michael Greenstone and Esteban Rossi-Hansberg for their invaluable advice, as well as to Vittorio Bassi, Olivia Bordeu, Tom Hierons, Sasha Petrov, Tommaso Porzio, Tanya Rajan, Jordan Rosenthal-Kay and Oscar Volpe and numerous seminar and conference participants. This project could not have been possible without the support of the Julia Villenas, Eugenio Cassesse and Robert Mpiira at Dyadic Research Impact; George Makajuma at the African Development Bank; Robert Kyukyu, Gerald Ahabwe at KCCA. I thank Caesar Oweitu for generously sharing ride-hailing data, Lorenzo Casaburi and Martina Jakob for sharing aggregated land tenure information based on parcel-level datal, and Adam Storeygard for sharing historical GoogleMaps API data. Data collection for this project received IRB approval from the University of Chicago (IRB23-0270), from Hospice Africa Uganda Research Ethics Committee (HAU-2024-84) and from the Uganda National Council for Science and Technology (SS2510ES). All errors are my own.

1 Introduction

Roads are key to urban productivity, determining how easily workers reach jobs and firms reach customers (Duranton and Turner, 2012; Lall et al., 2017). Yet the cities of the developing world are short of them: they have fewer major roads (Akbar et al., 2023) and devote less land to streets — about 16% of built-up area in Sub-Saharan Africa, against 21% across the world's cities (Angel et al., 2016).

Closing this gap runs into a bottleneck: widening or extending a road means acquiring land from the residents and firms already on it. In principle, eminent-domain powers, common to most legal systems, let governments acquire this private land for public use (Lindsay, 2012). In practice, paying for it is a first-order constraint for low-income-country governments as they struggle to raise revenue (Besley and Persson, 2014) and external donors typically finance construction but not land acquisition (World Bank, 2017). Land prices vary sharply across the city, often peaking in the central areas, so the costs may outweigh the benefits even for the most valuable improvements (Collier and Venables, 2016). Land acquisition costs therefore shape not just how much a city builds but also where, so evaluations that ignore land costs misjudge both.

This paper shows that the cost of acquiring land is decisive both for the net returns to urban road improvements and for where they are optimally built. I jointly estimate the spatially heterogeneous benefits and land costs of 114 km of road upgrades carried out between 2018 and 2026 in Kampala, Uganda, a Sub-Saharan African city near the bottom of the global urban speed distribution (Akbar et al., 2023) where roads are effectively the only motorized mode. Valued at the legally required market rate, the land consumed by these upgrades would have amounted to about 40% of total project costs. I find that the upgrades raised welfare substantially, but the gains hinged on the government acquiring most of the required land without compensation. Paying the market-value compensation would have sharply reduced those gains and, where the cost of raising public funds is high enough, reversed them outright. Because the most valuable improvements ran through the most expensive land, accounting for land costs reshapes not just whether but where a utilitarian planner should build.

In Kampala the cost of acquiring land produced a striking arrangement: owners were legally entitled to market-value compensation, yet most went unpaid. Citing land costs that would jeopardize projects whose construction the World Bank (WB) and African Development Bank (AfDB) financed but whose land they did not, the government asked most owners to cede strips of land one to three meters wide without compensation. Only few owners pursued the compensation they were still owed, an outcome that varied across the city's land-tenure regimes and with the cost of documenting a claim (World Bank, 2023).

To measure the net returns to these road improvements I assemble the most comprehensive data yet collected on both the land costs and the benefits of urban road improvements in a low-income city. I conduct two original surveys: one with 565 real-estate brokers, used to reconstruct land values in a city with no transaction registry, and one with 546 affected landowners along upgraded roads, used to recover what acquisition actually cost and whether owners claimed compensation. I combine these with more than 110,000 Google Maps speed queries and, through a partnership with a local ride-hailing company, data on traffic flows since 2019. With these data I estimate the local effects of the improvements and the key parameters of a quantitative spatial model that I build to capture their citywide consequences and to formalize the trade-off between transport benefits and the cost of land.

For land costs, the main task is prediction: I recover land values from the broker survey and extrapolate the probability of compensation across the city from variation in land tenure, owner, and neighborhood characteristics. For local benefits, I identify the causal effect of upgrades on speed with a staggered difference-in-differences design, using the bias-robust Sun and Abraham (2021) estimator and comparing roads before and after upgrading to not-yet-upgraded controls, so that identification rests only on completion timing being as good as random with respect to a road's speed dynamics. Timing is not predicted by road characteristics, and speed rises as construction advances rather than jumping at completion, consistent with the upgrades themselves driving the gains. For property values, I lead with a ceteris-paribus broker appraisal; transaction-based estimates corroborate it but are contaminated by the continuous-accessibility spillovers the spatial model is built to internalize, so the paper's welfare conclusions rest on the structural elasticities I estimate directly: the response of speed to road width and of commuting flows to travel time.

The analysis rests on four results. First, I recover the cost of acquiring land and show that it was both far below the legal benchmark and spatially uneven. Affected owners in the survey lost 78 m^2 of land on average and compensating every owner affected by the road upgrades would have amounted to about 40% of total project costs. In practice the government paid far less and what it paid was determined on two margins. First, only 20% of owners were offered compensation. Second, most were instead asked to cede land without compensation, and 79% of them did so without claiming what they were owed. Extrapolated citywide, realized land costs came to between 40 and 57% of the underlying market value. Compensation was concentrated among owners on the more formal of Kampala's three coexisting tenure regimes: freehold and public (including leasehold) land as opposed to mailo land.

Second, the upgrades delivered large local benefits. Traffic speeds rose by about 4.3 km/h at rush hour (22% of speed on control roads) and upgraded roads became more likely to carry through-traffic. These gains capitalized into property values: brokers assessed that an upgrade would raise the price of an otherwise-identical roadside property by about 28% (residential) and 31% (business), with owners' own assessments at least as large, implying a local capitalization of about USD 147.1 million. Transaction-based estimates pointed the same way but were noisier.

Third, road improvements do more than raise local property values: they reshape traffic and the location choices of residents and firms across the city, while the land they consume carries an opportunity cost in scarce residential and business use. To capture these citywide, general-equilibrium effects I build and estimate a quantitative spatial model of a closed city in which workers trade off commuting costs against rents (Redding and Rossi-Hansberg, 2017; Allen and Arkolakis, 2022), and where residential and road uses compete for land. Compensation affects net welfare (beyond distributional considerations) because payments to owners are subject to a fiscal wedge: the shadow price of public funds. I estimate the model's key elasticities from the same data as the reduced-form results, complemented with ride-hailing data. The estimated responsiveness of commuting flows to travel times ($\theta\kappa = 0.0319$) is consistent with the few existing low-income-country estimates (Balboni et al., 2020; Kreindler and Miyauchi, 2023), and the elasticity of travel speed to road width ($\xi = 0.367$) is roughly three times larger than existing estimates for high- and middle-income countries (Fajgelbaum and Schaal, 2020; Bordeu, 2024). I find that the realized upgrades cut average commuting time by 3.64% and raised modeled citywide property values by 0.84%, generating resident welfare gains of about 3.7 times construction costs at a 6% discount rate — a conservative figure, since the model prices only the accessibility channel and holds local amenities and productivities fixed. Paying all owners market value would have sharply reduced these gains, the more so the higher the cost of public funds. In practice, because only a minority of affected owners were paid in Kampala, the policy's net welfare gain is between USD 94 and 118 million, or roughly USD 106 per working-age resident. The realized road improvements align relatively well with the model's predictions: upgrades were concentrated around links with high predicted gains and low predicted land costs.

Fourth, I use the model to study how the price paid for land shapes the optimal allocation of improvements. At the heart of the analysis is the fiscal wedge: because compensation is financed from costly public funds ($\eta > 1$), paying for land is not a costless transfer but carries a deadweight loss. Additionally, because the wedge is levied on the price of urban land, highest in the dense center where widening helps most, it tilts investment toward cheaper, lower-benefit peripheries. In Kampala, a link's land cost and its marginal benefit of widening are strongly positively correlated (rank correlation 0.68). To quantify the misallocation resulting from payments under this wedge, I hold the construction budget at the policy's envelope and I solve for the allocation a utilitarian planner would choose, weighing each link's commuting benefit against the opportunity cost of the land it consumes and the fiscal loss on land payments, in general equilibrium. I then turn to two policy levers: the compensation rule, which sets the level and spatial pattern of land payments, and the financing rule, which sets how costly they are to fund.

The compensation rule reshapes where the planner builds. Because land values are highest in the dense center where road benefits are also greatest, compensating every owner at market value pulls investment away from exactly the high-benefit locations whenever public funds are costly; paying nothing removes this distortion but shifts the entire land burden onto owners. The claim-based rule Kampala adopted, under which only owners who document a claim are paid, sits between the two: because only a minority of owners claim, and because the funds saved from unpaid owners are largest on the expensive central links where the market-value bill is highest, it moves the allocation closer to the first best than market value does.

These allocation effects are driven by the cost of public funds: were revenue costless to raise, the compensation rule would be purely distributional. This points to a second policy lever: the donor restriction that confines external funds to construction and forces land payments onto costly domestic revenue. Donors justify this convention on institutional grounds, treating compensation as the government's sovereign responsibility (World Bank, 2001; African Development Bank Independent Development Evaluation, 2019), but I show that it can have substantial efficiency costs. Lifting this restriction lets donors finance compensation and the wedge on those payments. Yet, because the total donor envelope is fixed, that money can no longer go to construction, so the city may build less. Comparing the two regimes across the cost of public funds, I find that relaxing the restriction has very minimal impact when funds are cheap (low η) but can substantially raise net welfare if public domestic funds are sufficiently costly (high η). Some donors are now adjusting their rules along this direction: in 2023 the African Development Bank revised its safeguards to permit financing land compensation for its projects, a reform that postdated the project studied here and was motivated by "*lessons learned based on the Borrowers' systematic failures to mobilize liquidity in time*" (African Development Bank (2023), p.16).

With these results, I make two contributions to a growing literature at the intersection of development and urban economics (Collier and Venables, 2016; Glaeser and Henderson, 2017). First, I measure the net returns to road improvements in a low-income city, accounting jointly for citywide benefits and the cost of the land they consume. Doing so bridges two literatures that have largely developed apart: one prices the benefits of transport but treats the land it consumes as free or spatially uniform; the other takes the cost of land seriously but outside spatial equilibrium. Evidence on the returns to transport infrastructure in developing economies is largely rural or interurban (Baum-Snow et al., 2017; Donaldson, 2018; Asher and Novosad, 2020; Morten and Oliveira, 2024; Balboni, 2023) and sets aside the cost of acquiring land — negligible where rights-of-way are wide, but first-order inside a dense city. Within cities, the evidence remains partial. Speed effects have been estimated for the United States (Duranton and Turner, 2011; Currier et al., 2023) or across cities worldwide (Akbar et al., 2023), and property-value effects in low- and middle-income cities only locally (Gonzalez-Navarro and Quintana-Domeque, 2016). Within-city spatial models in these settings have focused on bus rapid transit (Balboni et al., 2020; Tsivanidis, 2026; Zarate, 2024; Kreindler et al., 2023); in Kampala, as in most low-income cities, roads instead serve all motorized travel. On the cost side, I extend the optimal-transport-network framework (Fajgelbaum and Schaal, 2020; Allen and Arkolakis, 2022; Bordeu, 2024) to incorporate two key elements. First, I model the opportunity cost of land, which is itself a central determinant of urban welfare (Harari, 2020). Second, I account for the fiscal cost of compensating owners, forces flagged for low-income infrastructure by Collier and Venables (2016) and for the United States by Brooks and Liscow (2023) but not previously quantified in a spatial-equilibrium evaluation. In doing so, I bring the misallocation lens of Hsieh and Klenow (2009) to the fiscal cost of public infrastructure: the distorted margin is where a government builds and the wedge is the cost of public funds rather than a zoning restriction. I show that the resulting allocative distortion is governed by a single sufficient statistic — the covariance across links between the price of land and the marginal benefit of widening — so misallocation is largest exactly where costly land and high returns coincide.

Second, I provide new evidence on how property rights and the cost of compensating landowners shape both the returns to and the spatial allocation of public investment. A long tradition recognizes a tension between protecting owners through secure rights and just compensation, and providing public goods (Acemoglu and Robinson, 2012; Posner and Weyl, 2017). This tension is traditionally invoked to justify eminent domain for public purposes (Munch, 1976; Shavell, 2010; Jeong et al., 2016) but I show that in low-income settings this logic can invert through a fiscal wedge: because of the well-documented difficulty of levying taxes in low-income countries (Besley and Persson, 2014; Knebelmann, 2019; Bergeron et al., 2023; Brockmeyer et al., 2023) and the scarcity of the domestic funds compensation must draw on (Ministry of Finance, Planning and Economic Development, 2024), public funds are costly to raise and disburse, so paying the market-value compensation that eminent domain requires can turn welfare-improving projects welfare-reducing. Instead, acquiring land below market value, at a real distributional cost to owners, may make the investments viable. In low-income cities, what the government actually pays for land is governed less by market value than by the property-right regimes that shape who claims compensation. Because documenting a claim is harder under Kampala's less formal tenure regimes, the weakness of property rights itself lowers the share of owners the government pays, cheapens the land it acquires, and, when public funds are costly, moves the optimal allocation closer to the first best. Importantly, this speaks to the allocation of a given road program and does not imply that stronger property rights would lower overall welfare: the analysis counts owners' uncompensated losses in full and abstracts from the potentially large benefits of secure rights for individual welfare and private investment (Besley and Ghatak, 2010; Hornbeck, 2010). Governments' own responses to this constraint are part of the story: to save money, the Kampala government moved to a claim-based rule, under which the land costs it faced were shaped not only by market value but also by the city's historical tenure regimes (Bird and Venables, 2020) — a pattern related to Holland (2023), who studies the role of strong property rights in shaping opportunistic behavior by private owners in Colombia. My main contribution here is to show that this shift — compensating only owners who formally claim — reshapes not just who is paid but the optimal map of roads itself.

The rest of the paper proceeds as follows. Section 2 describes the context and the data I collect. Sections 3 and 4 present the reduced-form evidence, estimating respectively the land acquisition costs and the local benefits of the improvements. Section 5 develops the quantitative spatial model and Section 6 estimates it on Kampala. Section 7 characterizes the net returns to the realized road upgrades. Section 8 solves for the optimal allocation of improvements under alternative compensation and financing rules. Section 9 concludes.

2 Context and Data

2.1 Context

Kampala Road Network: Kampala, Uganda's capital, is a fast growing city, with approximately 1.9 million inhabitants in 2024 and more than 22% of the country's GDP Rafa et al. (2017). Yet its road infrastructure is of low quality: of the 2,578 km of roads and paths recorded by OpenStreetMap (OSM), only 449 km (17%) are major roads.[1] This scarcity is consequential: across 1,200 cities worldwide, Akbar et al. (2023) show that the mileage of major roads is one of the main determinants of urban travel speed. Major roads carry traffic faster because they are wider (7.4 m against 4.5 m for non-major roads in Kampala) and of higher quality (33% paved against 3%). Width and coverage, however, are ultimately a matter of land: a denser, wider road network is one that devotes more urban land to roads. On this dimension Kampala falls short. Its roads are narrower than both the global and the sub-Saharan African average and overall less than 13% of its built-up area is occupied by roads and boulevards, against 21% globally, as described in Appendix Table A8 (Angel et al., 2016).

Road Improvement Projects: Both width and quality dimensions have been targeted by road improvement projects in the city over the past decade. I focus on 114 km of road improvements covering 2018–2026 as part of two externally funded projects: the Second Kampala Institutional and Infrastructure Development Project (KIIDP-2, World Bank) and the Kampala City Roads Rehabilitation Project

[1] Akbar et al. (2023) define major roads as roads falling under the motorway, primary, secondary or tertiary road classes on OSM.

(KCRRP, African Development Bank).[2] The projects, displayed on Figure 1 and spanning all five Kampala divisions, are part of a broader envelope totaling more than $470 million, which also funds capacity building, junction signalization, drainage, and non-motorized transport facilities. Based on AfDB (2014) cost estimates, inflated from 2006 to 2019 USD using the US CPI, road construction alone accounts for about 63.9 millions of this envelope ($186,575 per lane-km, assuming road projects (re)pave the average existing road (2 lanes) and add one). Importantly, while construction costs are funded by external sources, the government of Uganda is responsible for raising the funds needed for the acquisition of land, needed to add lanes, drainage channels, sidewalks, non-motorized transport tracks, and road shoulders that accompany all project roads.[3] I therefore measure road improvements as increases in effective road area – the total width of land devoted to road infrastructure, including both the carriageway and the adjacent right-of-way.

Land Acquisition for Infrastructure Projects in Kampala: The acquisition of this adjacent land is governed by Uganda's eminent domain framework, enshrined in Article 26 of the 1995 Uganda Constitution and Section 42 of the 1998 Land Act, which allows the government to acquire private land for public infrastructure against *fair* compensation, or pre-project market-value compensation.[4] Shortly after the first wave of investment (KIIDP2) began, the city government asked minimally affected owners to forfeit their small parcels without compensation, citing prohibitively high land costs under eminent domain threatening the project. Total disbursements for land acquisition therefore depend on the amount of land being acquired, the pre-policy market value of that land, the share of owners being offered compensation, as well as the share of owners consenting to *voluntary land take*.

This latter approach was exploratory, and while World Bank (2023) acknowledged that it may be a cost-effective strategy, it also documented concerns — raised in an Inspection Panel complaint alleging coercion during the resettlement process — that owners' consent to cede land may not have been fully voluntary. The complexity of Kampala's land tenure regimes compound this concern: land disputes and issues of multiple land claimants per plot must be resolved before compensation can be disbursed, further discouraging owners from pursuing their claims.

Coexistence of Multiple Land Tenure Regimes in Kampala: Land ownership in Kampala is governed by institutions inherited from Uganda's colonial history. In 1900, the Buganda Agreement divided land in Kampala between (i) the Buganda Kingdom (*mailo land*), (ii) the British Protectorate and (iii) remaining unallocated land. This initial partition led to the consolidation of several land tenure systems in Kampala throughout the twentieth century and Uganda's independence. Bird and Venables (2020) document how these tenure systems have shaped patterns of land use in the city.

A map of the different land tenure regimes in Kampala, built from shapefiles shared by the Kampala

[2] An additional 52 km of domestically funded road improvements were implemented over the same period, which I include in the reduced-form analysis. I do not have data on regular road maintenance.

[3] While official project documents classify interventions as any combination of reconstruction of degraded pavement, upgrading from unpaved to paved standard, or dualing (widening to dual carriageway), all three types may involve the acquisition of adjacent private land and respondents with affected land were interviewed in the vicinity of all three types of road upgrades.

[4] More details on the land acquisition timeline can be found in Appendix Section A.1.

Capital City Authority GIS unit is displayed in Figure 2.[5] Land at the center of the city was historically under the administration of the British Protectorate and is now dominated by freehold titles and long-term leases and hereafter referred to as *freehold*. Land to the north, and intermittently in the west and south, is *mailo* land – historically owned by large, often absentee landlords, but where tenants (*kibanja* holders) hold protections and rights over their structures. The remaining land, common in the west, south, and east, was historically unallocated and is now under the *public* authority, which may be leasing or selling individual parcels.

2.2 Main Data Sources

To study the land costs and local benefits of road improvements in Kampala, I gather detailed information on property values, the land acquisition process, and traffic speed by conducting two surveys and with more than 110,000 Google Maps API queries. I then combine these datasets with traffic flow data from a partnership with a local ride-hailing company and publicly available census data, to estimate key parameters of a model capturing the policy's spatial spillovers.

2.2.1 Property Values: New Broker Survey

As data on property values is not readily available in developing country cities like Kampala, I conduct a survey with 565 real estate brokers in 90 of Kampala's 96 *parishes* (neighborhoods) (Appendix Figure A1) between March 2024 and December 2025 to (1) characterize the spatial distribution of property values in the city – key ingredient to both the land payment calculations and the quantitative spatial model – and (2) shed light on the impact of road improvements on local property values. In the 3 months preceding the interview, the median broker had rented 6 properties and had sold 1 properties. The broker sample and data collection is described in more detail in Appendix B.1.2 and Appendix Table A3.

The survey has two main components. First, brokers were asked to report the five most recent properties of each type they rented and sold. From these answers, I build a retrospective panel of transactions (sales and rentals) including information on the transaction date (month, year), location (parish, distance to closest major road) and on a large set of property characteristics. The final database includes 5,007 properties (2,498 residential properties, 1,657 business properties and 852 bare lands), transacted between 2018 and 2025.

Property values in Kampala are high and spatially heterogeneous. For example, the median bare land sale for properties transacted in 2023 onward in my sample was 32,432 USD, corresponding to 120 USD per m^2 for residential properties and 153 USD per m^2 for business premises.[6] Market rates are very heterogeneous across Kampala and higher at the center, as evidenced by Figure 3 and Appendix Table A3. To build confidence in the elicited transactions, in Appendix Section B.1.2 I show they are positively correlated with online-scraped high-end property data and more consistent with publicly available summary statistics on Kampala property values.

[5] This map, shared by the Kampala Capital City Authority, should be understood as adapted from historical maps rather than representing parcel-level information. In Appendix A.2, I show that this map is consistent with alternative sources.

[6] For comparison, Bassi et al. (2026) recovered a median sales rate/m^2 for manufacturing premises in Kampala in 2019 of 141 USD.

In the second part of the survey (2024 wave only), I conduct a property appraisal exercise to gather additional reduced-form evidence of the local impacts of road improvement. Brokers estimate the sales price of a standardized property (business or residential), with characteristics varied one at a time. The characteristics include the property's walking distance to a busy main road (0, 5, 10, 15+ minutes) and the upgrade status of this road.

2.2.2 Land Acquisition: New Owner Survey

To learn about the specific land acquisition process for road upgrades in Kampala, I survey 546 owners whose land or property was affected by road upgrades along 47 roads in Kampala. My main sample is restricted to the 310 respondents who reported that the government officially reached out to them about their land or property being affected, either with a compensation offer or a formal request to consent. A further 236 surveyed owners report being affected but were not formally engaged, which I treat as a robustness sample (Appendix) since I do not observe their claim decision uniformly.

The average respondent in my main sample had owned their property for 35 years, and 43% of them had purchased the property, rather than inherited it (Appendix Table A4). The average respondent reported losing 78m^2 of land to the road upgrade.[7] As further discussed in Appendix Section B.2.3, the average owner in my survey is wealthier than the average Kampala resident, consistent with landowners being wealthier than non-landowners.

2.2.3 Traffic Speed: Google Maps API Data

I query Google Maps Direction API between March 2023 and January 2026 to consolidate a panel of 118,150 trips. The time coverage of queries is displayed in Appendix Figure A3. Each query returns the trip's origin, destination, time, distance, straight-line distance, traffic at query time, average duration, and a list of coordinates approximately tracing the path between origin and destination.

I sample two types of trips. First, to estimate the impact of road upgrades on local speed, I query a total of 2,175 trip instance on 82 unique trips (41 roads, both ways) for roads targeted by upgrades since 2018. Second, to assess whether a road is more likely to be used for a given trip after being upgraded, I query 115,975 trip instances between all pairs of Kampala parish centroids (9,120 unique trips) and map them to the underlying OSM road network.

Across the 2023-2025 period covered by the data, I find that the average speed in Kampala across all pairs of parishes (unweighted) was 23.3 km/h (20.5 km/h at rush hour and 25.7 km/h at non-rush hour). This corresponds to cities between the 10th and the 25th percentile of the world speed city distribution (Akbar et al. 2023). Improving road speed is thus a first order concern in Kampala.

I complement this new Google Maps data with queries from 2019 from Akbar et al. (2023) covering Kampala. These 640,880 trip instances are aggregated by origin and destination parish to recover average commuting time between parishes before the greatest part of the road upgrade policy was implemented.

[7] Out of 310 affected owners in the main sample, 259 were on the sides of roads with finished upgrades and 51 were on the sides of roads with ongoing upgrades. The average number of owners interviewed per road is 6.6 (min 1, max 24).

2.2.4 Additional Data Sources

I complement these datasets with several data sources. I recover the network of major roads in Kampala from the geo-referenced dataset published by the World Food Program (WFP), which follows the United Nations Spatial Data Infrastructure for Transport standards.[8] To study how workers will adjust their location choice in response to a change in local traffic speed, I partner with a local ride-hailing company running a popular app with more than 2.2 million rides per day and I describe this data in more detail in Appendix B.4. I also use the 2016 Kampala Travel Habit Survey (Appendix B.5) and publicly available data from the 2014 Population Census and 2019 Census of Business Establishments (Section 6.2).

3 Reduced-Form Evidence: Land Acquisition Costs

I leverage my surveys to establish four facts about the cost of acquiring land for road improvements in Kampala. First, under the eminent-domain regime with compensation at market value, the government would have disbursed 43.4 million USD for land — 68% of construction costs alone — and these land values are highly heterogeneous across space. Second, in practice, few owners were compensated. Third, high costs prevented many owners not initially offered compensation from claiming it, and 77% forwent compensation. Fourth, which owners are compensated is strongly tied to the underlying land tenure: mailo respondents are both less likely to be offered compensation and, when not offered, less likely to claim it.

3.1 Fact 1: High and Heterogeneous Land Payments with Market Value Compensation

Under the eminent-domain legal regime, owners would be compensated at (pre-policy) market value. In my survey, road improvements took 78 m^2 of land from the average interviewed owner (Appendix Table A4), corresponding to a band 2.8 m wide, or about one additional lane (average lane width 3.4 m).[9] Extrapolating to the 114 km of road upgrades studied, I recover that a total of approximately 5,073 owners were affected.[10]

The market value of this land is substantial. As detailed in Section 6.2 and Appendix C, from the broker survey I estimate that the average sales price per m^2 for residential properties was 105 USD in 2018 (208 USD for business properties). Taking a parish-level weighted average and assuming 3.4 km^2 of land acquired per km^2 of road upgrade (one additional lane), total land payments under market-value compensation would have summed to 43.4 million USD, or about 4,190 USD for the median owner in my survey.

These market rates are highly heterogeneous, central parishes far higher — as is standard in a monocentric city — with a parish-level standard deviation of 52 USD/m^2 (Figure 3).

[8] Map extracted in 2018/2019 (Source). Major roads have the following tags: *highways, primary/trunk, secondary, tertiary.*

[9] Effective land take is dispersed — reflecting whether widening is one- or two-sided, parcel frontage, and where the road/private boundary is drawn: the 25th, 50th, and 75th percentiles are 14, 37, and 74 m^2.

[10] The official count of affected owners for the KCRRP portion of the project (70 km) was 4,413 (Kampala Capital City Authority, 2025).

3.2 Fact 2: Few Owners Are Compensated

An owner can obtain compensation through two channels: a direct government offer (channel 1) or, absent an offer, an owner-initiated claim (channel 2). Overall, I estimate that only between 33% and 48% of owners have (or will) receive compensation for their affected land.

Through the first channel, only 34% of affected owners in my main sample (official contacted by the government) were offered financial compensation upfront (106 of 310 respondents). Of those offered, 94% reported having been paid by the time the survey was conducted, or 33% of the main sample, which I treat as a floor.[11] Consistent with the government's move toward *voluntary land take* as fiscal strain mounted after the initial upgrades (Section 2), the road-level share of owners offered compensation falls with the upgrade year ($\rho = -0.31$).

As the government withdrew from offering, the burden of obtaining compensation shifted toward the second channel. Yet among owners not initially offered compensation, only 22.6% claimed it, which I define as refusing to consent to forgoing their land for at least six months.

3.3 Fact 3: Claiming Compensation Is Costly (Channel 2)

So few owners pursue the second channel because, despite its potential reward — market-value compensation — claiming is costly: refusing to consent to *voluntary land take* means contesting the acquisition through administrative and legal steps, at substantial financial and social cost.

To quantify what drives claiming, I estimate a logit for the probability of an owner o in parish i claiming financial compensation on three owner characteristics capturing each: (i) a standardized wealth index (averaging eight household asset and human-capital indicators: education, electricity, floor type, bank account, television, computer, bicycle, and meals per day); (ii) documentation costs, proxied by a dummy for above-median self-reported cost of obtaining ownership documents; and (iii) social integration, proxied by an owner's (standardized) number of friends in the neighborhood.

I estimate the model by maximum likelihood on the 212 owners not offered compensation with non-missing claiming decisions and display the estimated coefficient in Appendix Table A6. The different channels move claiming in the expected direction. Wealthier owners are more likely to claim (log-odds 0.406, s.e. 0.211), consistent with pursuing a claim requiring resources. Owners facing above-median documentation costs are less likely to claim (log-odds -0.375, s.e. 0.391), consistent with the cost of proving ownership deterring claims. More socially integrated owners are *less* likely to claim (log-odds -1.169, s.e. 0.54), suggestive of social pressure to consent so the project can proceed in a timely manner.

These patterns line up with owners' stated reasons for consenting and displayed in Appendix Figure A5.[12]

[11] Reported receipt understates eventual compensation for two reasons. First, timing: the survey (Spring 2024) predates the release of additional Government-of-Uganda and African Development Bank funds that unlocked stalled projects, while acquisition and litigation were still ongoing. Second, informal payments go under-reported and are common: y% report bribes are at least frequent with public servants, and 49% report some neighbors were compensated (30% officially, 27% unofficially). I therefore also use the offered-or-claimed measure as an upper bound.

[12] Common reasons for consenting (local amenities (48%) and expected property-value gains (32%)) reflect support for the road, not a willingness to forgo compensation: owners enjoy these benefits whether or not they are paid, and only 4% believed the road would not be built if they refused. Owners thus see themselves as non-pivotal, so support for the road cannot rationalize giving up land for free — the operative deterrent is the claiming cost above. This perceived non-pivotality,

31% of them reported feeling powerless against the government, and only 24% thought they were very likely to obtain a higher payment by negotiating. The financial barrier is steep: the median reported cost of obtaining a copy of one's ownership documents — likely only part of the full claiming cost — is $811, against a median owed compensation of only $3,228 for owners not initially offered compensation. Social considerations also weigh: 64% said they consented because the upgrade benefits the community and 14% specifically cited pressure from their community.

3.4 Fact 4: Owners on More Formal Land Tenure Are More Likely to Be Compensated

I turn to *which* owners are compensated. Across every measure, owners on mailo land are less likely to be compensated. I regress each compensation measure on land-tenure type (relative to mailo), the standardized market value of the affected land, a standardized wealth index, and parish-level controls (population, share of owner-occupied households, area).

The results are presented in Table 1, on the full sample in odd columns and on the subsample of already-completed (past) upgrades in even columns, where financial procedures are less likely to be ongoing. Freehold owners are 36.5 (s.e. 9) percentage points more likely to access compensation through either channel than mailo owners, and owners on public land 28.8 (s.e. 6) percentage points more likely (columns 1–2). Columns 3–8 separate the channels. Owners on more formal tenure are more likely to be *offered* compensation upfront — freehold and public owners by 19 (s.e. 8.5) and 29 (s.e. 5.8) percentage points (columns 5–6) — and, when not offered, more likely to *claim* it: the logit coefficients on freehold and public are 1.77 (s.e. 0.56) and 0.68 (s.e. 0.41) (columns 7–8). The market value of the affected land also plays a role in owners' decision to claim the compensation. While the coefficient is imprecisely estimated, it is consistent with owners' increasing the probability they claim, the higher the potential reward.

Overall, this pattern is suggestive of strategic behavior by the government, which appears to anticipate that formal-tenure owners would be harder to dispossess without payment and offers them compensation upfront. The freehold advantage fades only for self-reported *receipt* (columns 3–4), where public tenure remains the significant predictor, consistent with the under-reporting of receipt noted in Fact 2.

Mechanisms. While a full account of why tenure formality matters is beyond the scope of this paper, the survey points to documentation costs, the perceived security of rights and social costs of claiming (Appendix Table A7): freehold owners are 19 (s.e. 9) percentage points less likely to face above-median documentation costs, and public owners report higher land security (0.54 s.d., s.e. 0.11). Additionally, on mailo land, the registered landowner and the *kibanja* holder (both considered as owners) would split the compensation, reducing individual benefits of claiming. Both public and freehold owners also report fewer friends in the village than mailo owners. Freehold and public owners are *less* wealthy than their mailo counterparts, so the tenure gradient is not driven by wealth.

Robustness. The tenure gradients hold on the completed-upgrade subsample (even columns of Table 1), are unchanged when the wealth and parish controls are dropped, are robust to adding further individual

unsurprising under eminent domain, is also why I do not treat holdout as a significant force.

controls (number of neighborhood friends, household-head status), to alternative parish-control sets, to alternative dependent-variable definitions (extending the sample to owners reported as affected but not officially reached out to, and imputing non-claim for owners with no recorded decisions), and to sample cuts (past upgrades only, or post-2018 only), as displayed on Appendix Figure A7.[13]

4 Reduced-Form Evidence: Benefits of Road Improvements

Having documented the cost of acquiring land, I now turn to the benefits road improvements deliver and I establish three reduced-form facts: faster local traffic, more trips drawn onto the improved segments, and higher nearby property values. These estimates are local and partial-equilibrium by construction; they motivate, and later discipline, the quantitative general equilibrium spatial model of Section 5.

4.1 Road Improvements Increase Local Traffic Speed

Empirical Strategy: To assess the extent to which road improvements affect local traffic speed, I regress the Google Maps estimated traffic speed in km/h[14] on road r at hour h on day d on a dummy variable equal to one if road r was upgraded by day d:

$$s_{h,d,r} = \alpha + \beta\, D_{d,r} + \beta^{\mathrm{p}} D^{\mathrm{p}}_{d,r} + \mathbf{X}'_r \psi + \gamma_h + \gamma_d + \epsilon_{h,d,r}. \tag{1}$$

I include day γ_d and hour-of-day γ_h fixed effects to control for citywide temporal variation. $\mathbf{X}'_r$ is either a vector of road-level characteristics (road class, pavement status, geographic controls) or road fixed effects, and I control for whether the road upgrade is in progress, $D^{\mathrm{p}}_{d,r}$. Because the speed data cover March 2023 to January 2026 while road upgrades began in 2018, the identifying variation in the two-way fixed-effects and staggered difference-in-differences specifications comes from the subset (19 of 41) of roads observed both before and after completion. In these specifications I collapse the trip-level data to a road $\times$ month panel, residualizing speed on day and hour-of-day fixed effects at the trip level and averaging within road-month cells. For the Sun and Abraham (2021) estimator I report the average treatment effect on the treated over the first 12 months following completion. Standard errors are clustered by road segment and day for trip-level regressions, and by road segment for the panel.

Identification Assumptions: I present cross-sectional, two-way fixed-effects, and staggered difference-in-differences specifications, estimated separately for rush hours (morning and evening peaks) and off-peak hours, in Table 2. The causal interpretation of β rests on assumptions whose strength differs across specifications. The cross-sectional OLS (columns 1 and 5) makes the strongest assumption: β is the ATT under selection on observables — conditional on road controls $\mathbf{X}_r$ and day/hour fixed effects, upgrade status is unrelated to other determinants of the speed *level*.

The strongest assumption is in the cross-sectional OLS specifications (columns 1 and 5), which compare upgraded to not-yet-upgraded roads, where β can be interpreted as the ATT under selection on observables.

[13] Because the gradients are unchanged on the extended sample, differential non-response is unlikely to drive them; any residual bias is of ambiguous sign, since more socially integrated owners are *less* likely to claim.

[14] Google Maps times are car-based, but minibus taxis and private cars — over 70% of home-to-work trips (Travel Habit Survey) — have similar speed technologies (Table A19), so car times proxy well for most commuters' speeds.

In other words, conditional on the road-level controls $\mathbf{X}_r$ and on day and hour-of-day fixed effects, upgrade status is unrelated to other determinants of the *level* of speed. The two-way fixed-effects and staggered difference-in-differences specifications (columns 2–4 and 6–8) include road fixed effects and therefore rely on a weaker assumption: they absorb time-invariant road heterogeneity and identify β from the *timing* of completion, requiring only that the timing of completion is as good as random with respect to a road's speed dynamics. Using not-yet-upgraded roads as the comparison group, rather than never-upgraded roads, weakens this assumption: it holds fixed the selection of roads *into* the program and leans only on timing.[15] Treating never-upgraded roads as controls would instead require no selection into the program on predicted speed gains, an assumption the program's appraisal document contradicts, since selection among candidate roads was based on economic viability (net present value and economic internal rate of return), traffic volumes, and congestion-reduction potential (World Bank, 2014).

The timing assumption underlying my preferred specifications is plausible to the extent that upgrade timing reflects construction sequencing and land-acquisition delays rather than expected speed gains, and I provide two pieces of supporting evidence. First, in Appendix Table A11 I regress the start and completion months of each upgrade on road characteristics: observable characteristics do not clearly predict the timing, particularly once geographic controls are included. The one exception is road length, which predicts the completion date but not the construction start date — consistent with longer roads taking longer to build rather than with selective entry into the program. Second, the event studies (Figure 4) show speed rising gradually over the construction period and stabilizing after completion, consistent with the road improving progressively as works advance rather than with a discrete jump at completion; the pre-construction periods, though limited in number, show no clear differential trend.

Results: I find that road upgrades consistently raise traffic speed (Table 2), both at rush and off-peak hours. My preferred estimate is the bias-robust Sun and Abraham (2021) interaction-weighted estimator at rush hour, 4.3 (s.e. 1.53) km/h (column 4). I favor it for two reasons. First, rush hour is when the commuting the model prices actually occurs. Second, the specification is robust to the negative weighting that biases two-way fixed effects when treatment timing is staggered and effects are heterogeneous. I use the same specification to calibrate the speed–width elasticity ξ (Section 6.1.1). The cross-sectional and two-way fixed-effects specifications are of similar magnitude at rush hour (columns 1–3). The effect is larger at rush hour than off-peak, consistent with upgrades relieving congestion where it binds: baseline speeds are lowest at rush hour (21.25 versus 24.01 km/h off-peak). Appendix Table A10 pools all hours of the day, yielding an average speed gain of 5.22 (s.e. 0.88) km/h with the full set of controls and 1.39 (1.07) km/h from the windowed Sun–Abraham estimator. Appendix Table A9 reports robustness to the control set and to the un-windowed Sun–Abraham estimator.

[15] Upgrading one road can alter speed on others, including not-yet-upgraded controls, violating SUTVA; but the bias is likely second-order, since the traffic any single upgrade redistributes is spread thinly and any one control absorbs little. I verify this: nearby completions do not predict still-control speeds, and controlling for exposure leaves the estimate essentially unchanged (Table A12).

4.2 Improved Roads are More Likely to Be Used by Commuters

Empirical Strategy: I further assess whether road upgrades lead to changes in traffic patterns by constructing a road-segment $\times$ month panel recording how many of the parish-to-parish trips from the Google Maps data approximately pass through each policy road segment. I regress both the number and the share of parish-to-parish trips on the upgrade completion dummy, with road-segment and month fixed effects, clustering standard errors by road segment and month.

Results: Results are presented in Table A13. Completed upgrades are associated with a 0.0059 (s.e. 0.0018) increase in the share of parish-to-parish trips through the segment, against a mean of 0.0199.[16] The effect holds but is noisier when restricting the sample to roads upgraded after April 2023 (columns 2 and 4), where the within-road variation is cleanest. This spatial reallocation of traffic is a key mechanism the model of Section 5 endogenizes.

4.3 Road Improvements Increased Local Property Values

Empirical Strategy: I provide evidence that road upgrades increased local property values using two complementary approaches. First, I exploit the stated-preference module of the broker survey, in which respondents reported the value of a standardized property at varying distances from a busy road, before and after a hypothetical road upgrade. I estimate the price difference attributable to the upgrade within $\{0, 5, 10, 15+\}$ minutes walk from the property by running a within-broker specification with broker fixed effects. I also use owners' self-assessed property values — at the time of interview and absent the upgrade — estimating the upgrade effect with owner fixed effects on the completed-upgrade sample. I treat the broker stated-preference estimate as my preferred *local* measure: its ceteris-paribus design holds the surrounding network fixed and is immune to spatial spillovers.

Second, I move to revealed preferences by leveraging the retrospective panel of actual transactions reported by brokers, linked to a parish-year (p, y) panel recording the timing of road upgrades. Because the transaction data includes property characteristics (floor type, number of rooms, building type, land dimensions for sales only), in a first stage, I residualize log prices on hedonic controls and month fixed effects, and in a second stage I estimate the treatment effect on residualized prices as follows:

$$\text{Stage 1:} \quad \log(P_{i,p,y}) = a + \mathbf{X}_i' b + e_{i,p,y} \tag{2}$$

$$\text{Stage 2:} \quad \tilde{p}_{i,p,y} = \alpha + \beta_{post} \mathbf{1}\{\text{post}_{p,y}\} + \gamma_y + \mathbf{Z}_p' \mu + \epsilon_{i,p,y}, \tag{3}$$

for property i with $\tilde{p}_{i,p,y}$ the residualized value of the first stage, $\text{post}_{p,y}$ a dummy equal to one after the road upgrade in parish p (from 2018 onwards) and $\mathbf{Z}_p$ a vector of neighborhood characteristics (sub-county fixed-effects, parish-level controls or parish fixed-effects). For sales, in the second stage, given the longer temporal coverage, I also implement the staggered difference-in-differences (DiD) estimators from Callaway and Sant'Anna (2021) and Sun and Abraham (2021), which are robust to heterogeneous treatment effects across cohorts. For rental transactions, in the second stage, I estimate pooled OLS

[16] These numbers cannot be interpreted as changes in total flows on these roads, which would require weighting parish-to-parish pairs by commuters, data unavailable in Google Maps.

specifications with increasing sets of location characteristics: sub-county fixed effects and parish-level socio-economic controls, parish fixed effects. With parish fixed effects, the identification assumptions to interpret β_{post} as the causal average treatment effect on the treated (ATT) of parish-level road upgrade on local rents are (1) the absence of anticipation effects – especially plausible for rentals – and (2) the orthogonality of the road upgrade timing and unobserved determinants of rental rates, conditional on receiving an upgrade. The uncertainties and delays surrounding road improvement projects in Kampala, which were partially driven by land acquisition issues, strengthen the plausibility of this assumption. Any residual anticipation would fall on sales, one reason I do not treat the transaction estimates as my preferred measure.

Stated-Preferences Results (Figure 5): In the hypothetical property exercise, brokers expect road upgrades to raise the value of properties by the road by 28% for residential and 31% for business properties, on average. The distribution of stated answers is right-skewed, so the typical (median) broker reports a smaller gain: a post-/pre-upgrade price ratio of 1.25 (residential) and 1.22 (business) (Appendix Figure A8). These estimates lie within the confidence interval of owners' self-assessed change in property values from the upgrade. Owners' stated-preference estimates are larger than brokers' (despite also capturing land losses from the road upgrade), a pattern consistent with the literature (Gonzalez-Navarro and Quintana-Domeque, 2016).

Revealed-Preference Results (Table A14): Transaction-based estimates confirm that road upgrades in a parish are associated with an increase in local (parish-level) property values. The estimated impact on rents is 0.107 log points (s.e. 0.056, p-value 0.0598) without parish fixed-effects, and 0.187 log points (s.e. 0.108, p-value 0.0877) when including parish fixed-effects. The rent panel is too short to support parish fixed-effects or staggered DiD estimators. Sale-price estimates follow the same pattern but are less precise, reflecting smaller samples and greater heterogeneity across sale transactions. Pooled OLS gives 0.095 (column 3), and the TWFE estimate attenuates toward zero (0.024), consistent with heterogeneous treatment effects under staggered adoption; the Callaway and Sant'Anna (2021) and Sun and Abraham (2021) estimators that correct for this recover larger effects of 0.145 and 0.284.

Spatial Spillovers: Because accessibility is inherently continuous, a road upgrade raises the value not only of adjacent properties but of those some distance away. This spillover is made visible by the property appraisal as displayed on Figure 5: while brokers expect the largest increase for properties immediately adjacent to the road (28% residential), they still expect a 15–25% increase for properties 5, 10, and even 15 minutes' walk away.[17] This spatial spillover violates SUTVA in the parish-level transaction design, therefore motivating using the ceteris-paribus stated-preference appraisal as my preferred estimate of the local effect. More fundamentally, spillovers of this kind cannot be recovered by any local regression; capturing the citywide, general-equilibrium effect of the upgrades — and hence their full welfare consequences — is what the structural model of Section 5 is built to do.[18]

[17] Kampala parishes are small: the average parish covers about 2.03 km^2, corresponding to a typical linear extent (approximated as the square root of the area) of roughly 1.26 km. A 15-minute walk (about 1.25 km at 4 km/h) is therefore comparable to a little less than the width of an entire parish.

[18] The local estimates capture both accessibility and improved neighborhood amenities (less dust and flooding, safer streets; Appendix Figure A9). The structural model isolates the *accessibility* channel and holds amenities fixed, so its welfare gains

5 Model

Road benefits spill across space, and land devoted to roads carries an opportunity cost in residential and business use. Because these forces put the citywide net returns beyond the reach of any local estimate, I propose a static general equilibrium model of a closed city, populated by a fixed number of households who choose where to live, work, and which route to commute on. Commuting happens on a road network, that connects neighboring locations. The road network takes up land, so that in each location, residential use, business use and road infrastructure compete for land. Private owners own residential and business land in each location. A city government determines the optimal road widening to maximize city residents' welfare, tying together the spatially heterogeneous benefits, opportunity costs of land use, and fiscal burden associated with land payments to private owners.

5.1 The Quantitative Spatial Model at Baseline

Environment: The city is composed of distinct locations, indexed by $i, j \in \mathcal{J}$. Locations are connected through a road network (undirected graph), where each location i is connected to $\mathcal{C}(i)$ other neighboring locations. Workers must use this road network to commute. Locations are endowed with fixed amounts of land H_i. Each location is also endowed with exogenous residential amenities B_i, productivity A_j, and land tenure regime Z_i. Each link between neighboring locations kl has a baseline amount of road infrastructure of width R^0_{kl} meters.

Land and floorspace: At baseline, total land in each location H_i is divided between private land H_i^{pr}, and public road land H_i^{pu}, such that:

$$H_i = H_i^{\mathrm{pr}} + H_i^{\mathrm{pu}} \quad \forall i. \tag{4}$$

Public road land H_i^{pu} is the sum of the area used by roads of width R_{mi} and length in l_{mi}, connecting i to neighboring locations $m \in \mathcal{C}_i$ and assuming that the road is split equally across the two locations,

$$H_i^{pu} = \sum_{m \in \mathcal{C}_i} R_{mi} \times \frac{\mathrm{l}_{mi}}{2} \quad \forall i. \tag{5}$$

Private land H_i^{pr} is split between residential and business uses and converted into residential H_i^r and business H_i^f floorspace according to

$$H_i^r = \gamma_i \mu_i^r H_i^{\mathrm{pr}} \tag{6}$$

$$H_i^f = (1 - \gamma_i) \mu_i^f H_i^{\mathrm{pr}}, \tag{7}$$

where γ_i is the share of private land allocated to residential use in i and μ_i^r and μ_i^f convert residential and business land into floorspace, consumed by households and used in firms' production function.

Land Owners: Private land H_i^{pr} is owned by $\bar{N}_i$ local immobile representative landowners who get utility from consuming the freely tradable good. N_i^{or} owners are near roads and N_i^{of} owners are far from

understate full local capitalization by the amenity component, making them conservative.

roads. At baseline, owners receive rental income l_i^o over their private floorspace $\frac{H_i^r+H_i^f}{\bar{N}_i}$, such that owners' welfare is concave in their rental income, with

$$W_i^o = \begin{cases} \bar{C}^o \ln l_i^o & \text{if } \psi^o = 1, \\ \bar{C}^o \dfrac{(l_i^o)^{1-\psi^o}}{1-\psi^o} & \text{if } \psi^o \neq 1, \end{cases} \tag{8}$$

with ψ^o exogenous and $\bar{C}^o$ a constant that renders owner utility money-metric by normalizing owners' marginal utility of income to one at a common reference income, so that a marginal dollar is valued equally across workers and owners at that reference. Owners richer than the reference therefore have marginal utility below one.

The number of owners in location i, $\bar{N}_i$, is proportional to the location's baseline rental value, so that owner income is equalized across owners at baseline and the curvature of W_i^o is well defined at the common baseline income.

Workers: The city is inhabited by a fixed mass $\bar{L}$ of mobile workers who freely choose where to live i, work j, their commuting route r, as well as how much to consume of the tradeable good C_{ij} and floorspace H_{ij} subject to their budget constraint:

$$\max_{i,j,C_{ij},H_{ij}^r} U_{ijr} = B_i \left(\frac{C_{ij}}{\beta}\right)^{\beta} \left(\frac{H_{ij}^r}{1-\beta}\right)^{1-\beta} \epsilon_{ij}. \qquad \text{s.t.} \quad \frac{w_j}{\tau_{ij}} = C_{ij} + H_{ij}^r q_i^r. \tag{9}$$

q_i^r is the residential rental rate in location i, B_i the residential amenity in location i, w_j the wage in location j and τ_{ij} the average iceberg commuting costs between i and j. Commuting costs τ_{ij} between their residence i and workplace j decrease workers' available income $\frac{w_j}{\tau_{ij}}$. ϵ_{ij} is a preference shock over the pair ij drawn from a Frechet distribution with shape parameter θ. In equilibrium, workers' expected welfare W^w is equalized over space, so that

$$W^{\text{w}} \equiv \left(\sum_{ij} \left(\frac{B_i w_j}{(q_i^r)^{1-\beta}\tau_{ij}}\right)^{\theta}\right)^{\frac{1}{\theta}}. \tag{10}$$

The larger θ, the more responsive are workers to local economic conditions given by $\frac{B_i w_j}{(q_i^r)^{1-\beta}\tau_{ij}}$, as their idiosyncratic preferences are less dispersed.

These preferences imply that the number of workers commuting from i to j in equilibrium is given by

$$L_{ij} = \bar{L}\frac{\left(\frac{B_i w_j}{(q_i^r)^{1-\beta}\tau_{ij}}\right)^{\theta}}{\sum_{mn}\left(\frac{B_m w_n}{(q_m^r)^{1-\beta}\tau_{mn}}\right)^{\theta}} = \bar{L}\frac{\left(\frac{B_i w_j}{(q_i^r)^{1-\beta}\tau_{ij}}\right)^{\theta}}{(W^w)^{\theta}}. \tag{11}$$

The local number of residents increases in the rent-adjusted amenities $\left(\frac{B_i}{q_i^{1-\beta}}\right)$ with elasticity θ. Similarly, the larger the nominal local wage w_j, the larger local labor supply. On the other hand, the larger the commuting costs τ_{ij}, the less likely are workers to select the ij option. The attractiveness of a pair ij therefore depends on the characteristics of its residential location i, its workplace location j and the commuting between the two locations. In a closed city, the attractiveness of this pair is relative to the attractiveness of all other available pairs in the city, summarized by the endogenous aggregate variable W^w. The number of residents L_i^r and the number of workers L_j^f in locations i and j can be further defined as

$$L_i^r = \sum_j L_{ij}, \qquad L_j^f = \sum_i L_{ij}. \tag{12}$$

Commuting: τ_{ij} is the average iceberg commuting costs of workers living in i and working in j. It is a function of the route-level commuting costs $\tau_{ij,r}$ on all possible routes $r \in r_{ij}$ between i and j, such that

$$\tau_{ij} \equiv \left[\sum_{r \in \nabla_{ij}} \tau_{ij,r}^{-\rho}\right]^{-\frac{1}{\rho}}. \tag{13}$$

Average commuting costs between i and j are decreasing with route-level commuting costs $\tau_{ij,r}$, at elasticity ρ. The larger ρ, the larger the decrease in average commuting costs τ_{ij} in response to a decrease in the lowest route-level commuting costs $\tau_{ij,r}$. As $\rho \to \infty$, all commuters use the least cost path.

The total costs of traveling through route $r \in \nabla_{ij}$ is a function of the cost of commuting d_{kl} on all individual edges kl on route r, such that

$$\tau_{ij,r} = \prod_{kl \in r} d_{kl}. \tag{14}$$

Edge-level commuting costs $d_{kl}, kl \in r$ are an exponential function of travel time, as in Ahlfeldt et al. (2015) and are given by

$$d_{kl} = \exp\left(\kappa \cdot \frac{\bar{t}_{kl}}{(R_{kl})^{\xi}}\right), \tag{15}$$

where κ is the elasticity of commuting costs on commuting times, $\bar{t}_{kl}$ is a link-specific constant that depends on the length of link kl. ξ governs the elasticity of commuting time on the level of road infrastructure. The larger ξ, the more speed will increase (and commuting time decrease) in response to an increase in road infrastructure, corresponding to road improvements in this model. There is no congestion in this model and commuting costs are symmetrical, s.t. $d_{kl} = d_{lk}\ \forall kl$. Workers must use roads to commute, consistent with a setting like Kampala where there is no subway or commuter train.

Given workers' preferences for specific routes, governed by ρ, the equilibrium expected commuting cost τ_{ij} depends on all edge-level commuting costs, represented by a matrix. Following Allen and Arkolakis

(2022), I write this relationship as the matrix τ of dimensions $N \times N$, where the (i, j) elements is τ_{ij},

$$\tau = \left((\mathbf{I} - \mathbf{A})^{-1} \right)^{-\frac{1}{\rho}}, \tag{16}$$

where $\mathbf{A} \equiv \left[d_{kl}^{-\rho} \right]$ is a matrix with (k, l) element $d_{kl}^{-\rho}$, which is non-zero for neighboring locations only.

Production: Perfectly competitive firms in each location j produce a freely traded final good using labor L_j^f and floorspace H_j^f, such that

$$Y_j = A_j \left(\frac{L_j^f}{\alpha} \right)^{\alpha} \left(\frac{H_j^f}{1-\alpha} \right)^{1-\alpha} \quad \forall j,$$

where A_j is the exogenous productivity in location j. This production function, together with firms' zero profit condition imply that

$$w_j = A_j \left(\frac{\alpha}{1-\alpha} \frac{H_j^f}{L_j^f} \right)^{1-\alpha} \quad \forall j. \tag{17}$$

$$q_j^f = A_j \left(\frac{1-\alpha}{\alpha} \frac{L_j^f}{H_j^f} \right)^{\alpha} \quad \forall j. \tag{18}$$

Private floorspace Market Clearing: The residential floorspace market is competitive so the equilibrium rental rates q_i^r equate the supply of residential floorspace H_i^r in each location to the demand for residential land pinned down by equation (9). In equilibrium, all workers spend a fraction $(1 - \beta)$ of their income on housing and

$$q_i^r = (1 - \beta) \frac{\sum_j L_{ij} \frac{w_j}{\tau_{ij}}}{H_i^r}. \tag{19}$$

Analogously, the business land market is competitive so the equilibrium rental rates q_i^f equate the supply of business H_i^f floorspace in each location to the demand for business floorspace pinned down by equation (18).

5.2 The Policy: Road Widening

A policy increasing road width between neighboring locations k and l by ΔR_{kl} has the following consequences on the economy, with baseline prices and quantities indexed by ^{0}.

Lower Commuting Costs: $\Delta R_{kl} > 0$ decreases commuting costs along that link d_{kl} and subsequently the average commuting costs between all pairs of locations τ_{ij}.

Lower Private Land Supply and floorspace: $\Delta R_{kl} > 0$ increases road area in each location k and l by $\Delta R_{kl} \frac{l_{kl}}{2}$ and reduces private land by the same amount, where I assume that each link is divided equally between the locations it connects. I assume that this land is taken from residential and business land in proportion to land use of each type at baseline, and that μ_i^r and μ_i^f convert land loss into residential and business floorspace losses. Therefore, $\Delta R_{kl} > 0$ decreases residential and business floorspace in location k by $(\Delta R_{kl}\frac{l_{kl}}{2})\gamma_i\mu_i$ and $(\Delta R_{kl}\frac{l_{kl}}{2})(1-\gamma_i)\mu_i$, respectively

Change in Owners' Rental Income: As commuting costs and both residential and business land supply decrease, owners' rental income changes through two channels. First, residential q_i^r and business q_i^f rental rates change, affecting owners both on the roadside and away from roads. Second, roadside owners in k and l experience a decrease in their floorspace, while owners away from roads do not.

Compensation to Roadside Owners: When widening takes land away from roadside owners in k and l, these owners may be compensated. I write the compensation rate as the baseline market rate of their land times the probability that they are in fact compensated,

$$p_m^{\text{land}} = \underbrace{\left(\gamma_m\mu_m^r q_m^{r,0} + (1-\gamma_m)\mu_m^f q_m^{f,0}\right)}_{\text{baseline land market rate}} \cdot \mathbf{P}_m, \qquad m \in \{k,l\}, \tag{20}$$

and consider three compensation rules, nested by the value of $\mathbf{P}_m \in [0,1]$:

1. **No compensation:** $\mathbf{P}_m = 0$.
2. **Market value:** $\mathbf{P}_m = 1$: every affected owner is paid the full baseline market rate.
3. **Claim-based:** $\mathbf{P}_m = \mathbf{P}_m(\text{claim})$: owners are paid only if they file a claim, which they do when the gain from claiming exceeds its tenure-specific cost.

Under the claim-based rule an owner is compensated if and only if she claims, so that $\mathbf{P}_m(\text{claim})$ is given by

$$\mathbf{P}_m(\text{claim}) = \Pr\left(a \underbrace{\frac{\sum_{n\in C_m}\Delta R_{mn}\frac{l_{mn}}{2}}{N_m^{or}}(\gamma_m\mu_m^r q_m^{r,0} + (1-\gamma_m)\mu_m^f q_m^{f,0})}_{\text{potential compensation from claim}} \geq \underbrace{f(Z_m)+\epsilon_m}_{\text{cost of claim}} \right), \qquad m \in \{k,l\}, \tag{21}$$

where ϵ_m is logistic with mean 0 and unit scale, $f(Z_m)$ is the deterministic claiming cost under tenure regime Z_m, and $a \geq 0$ means a larger potential payment raises the claiming probability.

A far-from-road owner's land is untouched, so their income moves only with rental rates, $l_m^{of} = q_m^r\frac{H_m^{r,0}}{N_m} + q_m^f\frac{H_m^{f,0}}{N_m}$. A roadside owner instead gives up a share of floorspace and may be compensated for it:

$$l_m^{or} = \underbrace{q_m^r\frac{H_m^{r,0}}{N_m} + q_m^f\frac{H_m^{f,0}}{N_m}}_{\text{baseline rental income}} - \underbrace{\left(q_m^r\gamma_m\mu_m^r + q_m^f(1-\gamma_m)\mu_m^f\right)\frac{\sum_n\Delta R_{mn}\frac{l_{mn}}{2}}{N_m^{or}}}_{\text{value of land lost}} + \underbrace{p_m^{\text{land}}\frac{\sum_n\Delta R_{mn}\frac{l_{mn}}{2}}{N_m^{or}}}_{\text{compensation}}, \tag{22}$$

where rental rates have adjusted to the new land supply and commuting costs, and roadside owners split lost land and compensation equally.[19]

[19] Equations (20)–(22) credit roadside owners with their *expected* compensation. Under the linear owner utility (main

5.3 Government's Problem

Conditional on the road infrastructure at baseline $\left\{R_{kl}^{0}\right\}$, the government chooses the road widening policy $\{\Delta R_{kl}\}$ to maximize the net present value (NPV) of a weighted average of residents' (workers and owners) expected welfare, net of costs. The government is utilitarian so each resident is assigned the same weight. In the status-quo policy, the government's construction expenditures can be up to $\bar{M}$ (equivalent to the studied Kampala policy) and the marginal cost of construction is 1 (funded externally), while land payments are funded domestically at marginal cost $\eta \geq 1$, the *shadow price of public funds*: the social cost of raising and delivering one dollar of compensation from scarce domestic revenue, compounding the deadweight loss of distortionary taxation, a scarcity (liquidity) premium on the constrained public budget, and leakage in delivery (Section 6.1). Costs are paid upfront, while benefits are accrued in each period, discounted at rate δ. The government's problem can be written as follows:

$$\max_{\{\Delta R_{kl}\}} \quad \mathbf{W}^{NPV} \equiv \frac{1}{\delta}\left[\underbrace{\sum_k \left[N_k^{or} W_k^{\text{or}}\left(\{\Delta R_{kl}\}\right) + N_k^{of} W_k^{\text{of}}\left(\{\Delta R_{kl}\}\right)\right]}_{\text{owners' welfare}} + \underbrace{\bar{L}\, W^{\text{w}}\left(\{\Delta R_{kl}\}\right)}_{\text{workers' welfare}} \quad \underbrace{- \eta\left(\sum_k p_k^{\text{land}} \sum_l \Delta R_{kl} \frac{l_{kl}}{2}\right)}_{\text{land payment at MC}=\eta}\right] - \underbrace{\sum_k \bar{p} \sum_l \Delta R_{kl} \frac{l_{kl}}{2}}_{\text{externally-funded construction at MC=1}} \tag{23}$$

s.t.

$$\bar{p} \cdot \sum_{kl} \frac{l_{kl} \Delta R_{kl}}{2} \leq \bar{M}, \qquad \text{[Limits on Construction]}$$

Equilibrium: Given the parameters $\left\{\beta, \theta, \kappa, \xi, a, \psi^o, \bar{C}^o, f(.), \rho, \alpha, \bar{L}, \eta, \bar{M}, \bar{p}, \delta\right\}$, location and link characteristics $\{B_i\}$, $\{A_i\}$, $\{H_i\}$, $\{\mu_i\}$, $\{\gamma_i\}$, $\{N_i^{or}\}$, $\left\{N_i^{of}\right\}$, $\{Z_i\}$, $\{\bar{t}_{kl}\}$, and baseline infrastructure $\{R_{kl}^0\}$, the equilibrium of the model is the set of prices $\{w_i\}\,\{q_i^r\}, \left\{q_i^f\right\}, \{p_i^{\text{land}}\}$ and quantities $\{L_{ij}\}$, $\{\mathbf{P}_i\}$, $\left\{H_{ij}^r\right\}$, $\left\{H_i^f\right\}$, $\{C_{ij}\}$, $\{\Delta R_{kl}\}$ such that (i) the government chooses $\{\Delta R_{kl}\}$ to maximize city residents' welfare ; (ii) the government's budget constraints are satisfied ; (iii) workers choose $\{i, j, r, C_{ij}, H_{ij}\}$ to maximize their utility ; (iv) owners consume all their income on the tradeable good and get compensated with probability $\mathbf{P}_i$; (v) residential and business floorspace markets clear ; (vi) the good market clears, and (vii) labor markets clear.

5.4 Optimal Policy

The road widening policy $\{\Delta R_{kl}\}$ that solves the government problem in equilibrium can be recovered by solving the Lagrangian of the government's problem, as is derived in Appendix F. The optimal policy $\{\Delta R_{kl}\}$ satisfies the following system of equations:

specification) this representative-owner aggregation is exact for welfare. Under the concave (log) utility of Appendix G.1 it abstracts from the risk premium an individual owner would attach to the compensation lottery so it *overstates* owners' welfare under the partial-compensation (claim-based) rule relative to its risk-adjusted value; because the land taken is a small fraction of owner wealth, this is second order.

$$\underbrace{\frac{l_{kl}}{2}\left(\eta(p_k^{\text{land}}+p_l^{\text{land}})+\delta\,\bar{p}(1+\lambda)\right)}_{\text{financial cost for the government}}-\underbrace{\phi_k^{H^r}\frac{\partial H_k^r}{\partial\Delta R_{kl}}+\phi_l^{H^r}\frac{\partial H_l^r}{\partial\Delta R_{kl}}+\phi_k^{H^f}\frac{\partial H_k^f}{\partial\Delta R_{kl}}+\phi_l^{H^f}\frac{\partial H_l^f}{\partial\Delta R_{kl}}}_{\text{loss of res and bus floorspace}}-\underbrace{\zeta_k^{l_k^{or}}\frac{\partial l_k^{or}}{\partial\Delta R_{kl}}+\zeta_l^{l_l^{or}}\frac{\partial l_l^{or}}{\partial\Delta R_{kl}}}_{\text{roadside owners' change in income}}$$

$$-\underbrace{\Phi_k^{p_k^{\text{land}}}\frac{\partial p_k^{\text{land}}}{\partial\Delta R_{kl}}+\Phi_l^{p_l^{\text{land}}}\frac{\partial p_l^{\text{land}}}{\partial\Delta R_{kl}}}_{\text{change in compensation rate}}=\underbrace{\frac{\kappa\xi\bar{t}_{kl}d_{kl}}{(R_{kl}^0+\Delta R_{kl})^{1+\xi}}(\psi_{kl}+\psi_{lk})}_{\text{speed gain benefits}}\quad\forall kl\quad(24)$$

The right-hand side is the per-period benefit of widening: the speed gain on link kl, valued by the shadow price of lower commuting costs $(\psi_{kl}+\psi_{lk})$. The left-hand side is its marginal cost, in four parts: (i) the financial cost — construction at $\bar{p}$ plus the fiscal burden of land payments, financed at η; (ii) the opportunity cost of the residential and business floorspace displaced, via shadow values ϕ^{H^r}, ϕ^{H^f}; (iii) the change in roadside owners' rental income, weighted by its shadow cost ζ; and (iv), under claim-based compensation, the change in the endogenous compensation rate, weighted by Φ. Appendix F derives each term.

5.5 Properties of the Optimal Policy

The optimal policy in Equation (24) balances each link's marginal benefit (the right-hand side: faster travel) against its marginal cost (the left-hand side: construction, the opportunity cost of the land it consumes, and the compensation paid to affected owners). I make two further assumptions to isolate the two main ways land shapes the optimal policy:

Assumption 1: *Market value compensation at the baseline market rate, independent of ΔR_{kl}:*

$$p_k^{\text{land}}=\gamma_k\mu_k^r q_k^{r,0}+(1-\gamma_k)\mu_k^f q_k^{f,0}.$$

Assumption 2: *Owners' utility is linear in their rental income ($\psi^o=0,\ \bar{C}^o=1$).*

Land first shapes the optimal policy through its opportunity cost, as stated in *Proposition 1* below.

Proposition 1: *Under Assumption 1, the welfare-maximizing road widening policy ΔR_{kl} is decreasing in the shadow value of residential ($\phi_k^{H^r}, \phi_l^{H^r}$) and business ($\phi_k^{H^f}, \phi_l^{H^f}$) floorspace at the endpoints k and l. A planner who ignores the opportunity cost of land use (wrongly assumes $\phi_H^r=\phi_H^f=0$) overinvests in locations with high opportunity cost of land use.*

I illustrate Proposition 1 in Appendix F.4 by solving the planner's problem in a simple monocentric economy on a 5×5 grid. A planner who optimizes on benefits alone, ignoring land use ($\phi_H^r=\phi_H^f=0$), concentrates widening on the central links that carry the most traffic; accounting for the opportunity cost of the land those roads consume shifts investment outward toward cheaper peripheral links (Appendix Figure A12, panels a–b). In a monocentric city, because the highest-benefit locations are also the most land-scarce (Appendix Figure A13), ignoring land use systematically overinvests at the center. This margin is largely absent from the quantitative spatial literature, which optimizes on benefits alone; by incorporating it throughout, I benchmark the rest of the analysis against a planner who already internalizes the opportunity cost of land, so that the distortions I isolate come solely from the cost of compensating owners for the land taken.

Second, compensating owners for that land carries a marginal social cost. Under market-value compensation (*Assumption 1*) and linear owner utility (*Assumption 2*), this cost reduces to a *fiscal wedge* $\equiv \frac{l_{kl}}{2}(\eta-1)\big(p_k^{\text{land}}+p_l^{\text{land}}\big)$: the government disburses η dollars for every dollar received by owners, so $(\eta-1)$ per dollar is pure deadweight loss. *Proposition 2* ensues:

> **Proposition 2:** *Under Assumptions 1 and 2, when the shadow price of public funds equals one ($\eta = 1$), the optimal road widening policy $\{\Delta R_{kl}\}$ is independent of the compensation rule p_k^{land} $\forall k$. When $\eta > 1$, the optimal policy is decreasing in the baseline land rates $(\gamma_k q_k^{r,0} + (1-\gamma_k) q_k^{f,0})$ at the endpoints k and l.*

Throughout, the *first best* is the $\eta = 1$ allocation: the map a utilitarian planner would choose if public funds were costless, holding the model's fundamentals and the construction envelope $\bar{M}$ fixed. Under linear owner utility it is independent of the compensation rule (Proposition 2) and coincides with the no-compensation allocation, so "$\eta = 1$," "no payment," and "first best" denote the same benchmark.[20]

The same simple simulated economy illustrates the impact of the fiscal wedge (Appendix Figure A12, panels c–d): at $\eta = 1$ market-value compensation is a costless transfer and the optimal map coincides with the no-compensation first-best, whereas at $\eta > 1$ investment contracts and tilts away from the expensive central links where the wedge bites hardest. Appendix F.4 reports further simulations, including alternative compensation rules.

Importantly, this fiscal wedge holds independently of the parameterization of owners' utility function. However, if *Assumption 1* is relaxed, alternative compensation rules may lead to different optimal policies R_{kl} even for $\eta = 1$ because compensation would also have distributional consequences accounted for by the central planner, as illustrated in Appendix G.1 for log owner utility.

Proposition 2 establishes that the wedge tilts the optimal map away from expensive land; the next result decomposes the welfare cost of that tilt and identifies what governs its size.

> **Proposition 3:** *Under Assumptions 1 and 2, the welfare cost of financing land at a shadow price of public funds $\eta > 1$, relative to the $\eta = 1$ first best, is the sum of (i) a* fiscal *cost equal to $(\eta - 1)$ times the land payments the planner makes, and (ii) an* allocative *cost from re-optimizing $\{\Delta R_{kl}\}$ under the wedge. The allocative cost is zero when land prices are spatially uniform.*

The two channels operate distinctively. The *fiscal* cost is a pure level loss: the government burns $(\eta-1)$ dollars for every dollar paid to owners, wherever it builds. The *allocative* cost arises only because the wedge is levied on a spatially heterogeneous price. Since the wedge raises a link's effective marginal cost in

[20] Under non-linear owner utility (Appendix G.1) the compensation rule matters even at $\eta = 1$ through distribution alone, and the first best is the welfare-maximizing no-compensation allocation. All "percent of first best" figures are measured against it.

proportion to its land price, it pushes investment toward cheap land. This reallocation is costly precisely when cheap land coincides with low benefits, that is, it is governed by the link-level covariance between the land price p_{kl}^{land} and the marginal benefit of widening MB_{kl} (the right-hand side of Equation 24). The allocative distortion vanishes if $\eta = 1$ or if $\text{Cov}(p^{\text{land}}, MB) = 0$, in which case the wedge is a uniform level shift that leaves the ranking of links unchanged. The allocative distortion is largest when public funds are costly and land is dearest where roads help most.

The *allocative* channel is the spatial-investment analog of the misallocation from wedges studied by Hsieh and Klenow (2009): a wedge on a heterogeneous price generates dispersion in effective marginal costs and pulls the allocation from the efficient one. What distinguishes the present case is that the benefits of widening are themselves spatially heterogeneous, so the relevant sufficient statistic is not the dispersion of the wedge alone but its *covariance* with the benefit: the distortion concentrates where costly land and high returns coincide.

6 Model Quantification and Net Returns to Realized Road Improvements

6.1 Estimating The Parameters of the Model

I use the data and variation from the reduced-form analysis to estimate key model elasticities: the elasticity of speed on road width ξ, the elasticity of commuting flows on commuting costs θ, the elasticity of commuting costs on commuting time κ, and the parameters governing $\mathbb{P}(\text{claim})$. I calibrate the remaining parameters on Kampala or using estimates from the literature.

6.1.1 Road Infrastructure Elasticity of Speed (ξ)

ξ governs the elasticity of traffic speed with respect to road infrastructure. I estimate ξ using the same Google Maps speed data and identification strategy as in Section 4.1, replacing the binary upgrade indicator with a continuous measure of road width: the number of lanes times 3.4 m (the length-weighted average lane width from OpenStreetMap), with upgraded roads gaining one additional lane upon completion, and running a log-log regression of speed on road width.

Because the upgrade effect concentrates at rush hour (Section 4.1), I estimate ξ separately by time of day, under both TWFE and the bias-robust Sun and Abraham (2021) estimator (Table A16). I calibrate it to the rush-hour, bias-robust estimate, $\xi = 0.367$ (s.e. 0.158; column 2) for the same two reasons as the speed effect: rush hour is when the commuting the model prices occurs, and Sun–Abraham is robust to the staggered-design bias. Consistent with that bias, the estimate is smaller than the TWFE elasticities (columns 1 and 3).

Estimated at rush hour, ξ captures both the free-flow and congestion-relief margins of widening. It nets out the short- and medium-run re-routing of traffic onto the improved links, so it is not a free-flow engineering elasticity but the speed change that survived drivers' adjustment: an ATT applied to its own estimation sample, and thus exactly the object relevant for evaluating the realized program. I apply it as a common structural elasticity across links in the counterfactuals. *Long-run* congestion feedback is not captured: as cheaper commuting leads workers to relocate and travel farther, network flow rises and may erode the speed gains, an effectively lower ξ (Duranton and Turner, 2011).[21]

Interpretation: a bundled effective-capacity elasticity. The estimated ξ is best read as the elasticity of speed to *effective* road capacity, and I code every completed upgrade as a one-lane increase in effective width. While projects were officially categorized as dualling (27% of projects), reconstruction (59%), and paving (44%), land was acquired for upgrades of all types and affected owners in my survey line all upgrade types alike. Moreover, a degraded carriageway operates well below its nominal capacity: where potholes force two-way traffic to alternate over a single passable lane, reconstruction restores a second usable lane, which is an effective-capacity gain even absent nominal widening. Because only 2 of the 19 switcher roads are dualling projects, I cannot separately identify nominal-width from surface channels. In other words, the estimated ξ captures their combined effect, and the ΔR_{kl} is read throughout as an effective-capacity increase. This capacity interpretation is supported by two features of the estimates. First, the effect is smaller and not consistently significant off-peak, unlike at rush hour — precisely when added capacity relieves congestion. Second, road capacity in Kampala is exceptionally scarce, placing the city on a steep segment of the speed–capacity relationship where the marginal return to capacity is high. Table A8 shows that Kampala devotes a smaller share of built-up land to roads than the world (21%) or even the Sub-Saharan (16%) average, has the narrowest post-1990 streets and the highest share of sub-4 m roads among the comparison cities, and almost no roads wide enough to carry fast through-traffic (Angel et al., 2016). This scarcity rationalizes an elasticity roughly three times the high-income benchmarks ($\xi = 0.10$ in the U.S., Fajgelbaum and Schaal 2020; $\xi = 0.13$ in Santiago, Bordeu 2024) as a sensible response to relieving an unusually binding capacity constraint, rather than as an artifact of the bundled treatment. In Appendix G.2 I show that at half the estimated elasticity ($\xi = \frac{1}{2}\hat{\xi}$) the paper's comparative conclusions hold and the misallocation from land payments under $\eta > 1$ strengthens at $\xi = \frac{1}{2}\hat{\xi}$ as a lower ξ raises the relative bite of land costs. The *level* of the realized net gain is more sensitive, resting on owners going largely uncompensated.

6.1.2 The semi-elasticity of commuting flows on commuting costs

I turn to the parameters governing how workers' location choice responds to increases in local road speed. In the short term, fixing residence and workplace, workers will be more likely to choose a route using these faster roads, consistent with evidence from Google Maps that upgraded roads are more likely to be used by commuters (Table A13). In equilibrium, workers will change their residence and workplace, affecting commuting flows L_{ij}. The magnitude of this response is governed by θ, the semi-elasticity of commuting flows on commuting costs, and κ, the elasticity of commuting costs on commuting times: the larger $\theta\kappa$, the larger the impacts of road improvements on workers' welfare.

I estimate $\theta\kappa$ using a Poisson pseudo-maximum-likelihood gravity equation on data from a local ride-hailing company, which records origin-destination commuting flows at the parish level for a random sample of weeks between 2019 and 2024. Regressing flows on predicted commuting times with origin × year-month

[21] Estimating the congestion feedback directly would require link-level traffic counts paired with speeds on the same links, which are not jointly available for Kampala.

and destination × year-month fixed effects yields $\hat{\theta\kappa} = 0.0319$ (s.e. 0.0039). More details on the data and empirical specifications can be found in Appendix F.5.1.

To take $\theta\kappa$ to the model I separate its two components. I identify κ from mode-choice variation in the 2016/2017 Travel Habit Survey, obtaining $\hat{\kappa} = 0.0089$ and hence $\theta = 4.36$ (Appendix F.5.1). This estimate is imprecise ($p = 0.251$, $N = 142$) but is of the same order of magnitudes as existing estimates, like $\kappa = 0.01$ for Berlin (Ahlfeldt et al., 2015) and $\kappa = 0.013$ for Bogota (Tsivanidis, 2026). In Appendix G.2 I show that the results are robust to taking $\kappa = 0.01$ from the literature instead.

6.1.3 Parameters Governing Compensation

Of the three compensation rules—no compensation ($\mathbf{P}_i = 0$), market value ($\mathbf{P}_i = 1$), and claim-based ($\mathbf{P}_i = \mathbf{P}_i(\text{claim})$)—only the last requires estimation. I estimate the empirical counterpart of Equation (21):

$$\mathbf{P}_{o,m}(\text{claim}) = \Pr\big(\alpha\, q^0_m\, \Delta H_{o,m} \geq \mathbf{Z}'_{o,m}\mu_Z + \mathbf{X}'_o\mu_X + \epsilon_{o,m}\big)\,, \tag{25}$$

where q^0_m is the pre-policy market rate in parish m, $\Delta H_{o,m}$ the land taken from owner o, $\mathbf{Z}_{o,m}$ the tenure regime (freehold and public, with mailo the reference), $\mathbf{X}_o$ owner and parish controls, and $\epsilon_{o,m}$ standard logistic. Estimated by maximum likelihood (Table 1, columns 7-8), the parameters reproduce a robust pattern even as the point estimates move across the two estimation samples: claiming rises with the potential compensation ($\alpha > 0$, the sensitivity to the value of affected land), and freehold and public-land owners claim more than mailo owners (the reference land tenure). This ordering is robust across different sets of owner-level controls, parish-level controls, alternative outcome definitions and sample restrictions, as displayed on Appendix Figure A7.

To take them to the model, I predict $\mathbf{P}_i(\text{claim})$ for each parish from its tenure shares, a wealth indicator (2014 Census), and parish-level controls, holding the remaining owner-level variables (including the market value of the taking) at their means; these predictions are highly correlated across specifications (Spearman 0.85–0.96; Appendix Figure A6), so the spatial ranking reflects tenure composition. $\mathbf{P}_m(\text{claim})$ is *endogenous to the policy*: by Equation (21) it rises with the potential compensation, and hence with the amount of land the upgrade takes, ΔR_m. The implied parish-level expected land-acquisition cost per m^2 of land taken is mapped across the alternative compensation definitions in Appendix Figure A15.

When evaluating Kampala's *realized* upgrades (2018–2026) rather than a single counterfactual rule, the government applied a mix of the two paid rules — offering market value on some projects and relying on claims on others (Section 2). I observe both margins (Table 1: the offer/receipt margins in columns 3-6, the claim logit in columns 7–8). Because the prevailing rule shifted over time and is not recorded for every road, I do not impose a single rule. Instead I bound the compensation actually paid between a lower bound $\mathbf{P}^{\text{offered}}_m$ (the share of owners offered compensation) and an upper bound $\mathbf{P}^{\text{ever}}_m$ (the share ever compensated), each predicted at the parish level from a linear probability model (columns 5-6 and 1-2, respectively).

6.1.4 Calibrated Parameters

I calibrate the remaining parameters to Kampala, as listed in Table 3 together with their sources. Owners' utility is linear in income in the main specification ($\psi^o = 0$) and I solve the model under non-linear owner utility in Appendix G.1.

The number of landowners is recovered from baseline rents under a free-subdivision assumption (Appendix F.5.5), which implies a baseline owner income about 2.44× the average worker wage. This income gap is the source of the redistributive cost of compensation under non-linear owner utility.

Discount rate δ: Construction and land expenditures are made upfront (a stock) while the benefits arise over time (a flow), so a discount rate is needed to express both in the same units. I take $\delta = 0.06$ as my central social discount rate, following World Bank (2016)'s guidance of roughly twice the expected long-run rate of per-capita consumption growth, or 6% given Uganda's recent per-capita growth of about 3% (3.2% in 2024).[22] As a plausible high rate I use $\delta = 0.12$, which reflects sustained catch-up growth (about 6% per capita) and is the rate used in the economic appraisal of the program itself: the KIIDP-2 Project Appraisal Document evaluates the investment at a 12% discount rate and treats a 12% economic rate of return as the viability threshold (World Bank, 2014). I report results across the full range $\delta \in [0.03, 0.15]$.

The shadow price of public funds η: I take η to be the *shadow price of public funds*: the social cost, in units of general government revenue, of delivering one dollar of compensation to a landowner. η is broader than the marginal cost of public funds (MCPF) from distortionary taxation alone and is instead defined as the local value of the multiplier on a cash-constrained government's budget. It compounds two main forces. First, the deadweight loss of taxation: MCPF estimates for Africa average between 1.2 and 1.4 (Auriol and Warlters, 2012; Dahlby, 2008). Second, as evidenced by the government's chronic cash constraints and its reliance on concessional external finance for construction, a scarcity (liquidity) premium on domestic public capital: low-income governments face binding borrowing constraints and much tighter fiscal space than high-income governments (Kaminsky et al., 2004; Ilzetzki and Végh, 2008; Frankel et al., 2013; Reinhart et al., 2003), so a dollar spent on land is a dollar of scarce domestic revenue diverted from high-return public uses.[23] Additionally, η may be inflated by leakage in delivery: weak public-investment management dissipates a large share of public spending before it reaches its purpose (International Monetary Fund, 2015; Reinikka and Svensson, 2004), and Kampala's local fiscal capacity is especially limited (Manwaring and Regan, 2024). While captured funds are partly transferred rather than lost, the rent-seeking they attract and the monitoring they require are themselves real resource costs (Krueger, 1974; Olken, 2007).

Together these imply a value of η well above the taxation-only benchmark and I report results across the full range $\eta \in [1, 8]$, with $\eta = 4$ as a central value motivated by the government's own decisions

[22] Source: World Bank.

[23] Uganda's government pays a weighted-average 14.4% on its domestic debt against 2.4% on its predominantly concessional external debt (Ministry of Finance, Planning and Economic Development, 2024). While this nominal gap, about six-fold, also embeds local-currency inflation and exchange-rate risk and therefore overstates the pure scarcity premium, I read it as consistent evidence that domestic public funds are considerably dearer than the funds financing construction, consistent with η much greater than 1.

(Section 7).[24]

6.2 Kampala Geography and Baseline

Locations: I split Kampala in 68 locations, corresponding to the city's parishes with very small parishes merged to their neighbor as represented on Figure A14.[25]

Nodes and edges: I define nodes in the network as the centroids of the city's 68 parishes. Nodes in contiguous parishes are connected by edges, corresponding to representative roads. The length of each edge is the centroid-to-centroid distance multiplied by a global detour factor (the median ratio of Google Maps trip distance to centroid distance across all matched edges, approximately 1.755), so that l_{kl} reflects realistic road distances rather than straight-line distances. To measure road infrastructure on each edge, I map each road of the underlying OpenStreetMap road network to the closest edge and compute the average road width along this edge, weighted by the length of the underlying roads.

Baseline (Pre-policy) Residential Population $\left\{L_i^{r0}\right\}$ and Employment $\left\{L_i^{f0}\right\}$ come from the 2014 Population and Housing Census restricted to working-age population, and 2019 Census of Business Establishments from the Uganda Bureau of Statistics (UBOS), respectively.[26] As evidenced by panels (a) and (b) of Figure A14, which display the distribution of residents and employment throughout Kampala at baseline, Kampala is a relatively monocentric city, with employment being concentrated near the central business district (CBD) at the city center.

Infrastructure-free travel time $\bar{t}_{kl} = t_{kl}^0 \times (R_{kl}^0)^{\xi}$ is recovered from equation (15), where t_{kl}^0 is the observed baseline travel time on edge kl, averaged over 107,956 pre-policy (2019) Google Maps trip instances along edges connecting neighboring parishes (Akbar et al., 2023), and R_{kl}^0 is the length-weighted average width of all major OSM 2019 road segments assigned to the link. To ensure $\bar{t}_{kl}$ varies smoothly across the network, I predict it from link length, distance to the CBD, and local road length rather than mixing observed and imputed values (Appendix F.5.3). Equipped with $\bar{t}_{kl}$, R_{kl}^0, ξ and κ, I recover τ_{ij}^0 from Equations (13), (14) and (15).

Baseline rental rates are recovered from the broker-survey transactions in several steps: I convert sale prices to a per-m^2 floorspace rate, strip out the reduced-form effect of upgrades to obtain pre-policy values, predict a parish-year rate surface (imputing parishes without transactions from their neighbors), and convert sales to rents using property-type cap rates. Appendix C details each step and Figure A14

[24] η, as a shadow price, is endogenous to the government's *aggregate* fiscal problem and may rise with total public spending. I hold it fixed in the analysis because the program's land bill is small relative to the government's overall budget, so I treat the program as a price-taker in the market for public funds. This means that in the most aggressive counterfactuals (city-wide market-value compensation), where the land bill is large enough that η would rise with it, holding η fixed implies a lower bound on the relative fiscal penalty and subsequent misallocation.

[25] Kampala includes 96 distinct parishes in total, including some very small ones. For computational purposes, I merge parishes with area less than 1 km^2 with their smallest neighbor.

[26] The 2014 Population census is the most recent available before the start of the policy. Total population is rescaled to match the 2019 KCCA Statistical Abstract estimate of 1.75 million residents (1.5 million in the 2014 Census); the spatial distribution across parishes is held at 2014 Census shares. Employment count is rescaled s.t. $\sum_i L_i^{f0} = \sum_i L_i^{r0}$.

panels (c) and (d) display the resulting residential and business rental rates by location at baseline. Market-value compensation pays owners the per m^2 sale value of the land taken, valued per period at δ.

Baseline Road Land $\{H_i^{pu0}\}$ and Private Land $\{H_i^{pr0}\}$. Land used by road H_i^{pu0} in each parish is measured as the ground area occupied by roads in the OpenStreetMap 2019 road data, for all road segments clipped to the parish. Private land area H_i^{pr0} is recovered by subtracting H_i^{pu0} from the parish total area. The share of residential γ_i and firm $(1-\gamma_i)$ land in each parish comes from the GHS (Global Human Settlement) satellite land-use classification.[27] Residential and business land areas are obtained by multiplying H_i^{pr0} by γ_i and $(1-\gamma_i)$, respectively.

Floorspace Production. Agents do not consume raw land but *floorspace*: residential floorspace H_i^R and business floorspace H_j^F. In each location, floorspace is produced from privately owned ground area H_i^{pr0} through a linear, location-specific technology whose productivity μ_i captures how intensively land is developed:

$$H_i^R = \mu_i^R \gamma_i H_i^{pr0}, \qquad H_j^F = \mu_j^F (1-\gamma_j) H_j^{pr0}. \tag{26}$$

I recover the productivity terms μ_i^R, μ_j^F residually from floorspace market clearing, so that by construction the model reproduces observed broker rents at baseline (Appendix F.5.4). In counterfactuals I hold the floorspace technology μ fixed: development density does not respond to the road policy, so floorspace adjusts only through the land roads take.[28] The average pre-policy monthly floorspace rental rate is 0.69 USD/m^2 for residential and 1.32 USD/m^2 for business premises.

Recovering fundamentals and wages: Given $\{L_i^{r0}, L_j^{f0}, q_i^{R0}, q_j^{F0}, \tau_{ij}^0\}$, I recover location-specific productivity A_j and residual amenities B_i by inverting the model following Ahlfeldt et al. (2015): I solve numerically for the $2N$ fundamentals (A_j, B_i) that match the model-predicted spatial distributions of residents and workers to their observed counterparts.[29] I rescale A_j so that the average wage matches the 2019 mean monthly wage of \$147 (KCCA 2019), and normalize B_i so that a \$1/month lump-sum transfer changes worker welfare by exactly one unit, which pins down the welfare units used in the cost-benefit analysis.[30] From the firm first-order condition, baseline wages are then pinned down by productivity and observed business rents, $w_j^0 = \alpha\, A_j^{1/\alpha} \left(\frac{1-\alpha}{q_j^{f0}}\right)^{(1-\alpha)/\alpha}$. Figure A14 panels (e), (f) and (h) show the maps of residential amenities, productive amenities, and baseline wages.

7 Net Returns to the Realized Road Upgrades

I evaluate the welfare effects of the externally funded road upgrades in Kampala between 2018 and 2026 by solving for the new spatial equilibrium given these realized improvements. I map each road project to the closest link in the model's aggregated network, assuming an average width increase of 3.4 meters (one

[27] γ_i is capped at 0.95 so that every parish retains at least 5% of business land. Unclassified land is assigned to residential and firm uses proportionally to classified land.

[28] This is the standard assumption of exogenous floorspace supply (Allen and Li, 2015; Monte et al., 2018; Tsivanidis, 2026).

[29] The system is solved using Newton's method.

[30] Because worker indirect utility is homogeneous of degree one in income, this renders W^{w} money-metric, so ΔW^{w} is the equivalent variation—additive to owners' income and net of costs in Equation (23).

additional lane), which assigns upgrades to 92 undirected links, adding a total road area of 387,946 m^2. The corresponding change in link width is displayed on Appendix Figure A16. I present results for both a lower and an upper bound on compensation, as described in Section 6.1.3. Road improvement benefits come through increased speed and rely on the width and commuting elasticities (ξ, $\theta\kappa$) estimated in Section 6. I solve the model for the resulting equilibrium changes in location, commuting times, property values and welfare.

This section makes three points. First, the upgrades generated large net welfare gains beyond local changes in property values. Second, the marginal net gains from road upgrades are spatially uneven and I uncover a positive correlation between the realized allocation and links for which the model estimates high marginal benefits from the upgrade and low land costs. Third, the size of the net gain depends on how land was paid for: holding the realized map fixed, a more generous compensation rule lowers the net return, and market-value compensation could have led to negative net welfare gains of the projects once costs are accounted for, an issue that the implemented compensation approach sidesteps.

7.1 Changes in Commuting Time, Property Values and Gross Welfare

The realized upgrades benefited workers by decreasing commuting costs: average commuting time fell by 3.64% in general equilibrium, and citywide residential and business rents rose by 0.84% and 0.33%, respectively, raising landowners' rental income.

In aggregate, these benefits translated into residents' and owners' welfare gains of USD 235.7 million, when evaluated at a 6% market discount rate. Appendix G.2 shows how these returns change at a more conservative elasticity ($\xi = \frac{1}{2}\hat{\xi}$): the realized net gain falls and comes to rest on owners going largely uncompensated, while the comparative conclusions sharpen.

These citywide gross benefits are about 1.6 times the increase in the value of properties lining the upgraded roads (evaluated at USD 147.1 million from the broker appraisal-derived reduced-form estimates), illustrating the need for a model to capture citywide impacts of transportation infrastructure improvements.[31]

7.2 Spatial Heterogeneity and Alignment with Model-Predicted Returns

These averages mask heterogeneity along two dimensions. First, because of the model is one of a closed city, rental income did not increase in all locations: as workers moved out of some locations, rental rates there fell. Appendix Figure A17 displays the change in total rental income by parish and highlights this pattern: 39 locations gained while 29 lost. Second, within a given parish, not all owners benefited from higher rental rates, as some close-to-road owners also lost land. I estimate that the average close-to-road owner's net change in rental income is -3.56%, against 0.36% for the average away-from-road owner.

Despite this heterogeneity, the realized upgrades line up reasonably well with where the model says the returns were highest. I regress realized upgrade outcomes on the model's (standardized) marginal benefits and costs, controlling for link length (Table 4). I find that a link with larger marginal benefits, measured

[31] The citywide welfare number is conservative on two counts. First, it comes from a closed-city model that holds Kampala's total population fixed, so it omits the additional gains an open city would realize as improved accessibility raises the city's attractiveness and draws in new residents and firms. Second, it prices only the access benefits of road improvements, excluding the local amenity gains (reduced dust and flooding, improved safety) that the broker appraisal reflects.

by the RHS of Equation 24, is more likely to have been upgraded (0.079, s.e. 0.042 in column 1). Similar patterns can be observed when benefits are measured by changes in residents' welfare (Appendix Table A21) Conversely, links with a higher market value of affected land, a higher predicted share of owners claiming compensation (columns 2 and 6), and a higher share of freehold land (columns 3-4 and 7-8) are less likely to have been upgraded. This mirrors the program's official selection process: its appraisal document lists the availability of funds to compensate affected owners among the criteria governing which sub-projects proceed (World Bank, 2014).

The signs match the model and hold across specifications (benefits positive, land costs negative), though significance is not uniform: the public-land coefficient is never significant ($p > 0.17$). I read the realized map as consistent with the model, not a formal test of it. The imperfect fit reflects three features of the setting: the model aggregates individual roads to the nearest parish-to-parish link, so a project need not map to one link; it allows continuous widening while the government's choice is closer to discrete (*which* road to upgrade), weakening the intensive-margin results (columns 5–8); and the government weighs political and distributional considerations the model omits.

7.3 Compensation Rules Affect the Net Policy Returns

The size of these net gains, however, depended on how the government compensated the owners whose land it took. Holding the realized map fixed, the more generous the compensation rule, I show that the lower the net return from the policy, especially when (i) the shadow price of public funds η is large, so that any compensation carries a fiscal loss, and (ii) the discount rate δ is large, since benefits accrue over time.

Figure 6 displays the net present value (NPV) of the realized upgrades across values of $\eta \in [1.0, 8.0]$ and discount rates $\delta \in [3\%, 15\%]$ under the three modeled compensation rules on the top row (no payment, claim-based, market-value) and for the recovered compensation rates (lower and upper bounds) for the realized upgrades on the bottom row.

The more generous the compensation rule, the lower the net returns of the policy, especially for high values of η. The compensation rule can even switch the sign of the net returns of the policy. For example, at $\delta = 0.06$ and $\eta = 6$, market-value compensation turned the NPV negative. On the bottom row, I display the NPV net returns of the policy for the lower and upper bounds of the compensation rates recovered for Kampala from the owner survey (Section 3) and I estimate that for $\delta = 0.06$ and $\eta = 4.0$, the policy's net returns were between USD 94 million and USD 118 million.[32]

7.4 What the Government's Compensation Decision Suggests about η

The government's own decision partway through the project to move away from the market-value compensation it had initially committed to, in favor of the claim-based rule, is consistent with a high shadow

[32] Because the main specification uses linear owner utility (marginal utility of income $\equiv 1$), the fall in net NPV when moving from no compensation to market value reflects *only* the fiscal wedge $(\eta - 1)$ on land payments, not redistribution. Appendix Figure A23 repeats the exercise under log owner utility, which additionally values a dollar paid to (richer-than-average) owners at less than a dollar. Comparing the two, about 84% of the no-compensation-to-market-value decline in net NPV at $\eta = 4$ and $\delta = 0.06$ comes from the cost of public funds, with the remainder from the distributional cost of transferring to owners.

price of public funds. Its stated motivation — that it could not afford to compensate every owner — is itself a statement about the scarcity (liquidity) component of η, so I read the switch as corroborating a value in the upper part of the range motivated in Section 6 rather than as an independent estimate of it.

As illustrated on Figure 6, holding the realized map fixed, market-value compensation turns the project's net return negative once η exceeds 4.8 at a 6% discount rate (and 2.3 at 12%), so the switch is consistent with the government perceiving an η at least this high, under the maintained assumption that its decision tracks the model's net-return calculation.

This logic is consistent with η lying in the middle of the literature range. I therefore adopt $\eta = 4$ as the central value and $\eta = 6$ as an upper case in the optimal-policy analysis that follows, while continuing to report results across the full range $\eta \in [1, 8]$.

7.5 Distributional Consequences of the Compensation Rules

Importantly, while compensating all affected owners at market value can make the net returns of the policy negative, it does significantly improve the welfare of these owners. Appendix Figure A18 plots the welfare change of close-to-road owners against the amount of land taken per owner, under the different compensation rules, with each dot a representative close-to-road owner in an affected parish. In the absence of compensation (blue), 0% of affected (close-to-road) owners benefit from the change in rental income alone. Under market-value compensation to all owners, 33% of affected owners would have benefited.

8 Optimal Allocation of Road Improvements

The realized upgrades were worth building and were positively selected, but what the government paid for land was already first-order for their returns. Because compensation amounts are heterogeneous across the city, in addition to this *level effect*, payments for land also have an *allocative* impact: they can reshape where investment should go, and the cost of public funds governs how large that reshaping is. In this section I study the decision of a planner choosing where to allocate the same construction budget the donors delivered for road widening, so as to maximize residents' net welfare gains from road improvements — and how that allocation depends on compensation rules and funding.[33]

I do so in three steps. I first characterize the marginal net returns to widening one link at the time to highlight the trade-offs introduced by accounting for land costs. I then solve for the welfare maximizing (*optimal*) allocation under the eminent domain with market value compensation regime and show that it is shaped by the cost of public funds η (and the discount rate δ). I solve for the optimal policy as a single nonlinear program (planner's objective subject to the equilibrium conditions), using multistart to address non-convexity; see Appendix F.2 for details. Finally, I study two policy levers in turn: the *compensation rule*, which sets the level and spatial pattern of land payments, and the *financing rule*, which sets how costly those payments are to fund.

[33]Throughout, I hold the maximum total road upgrade to what was actually delivered by external donors, so that each counterfactual reallocates at most the same aggregate construction across the network. Fixing total kilometers shuts down the channel by which a cheaper rule simply lets the government build *more* and isolates the channel by which a rule changes *where* it builds.

8.1 Marginal Costs, Benefits, and Net Returns of Road Widening

To characterize how the marginal returns to road improvements vary across Kampala's geography, I conduct a link-level experiment: for each of the 172 undirected links in the commuting network between neighboring parishes, I widen the link by 3.4 meters (one lane) and solve for the new spatial equilibrium. I then evaluate the terms of Equation (24)—which summarizes the marginal costs and benefits of road widening—for each link, normalizing by link length. This marginal net return, marginal benefit minus marginal cost, is the object a planner equalizes across links when choosing the optimal map.

Marginal benefits and costs. Figure 7 maps the marginal benefits and costs of each link upgrade. Marginal benefits come from the reduction in commuting costs along link kl (panel a). They are highest for links near the CBD and decline with distance, consistent with central links serving a disproportionate share of commuting flows. The fiscal cost of land under market-value compensation (panel b) is also highest at the center.[34] The rank correlation between the marginal benefit and the land cost of a link upgrade (per meter) is 0.68 (Appendix Figure A19).[35] This correlation is the covariance identified by Proposition 3 as the driver of the misallocation channel: land is dearest exactly where roads help most, so the fiscal wedge distorts the map most where the returns are highest.

The fiscal burden of land payments can lead to net losses. This positive correlation is the central tension for a planner for whom land payments are not mere transfers but carry a loss ($\eta > 1$). It raises a natural question: can the fiscal burden of compensating owners at market value outweigh the benefits of widening? I compare the net marginal return under no compensation (panel c) with that under market-value compensation at $\eta = 4$ (panel d). In the absence of land payments, the links with the largest net returns (dark blue) ring the CBD, while upgrades *in* the CBD have null or slightly negative net returns because floorspace is so valuable for production there. Under market-value compensation, upgrades at and around the CBD have negative net returns, because this is also where the market value of land is highest. Overall, 34% of links (58 out of 170) that are net positive under no compensation flip to net negative under full market-value compensation.

8.2 The Cost of Public Funds Shapes the Optimal Map

Land acquisition costs shape the optimal allocation through two channels, both governed by η. The welfare-relevant wedge on a land payment, defined in Section 5, is $(\eta - 1)\, p_k^{\text{land}}$. Through a *level* effect, a higher η makes all widening more expensive and can shrink total investment; through a *spatial* (substitution) effect, the wedge grows in proportion to the local land price, so it falls unevenly and tilts investment away from expensive locations. At $\eta = 1$ the wedge is zero everywhere, however dispersed land prices are: payments are a pure transfer, and under linear owner utility the optimal map coincides with the no-compensation, first-best allocation, depending on the compensation rule only up to distribution. At $\eta > 1$ it tilts away from high-land-value links, the more so the larger η, as the deadweight on the expensive central links

[34] Construction costs per square meter do not vary across links. Under the claim-based compensation rule an additional marginal cost is the change in the endogenous compensation rate, which I abstract from here.

[35] The full marginal cost of widening link kl also comprises the opportunity cost of displaced residential and business floorspace and the welfare loss for roadside owners whose land is taken.

grows relative to the cheap peripheral ones (Proposition 3). This is the spatial analogue of a broad class of misallocation results (Hsieh and Klenow, 2009): a wedge on a heterogeneous price disperses the effective marginal cost of an input, here land, and that dispersion pulls the allocation away from the efficient one.

Appendix Figure A20 traces the optimal allocation as η rises from 1 to 8 at $\delta = 0.06$. At $\eta = 1$ the planner builds the first-best map, concentrated in the center. As η rises, two things happen. The planner builds less: at $\eta = 8$ it adds only 78% of the potential road area (about 388,000 m^2 at $\eta = 1$). And investment shifts unevenly, since the wedge scales with the local land rate: at moderate η (e.g. $\eta = 2$) total construction is unchanged but reallocated, with some links widened more than the first best (blue) and others less (red) in the bottom row.[36]

Thus η is not a knob on the *size* of the program alone; it bends its *shape*. Everything that follows operates through this wedge: the compensation rule affects the allocation (beyond distribution) only when $\eta > 1$, and the financing rule matters precisely because it changes which costs the wedge applies to.

8.3 Policy 1: Alternative Compensation Rules

The first lever a government has control over is the compensation rule. I consider three: no compensation, which shifts the entire land burden onto owners and corresponds to uncompensated expropriation, still widespread in many parts of the world (Tagliarino, 2017); market-value compensation, the benchmark eminent domain requires and the rule Kampala applied to the first roads in the program; and claim-based compensation, under which only owners who formally document and pursue their claim are paid, the rule the government moved toward for the latter part of the project. The three rules differ both in how much the government pays in aggregate and in how the payment is distributed across space. In turn, they reshape the optimal map whenever public funds are costly.

Moving between compensation rules operates through the same two channels as the wedge itself (Proposition 5.5): an *income* effect that changes the total land bill, and a *substitution* effect that changes how the wedge covaries with benefits across space (Appendix Figure A22 decomposes both for each link). Going from market value to the claim-based rule, the income effect is positive, since fewer than all owners claim and the total bill falls. The substitution effect is in principle ambiguous at central locations, which combine a higher claim probability (raising cost) with a higher market value (so more is saved when owners do not claim); quantitatively, in Kampala the latter dominates, so the claim rule corrects the wedge most where land is dearest, which is also where marginal benefits are highest (Figure 7). The claim-based rule therefore moves the allocation closer to the first best than market value does.

Figure 8 shows the optimal allocation under each rule at $\eta = 4$ and $\delta = 0.06$. Relative to no compensation (*no pay*), market value does not lower widening uniformly: links through expensive residential or business areas (Figure A14, panels c–d) receive fewer upgrades while those away from them receive more, and at higher η the total also falls (302,000 m^2 of added width at $\eta = 8$ against 388,000 m^2 under no compensation). The claim-based rule reverses much of this avoidance: expensive central areas are no longer

[36] Appendix Figure A21 repeats the exercise across the discount rate δ, which converts per-period resident welfare into an NPV. A higher δ does not affect the wedge (payments are already a stock) but lowers the NPV of benefits, letting the wedge bite earlier: the larger δ, the lower the incentive to widen (230,000 m^2 at $\delta = 0.15$ against 388,000 m^2 at $\delta = 0.06$ for $\eta = 4$), and the larger the role of distortions from a high η.

shunned, because that is where the funds saved from unpaid owners are largest. The gains from market value to claim-based compensation (NPV of USD 685 million versus USD 874 million) come both from lower fiscal losses and from an allocation closer to the first best (no pay, or $\eta = 1$).

Moving to a claim-based regime substantially changes the geography of investments: expensive business and residential areas are not avoided anymore despite relatively higher claim probability (Appendix Figure A22) because that is also where the saved funds from not paying at market value are the largest. The net gains from moving from market value compensation (NPV = USD 685 million) to claim-based compensation (NPV = USD 874 million) therefore both come from lower fiscal losses and from reaching an allocation closer to the first best (no pay, or $\eta = 1$).[37]

Property rights and the second best. Under the claim-based rule the government pays only owners who document a claim, and claiming is harder on Kampala's less formal tenure, mailo land, where ownership is layered and titles are contested. Weaker property rights therefore lower the share of owners who must be paid, cheapen the land the planner acquires, and, because the fiscal wedge bites hardest on expensive central land, pull the optimal allocation closer to the first best. This is a second-best result: when public funds are costly ($\eta > 1$), the imperfect enforcement of compensation offsets rather than compounds the distortion from the wedge. Importantly, owners who cede land uncompensated bear a real cost, which the utilitarian objective nets against the citywide gains. Governments weighing owners' losses more heavily would find it less attractive to build in locations where owners are less likely to be compensated.

Crucially, this ranking does not hinge on how owners' welfare is weighted. Under the linear, money-metric benchmark in the main specification, a dollar to an owner is valued exactly as a dollar to a worker, so the advantage of the no-compensation and claim-based rules over market value is driven purely by the fiscal wedge (the deadweight loss of raising public funds to finance compensation), with owners' uncompensated losses counted at full weight. Allowing diminishing marginal utility of owner income (log utility; Appendix G.1) only *reinforces* the ranking, because owners are richer than average and market-value compensation then carries an additional regressive cost and compensation lowers net returns even at $\eta = 1$. With log owner utility, the fiscal wedge accounts for 84% of the no-compensation-to-market-value welfare decline at $\eta = 4$, $\delta = 0.06$, while the remainder is distributional.[38] What the analysis does *not* price is the dynamic cost of insecure property rights (e.g. the investment and holdup effects of owners anticipating uncompensated takings (Besley and Ghatak, 2010; Hornbeck, 2010)) so it speaks to the static allocation of a given program, not to the desirability of weak property rights as such.

[37] The claim-rule advantage does not hinge on the imprecisely estimated sensitivity of claiming to the value at stake (α; Table 1, columns 7–8), through which $\mathbf{P}$(claim) responds to how much the planner builds. Freezing $\mathbf{P}$(claim) at its realized, one-lane predicted level leaves the share of the first best the rule captures essentially unchanged — 88% against 89% (Table A23). The claim-rule advantage thus reflects the spatial *level* of claiming across the city's tenure regimes, not the way claiming responds to the policy.

[38] The model finances compensation from general domestic public funds at cost η. The main (linear) result is invariant to the incidence as with equal weights the transfer nets out, leaving only the fiscal wedge ($\eta - 1$). Only the log-utility distributional reinforcement depends on it: broad-based financing (the de jure arrangement, land acquisition being the central government's responsibility) is regressive and reinforces the ranking, whereas a burden concentrated on owners (e.g. if pushed onto local property-tax revenue) would be progressive among owners and could offset it for low η. This incidence is hard to trace; the headline does not depend on it.

8.4 Policy 2: Alternative Financing Rules

That the allocative effects above all stem from the cost of public funds points to a second lever that operates directly on η: the restriction confining donor funds to construction, which forces land payments onto costly domestic revenue. Donors justify this restriction on institutional rather than efficiency grounds, treating land compensation as the government's sovereign responsibility and a test of its commitment to the project (World Bank, 2001; African Development Bank Independent Development Evaluation, 2019). To assess its distortive potential, I compare two funding rules, holding total external budget and maximum construction fixed as before:

- **Status-quo:** external donors fund construction (up to the realized envelope $\bar{M}$), while *all* land payments are funded domestically at the shadow price of public funds η. This is the regime under which the Kampala projects were delivered.
- **Donor land financing:** external funds may cover both construction *and* land acquisition at marginal cost 1, up to the same envelope $\bar{M}$, with any spending beyond it (necessarily land, since construction is capped at $\bar{M}$) subject to η. Relaxing the restriction removes the wedge on (donor-financed) land payments but, because the envelope is fixed, money spent compensating owners can no longer go to construction, so the city may build less.

In Figure 9, I compare these two scenarios across different values of η to trace out the allocative value of letting donor money pay for land. When η is low (panel (a), $\eta = 2$) or moderate (panel (b), $\eta = 4$), the two regimes nearly coincide, since the wedge donor land financing removes is small. As η rises to high values (panel (c), $\eta = 6$), the unrestricted scenario achieves substantially higher net returns (NPV USD 551 million against NPV USD 492 million under restricted funds use), despite building less (216,000 m^2 against 388,000 m^2), because the construction it forgoes is inframarginal and the savings from having to raise less domestic funds is large.

Both regimes take as given that donors finance construction at marginal cost 1. Were construction, like land, financed from domestic revenue at η, the fiscal markup would fall on the entire road bill rather than on land alone, so fewer links would clear the cost–benefit test and the optimal network would fall short of the donor-financed allocation (net returns of USD 244 million at $\eta = 6$, against USD 492 million under the status quo and USD 551 million under donor land financing).

9 Conclusion

Improving road infrastructure in fast-growing cities is an investment priority, yet little is known about the net returns to these investments in practice, and even less about the role land plays in shaping them. In this paper, I fill this gap by collecting new data on land ownership and compensation around a road-upgrading program in Kampala and embedding them in a quantitative spatial model. I show that the benefits of road improvements are large, but so are the costs associated to the acquisition of the underlying land, and that the returns depend on how land is acquired and paid for, not only on how much road is built. Because compensation is financed from costly public domestic funds, paying for land at market value would turn the program's realized welfare gains negative; realized gains are positive largely because most owners, especially those with less formal property rights, go uncompensated.

This same cost of public funds shapes where it is optimal to build. Compensating owners at market value pulls investment away from the high-value center and is the most distorted regime, whereas a claim-based rule delivers both higher net returns and an allocation closer to the first best. Financing rules matter for the same reason: the donor convention of paying for construction but not land forces compensation onto costly domestic revenue, and bites hardest where those funds are dearest. Allowing external funds to cover land, as some donors' eligibility rules now permit (African Development Bank, 2020), can raise welfare. Ultimately, land acquisition in Kampala is a textbook case of the theory of the second best: enforcing market-value compensation to correct one imperfection need not improve welfare in the presence of another, the cost of raising domestic funds. Accounting for existing property rights and the conditions of land acquisition is therefore central to designing infrastructure investment in rapidly growing cities.

Last, the mechanism I isolate is specific to neither roads nor Kampala, but may distort where schools, clinics, transit, or utilities are built in crowded urban areas. More generally, it applies wherever a government finances a spatially fixed input from costly public funds and that input is priced highest exactly where the public investment is most valuable. The distortion is largest when three conditions coincide: public funds are costly to raise, the input's price is dispersed across space, and that price is aligned with the returns to the investment.

References

Acemoglu, Daron and James Robinson, *Why Nations Fail: the Origins of Power, Prosperity, and Poverty*, New York: Crown Publishers, 2012.

African Development Bank, "Bank Group Policy on Eligible Expenditures (Version II)," 2020. African Development Bank Group, Abidjan.

_ , "Integrated Safeguards System: Operational Safeguards," 2023. African Development Bank Group, Abidjan.

African Development Bank Independent Development Evaluation, "Evaluation of Counterpart Funding at the African Development Bank Group," Technical Report, African Development Bank Group 2019.

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf, "The Economics of Density: Evidence From the Berlin Wall," *Econometrica*, 2015, *83*, 2127–2189. Publisher: Econometric Society.

Akbar, Prottoy A., Victor Couture, Gilles Duranton, and Adam Storeygard, "The Fast, the Slow, and the Congested: Urban Transportation in Rich and Poor Countries," *NBER Working Paper w31642*, 2023.

Alani, Jimmy, Lorenzo Casaburi, Martina Jakob, Joseph Mivule, and Cissy Namuddu, "Land Subdivision, Tenure, and Land Use Change: Evidence from Uganda," 2026. Work in Progress.

Allen, Treb and Costas Arkolakis, "The Welfare Effects of Transportation Infrastructure Improvements," *The Review of Economic Studies*, 02 2022, *89* (6), 2911–2957.

Angel, Shlomo et al., *Atlas of Urban Expansion—2016 Edition, Volume 2: Blocks and Roads*, New York; Nairobi; Cambridge, MA: New York University, UN-Habitat, and Lincoln Institute of Land Policy, 2016.

Asher, Sam and Paul Novosad, "Rural Roads and Local Economic Development," *American Economic Review*, March 2020, *110* (3), 797–823.

Auriol, Emmanuelle and Michael Warlters, "The marginal cost of public funds and tax reform in Africa," *Journal of Development Economics*, 2012, *97* (1), 58–72.

Balboni, Clare, "In Harm's Way? Infrastructure Investments and the Persistence of Coastal Cities," *Working Paper*, 2023.

_ **, Gharad Bryan, Melanie Morten, and Bilal Siddiqi**, "Transportation, Gentrification, and Urban Mobility: The Inequality Effects of Place-Based Policies," *Working Paper*, 2020.

Bassi, Vittorio, Matthew E. Kahn, Nancy Lozano Gracia, Tommaso Porzio, and Jeanne Sorin, "Jobs in the Smog: Firm Location and Workers' Exposure to Pollution in African Cities," April 2026.

Baum-Snow, Nathaniel, Loren Brandt, J Vernon Henderson, Matthew A Turner, and Qinghua Zhang, "Roads, railroads, and decentralization of Chinese cities," *Review of Economics and Statistics*, 2017, *99* (3), 435–448.

Bergeron, Augustin, Gabriel Z. Tourek, and Jonathan L. Weigel, "The State Capacity Ceiling on Tax Rates: Evidence from Randomized Tax Abatements in the DRC," September 2023, (31685).

Besley, Timothy and Maitreesh Ghatak, "Property Rights and Economic Development," in Dani Rodrik and Mark Rosenzweig, eds., *Dani Rodrik and Mark Rosenzweig, eds.*, Vol. 5, Elsevier, 2010, chapter Chapter 68, pp. 4525–4595.

_ **and Torsten Persson**, "Why Do Developing Countries Tax So Little?," *Journal of Economic Perspectives*, November 2014, *28* (4), 99–120.

Bird, Julia and Anthony J. Venables, "Land tenure and land-use in a developing city: A quantitative spatial model applied to Kampala, Uganda," *Journal of Urban Economics*, 2020, *119* (C).

Bordeu, Olivia, "Commuting Infrastructure in Fragmented Cities," *Working Paper*, 2024.

Brockmeyer, Anne, Alejandro Estefan, Karina Ramirez Arras, and Juan Carlos Suarez Serrato, "Taxing Property in Developing Countries: Theory and Evidence from Mexico," *NBER Working Paper W28637*, 2023.

Brooks, Leah and Zachary Liscow, "Infrastructure Costs," *American Economic Journal: Applied Economics*, April 2023, *15* (2), 1–30.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Collier, Paul and Anthony J. Venables, "Urban infrastructure for development," *Oxford Review of Economic Policy*, 2016, *32* (3), 391–409.

Currier, Lindsey, Edward L. Glaeser, and Gabriel Kreindler, "Infrastructure Inequality: Who Pays the Cost of Road Roughness," *NBER Working Paper w31981*, 2023.

Dahlby, Bev, *The Marginal Cost of Public Funds: Theory and Applications*, Cambridge, MA: MIT Press, 2008.

Donaldson, Dave, "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, April 2018, *108* (4-5), 899–934.

Dunning, Iain, Joey Huchette, and Miles Lubin, "JuMP: A Modeling Language for Mathematical Optimization," *SIAM Review*, 2017, *59* (2), 295–320.

Duranton, Gilles and Matthew A. Turner, "The Fundamental Law of Road Congestion: Evidence from US Cities," *American Economic Review*, October 2011, *101* (6), 2616–52.

_ and _ , "Urban Growth and Transportation," *Review of Economic Studies*, 2012, *79* (4), 1407–1440.

Fajgelbaum, Pablo D. and Edouard Schaal, "Optimal Transport Networks in Spatial Equilibrium," *Econometrica*, 2020, *88* (4), 1411–1452.

Frankel, Jeffrey A., Carlos A. Végh, and Guillermo Vuletin, "On Graduation from Fiscal Procyclicality," *Journal of Development Economics*, 2013, *100* (1), 32–47.

Glaeser, Edward and J. Vernon Henderson, "Urban economics for the developing World: An introduction," *Journal of Urban Economics*, March 2017, *98*, 1–5.

Gonzalez-Navarro, Marco and Climent Quintana-Domeque, "Paving Streets for the Poor: Experimental Analysis of Infrastructure Effects," *The Review of Economics and Statistics*, May 2016, *98* (2), 254–267.

Harari, Mariaflavia, "Cities in Bad Shape: Urban Geometry in India," *American Economic Review*, August 2020, *110* (8), 2377–2421.

Holland, Alisha C., "Roadblocks: How Property Rights Undermine Development in Colombia," *American Journal of Political Science*, 2023, *67* (3), 639–655.

Hornbeck, Richard, "Barbed Wire: Property Rights and Agricultural Development," *The Quarterly Journal of Economics*, 2010, *125* (2), 767–810.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 11 2009, *124* (4), 1403–1448.

Ilzetzki, Ethan and Carlos A. Végh, "Procyclical Fiscal Policy in Developing Countries: Truth or Fiction?," NBER Working Paper 14191, National Bureau of Economic Research 2008.

International Monetary Fund, "Making Public Investment More Efficient," IMF Policy Paper, International Monetary Fund, Washington, DC June 2015. https://www.imf.org/external/np/pp/eng/2015/061115.pdf.

Jeong, David, DD Gransberg, Heedae Park, and Nikhil Shelar, "Barriers to right-of-way acquisition and recommendations for change," Technical Report, Minnesota Department of Transportation, Research Services & Library 2016.

Kaminsky, Graciela L., Carmen M. Reinhart, and Carlos A. Végh, "When It Rains, It Pours: Procyclical Capital Flows and Macroeconomic Policies," in Mark Gertler and Kenneth Rogoff, eds., *NBER Macroeconomics Annual 2004, Volume 19*, Cambridge, MA: MIT Press, 2004, pp. 11–53.

Kampala Capital City Authority, "The Kampala City Roads Rehabilitation Project (KCRRP): Progress Report (BQPR No. 11), 1st Quarter FY2025/2026," Technical Report, Kampala Capital City Authority September 2025. Project ID No. P-UG-DB0-016, African Development Bank.

Knebelmann, Justine, "Taxing Property Owners in Dakar," *IGC Policy Brief 50415*, 2019.

Kreindler, Gabriel, Arya Gaduh, Tilman Graff, Rema Hanna, and Benjamin A. Olken, "Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta," NBER Working Papers 31369, National Bureau of Economic Research, Inc June 2023.

Kreindler, Gabriel E. and Yuhei Miyauchi, "Measuring Commuting and Economic Activity inside Cities with Cell Phone Records," *The Review of Economics and Statistics*, 07 2023, *105* (4), 899–909.

Krueger, Anne O, "The Political Economy of the Rent-Seeking Society," *American Economic Review*, June 1974, *64* (3), 291–303.

Lall, Somik Vinay, J. Vernon Henderson, and Anthony J. Venables, *Africa's Cities: Opening Doors to the World*, Washington, DC: World Bank, 2017.

Lindsay, Jonathan Mills, "Compulsory Acquisition of Land and Compensation in Infrastructure Projects," PPP Insights Vol. 1, Issue 3, World Bank, Washington, DC 2012.

Manwaring, Priya and Tanner Regan, "Public Disclosure and Tax Compliance in a Low-Capacity State," 2024. https://tannerregan.github.io/tanner_regan/Manwaring_Regan_unp2024.pdf.

McFadden, Daniel, "The measurement of urban travel demand," *Journal of Public Economics*, 1974, *3* (4), 303–328.

Ministry of Finance, Planning and Economic Development, "Debt Sustainability Analysis Report FY2023/24," Technical Report, Government of Uganda December 2024.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg, "Commuting, Migration, and Local Employment Elasticities," *American Economic Review*, December 2018, *108* (12), 3855–90.

Morten, Melanie and Jaqueline Oliveira, "The Effects of Roads on Trade and Migration: Evidence from a Planned Capital City," *American Economic Journal: Applied Economics*, April 2024, *16* (2), 389–421.

Munch, Patricia, "An Economic Analysis of Eminent Domain," *Journal of Political Economy*, 1976, *84* (3), 473–497.

Olken, Benjamin A., "Monitoring Corruption: Evidence from a Field Experiment in Indonesia," *Journal of Political Economy*, 2007, *115* (2), 200–249.

Posner, Eric A. and E. Glen Weyl, "Property Is Only Another Name for Monopoly," *Journal of Legal Analysis*, 04 2017, *9* (1), 51–123.

Rafa, Mickey, Jonathan D. Moyer, Xuantong Wang, and Paul Sutton, "Estimating District GDP in Uganda," Report prepared for USAID, Frederick S. Pardee Center for International Futures, Josef Korbel School of International Studies, University of Denver 11 2017.

Redding, Stephen J. and Esteban Rossi-Hansberg, "Quantitative Spatial Economics," *Annual Review of Economics*, 2017, *9* (Volume 9, 2017), 21–58.

Reinhart, Carmen M., Kenneth S. Rogoff, and Miguel A. Savastano, "Debt Intolerance," *Brookings Papers on Economic Activity*, 2003, *2003* (1), 1–74.

Reinikka, Ritva and Jakob Svensson, "Local Capture: Evidence from a Central Government Transfer Program in Uganda," *Quarterly Journal of Economics*, 2004, *119* (2), 679–705.

Shavell, Steven, "Eminent Domain versus Government Purchase of Land Given Imperfect Information about Owners' Valuations," *Journal of Law and Economics*, February 2010, *53* (1), 1–27.

Shepard, Donald, "A two-dimensional interpolation function for irregularly-spaced data," in "in" ACM Conferences ACM 1968, p. 8.

Silva, J. M. C. Santos and Silvana Tenreyro, "The Log of Gravity," *The Review of Economics and Statistics*, 11 2006, *88* (4), 641–658.

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

Tagliarino, Nicholas K., "The Status of National Legal Frameworks for Valuing Compensation for Expropriated Land: An Analysis of Whether National Laws in 50 Countries/Regions across Asia, Africa, and Latin America Comply with International Standards on Compensation Valuation," *Land*, 2017, *6* (2), 37.

Treb, Arkolakis Costas. Allen and Xiangliang Li, "Optimal city structure," *Mimeo*, 2015.

Tsivanidis, Nick, "Evaluating the Impact of Urban Transit Infrastructure: Evidence from BogotÃ¡'s TransMilenio," *American Economic Review*, February 2026, *116* (2), 418â63.

Wächter, Andreas and Lorenz T Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical programming*, 2006, *106* (1), 25–57.

World Bank, "Operational Policy 4.12: Involuntary Resettlement," 2001. World Bank Operational Manual (rev. 2013), Washington, DC.

_ , "Uganda — Second Kampala Institutional and Infrastructure Development Project (KIIDP-2): Project Appraisal Document," Project Appraisal Document Report No. PAD800, World Bank, Washington, DC February 2014. Project ID P133590. http://documents.worldbank.org/curated/en/504911468115450273.

_ , "Discounting Costs and Benefits in Economic Analysis of World Bank Projects," Guidance Note, World Bank, Operations Policy and Country Services (OPSPQ), Washington, DC 2016.

_ , "Environmental and Social Standard 5: Land Acquisition, Restrictions on Land Use and Involuntary Resettlement," 2017. Environmental and Social Framework, Washington, DC.

_ , "Implementation Completion and Results Report (IDA-53840) on a Credit to the Republic of Uganda for the Second Kampala Institutional and Infrastructure Development Project," Implementation Completion and Results

Report ICR00005839, World Bank, Urban, Resilience and Land Global Practice, Eastern and Southern Africa Region December 2023.

Zarate, Roman David, "Spatial Misallocation, Informality and Transit Improvements: Evidence from Mexico City," *Working Paper*, 2024.

Figure 1: Road Improvements in Kampala

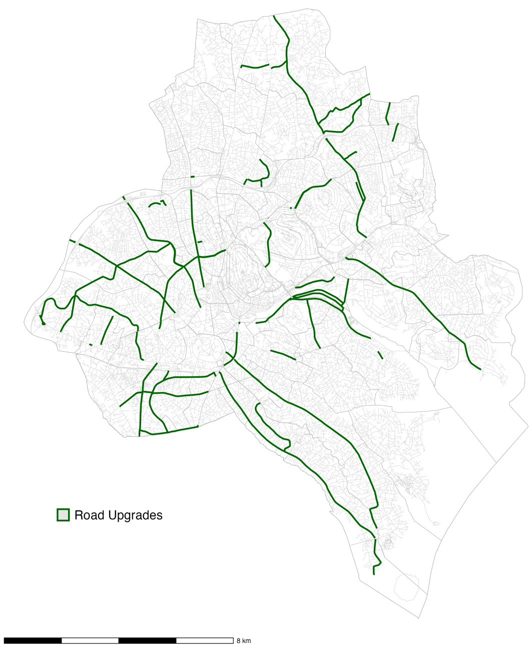

Notes: Road improvements in Kampala since 2018 that were funded by external donors (World Bank and African Development Bank).

Figure 2: Property Right Regimes in Kampala

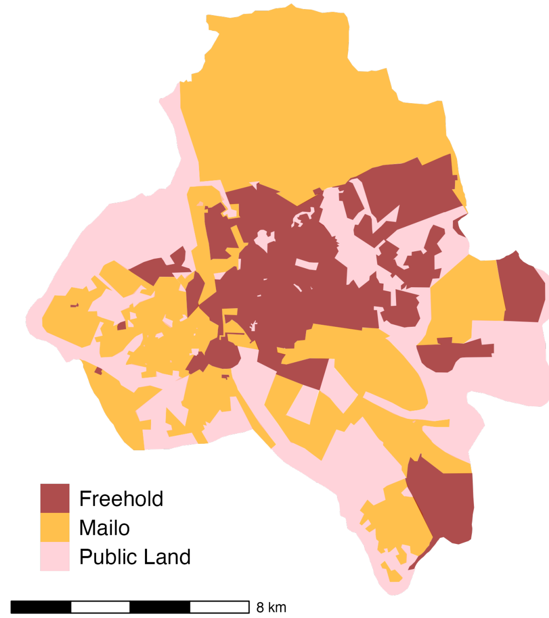

Notes: This map displays the three dominant land tenure systems in Kampala, together with the parish administrative boundaries. Source: GIS Unit at KCCA.

Figure 3: Pre-policy (2018) Land Sale Rates By Parish

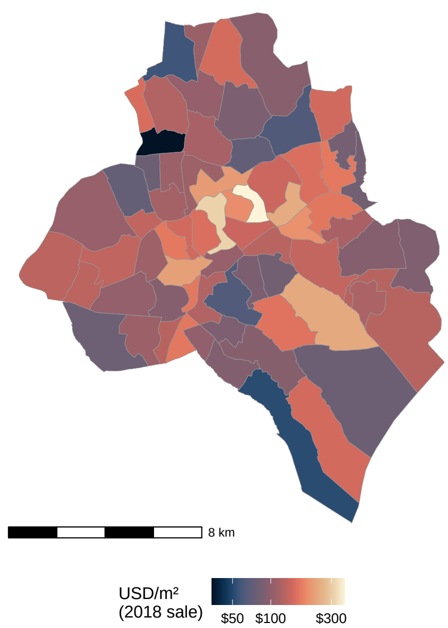

Notes: 2018 residential and business land rates in USD/m^2 are recovered from the broker survey following the approach described in Appendix Section C. The maps displayed the blended rate across residential and business land, with the weights corresponding to the share of residential and business land in each parish, respectively. Parish-level rates (N=96) are consolidated into the spatial units used in the main quantitative analysis (N=68).

Figure 4: Event Study: Road Upgrades on Traffic Speed, by Time of Day

(a) Rush hours (morning + evening)

(b) Off-peak hours

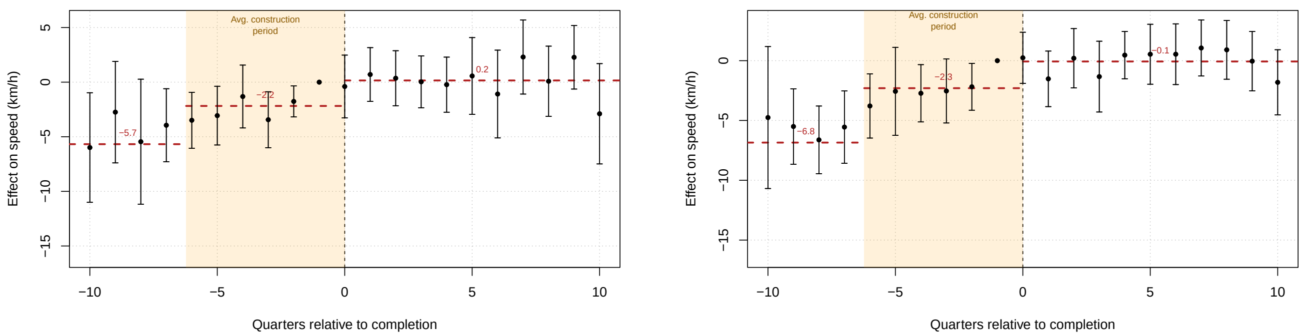

Notes: Each panel plots quarterly Sun and Abraham (2021) interaction-weighted estimates of the effect of road-upgrade completion on traffic speed (km/h), relative to the quarter before completion (period −1), over a window of ±10 quarters. The dependent variable is the road × month average speed, residualized on day and hour-of-day fixed effects. Panel (a) restricts to rush-hour trips (morning and evening peaks pooled); panel (b) to off-peak trips. The vertical dashed line marks the completion quarter (period 0); the shaded area indicates the average construction period preceding it. The horizontal dashed (red) segments mark the average event-study coefficient within the pre-construction, construction, and post-completion windows respectively, illustrating the gradual increase in speed across the three regimes (the effect should be read relative to the construction and pre-construction levels, not only the reference period). Standard errors are clustered by road segment; vertical bars show 95% confidence intervals.

Figure 5: Road Upgrades Increased Local Property Values (Stated-Preferences)

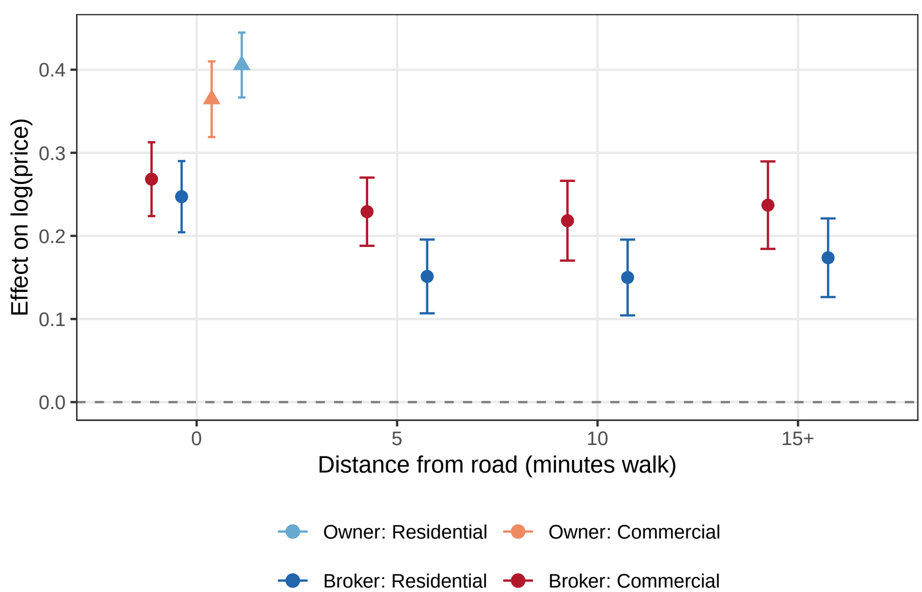

Notes: This figure displays the estimated coefficients of road upgrade (dummy) on log property values. Broker estimates (circles) are from a within-broker before/after design using the hypothetical property valuation module of the broker survey, in which brokers reported the value of a standardized property before and after a hypothetical road upgrade at distances of 0, 5, 10, and 15+ minutes walk. Each point is from a separate regression with broker fixed effects; standard errors are clustered at the broker level. Owner estimates (triangles, at distance 0 only) are from a within-owner before/after design using self-assessed property values from the owner survey, restricted to owners with a completed road upgrade nearby; standard errors are clustered at the respondent level. Bars represent 95% confidence intervals.

Figure 6: Net NPV of Realized Upgrades by Compensation Rule

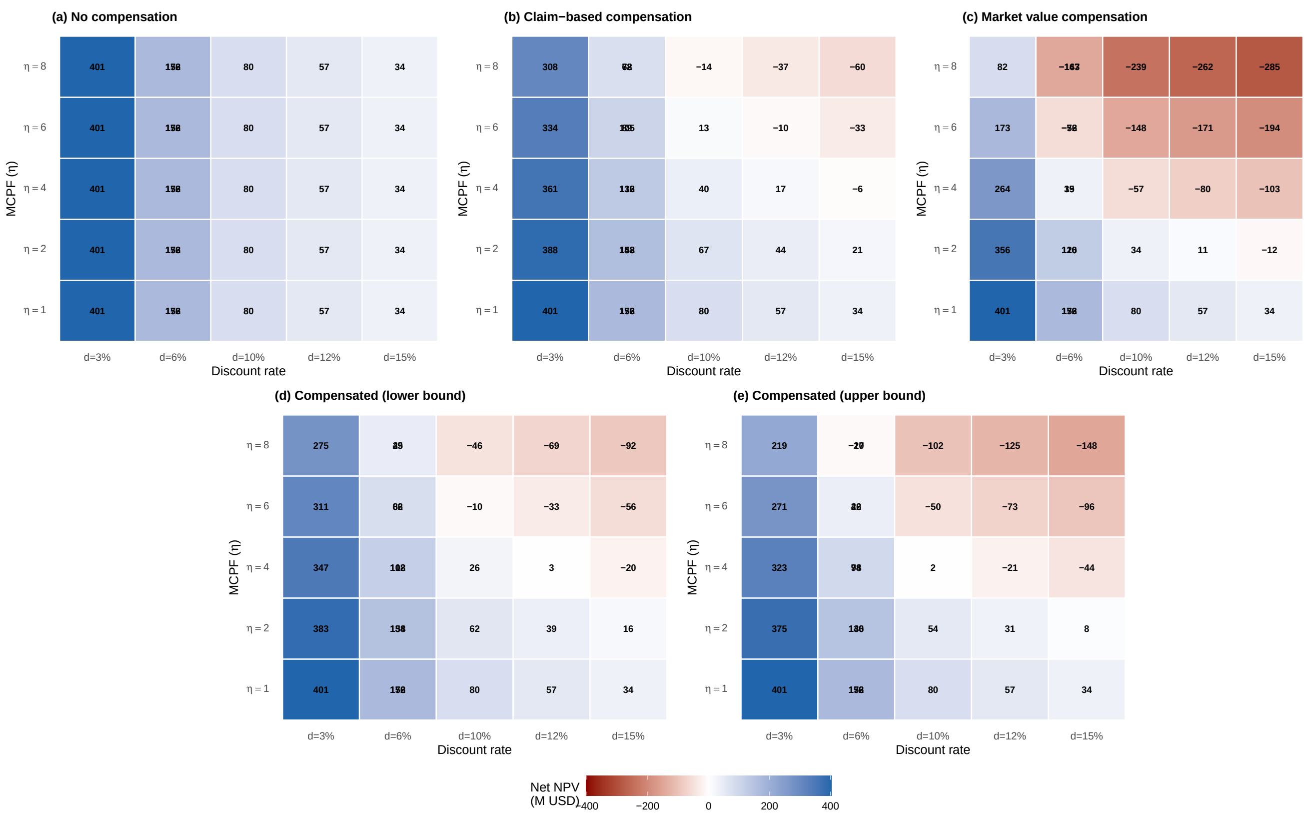

Notes: Net NPV = NPV of welfare gains − construction costs − $\eta\times$ land compensation. Each cell shows the net NPV (million USD) for a given (η, d) pair. Blue = positive, red = negative. Panel (a) represents the case without land compensation. Panel (b) represents the case with claim-based compensation. Panel (c) represents the case with market-value compensation (legally mandated under the eminent domain legal regime). Panel (d) represents the hybrid compensation approach observed in the data where some landowners were offered market value compensation and some had to claim it.

Figure 7: Marginal Costs, Benefits, and Net Returns of Link-Level Road Upgrades

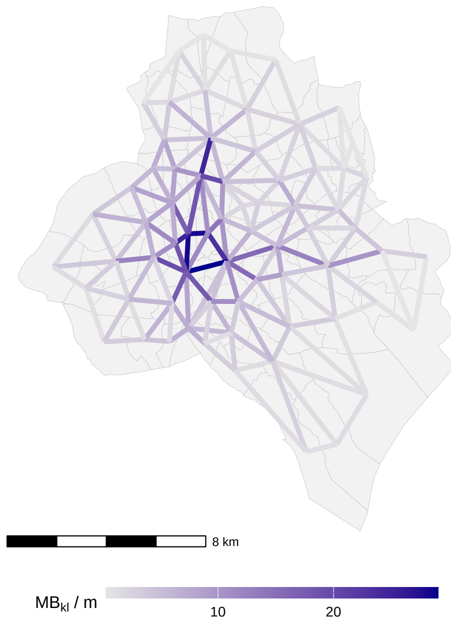

(a) Marginal Benefit (%$\Delta\tau$ / m)

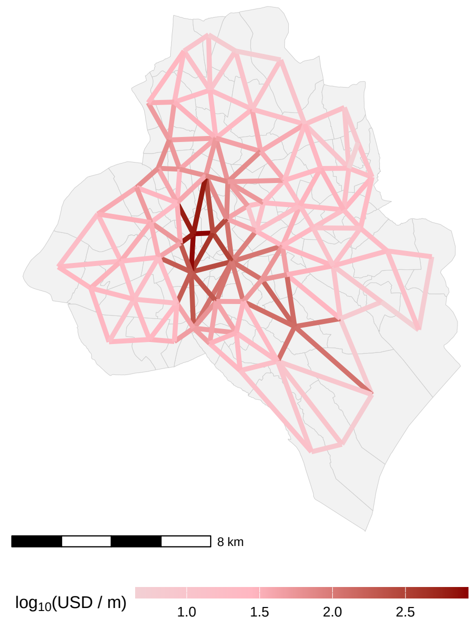

(b) Land Cost, Market Value (USD / m)

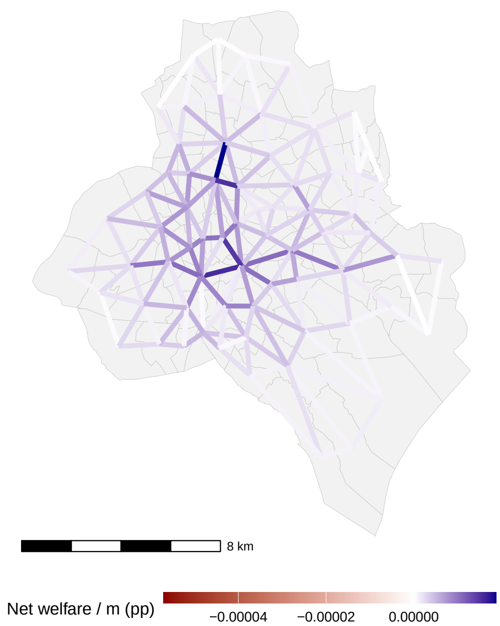

(c) Net Return, No Compensation

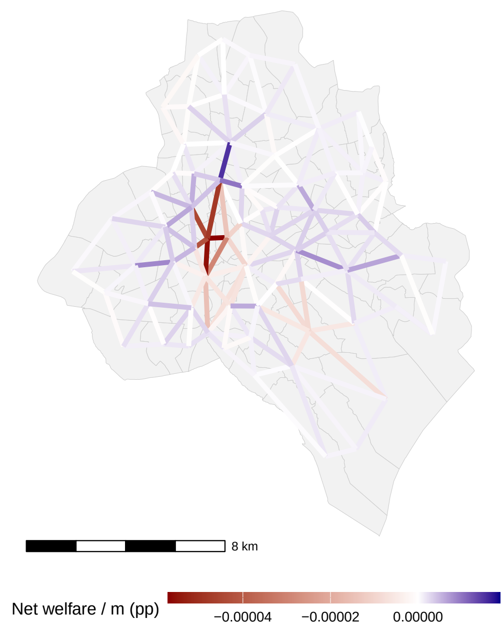

(d) Net Return, Market-Value Compensation

Notes: Each link connects two neighboring parishes. I widen each link by 3.4m (one lane) and solve for the new spatial equilibrium. Panel (a): marginal benefit (commuting cost decline per meter). Panel (b): market-value land cost per meter (log scale). Panels (c) and (d): NPV of net welfare gains per meter under no compensation and under market-value compensation with $\eta = 4$, on a shared color scale (blue = positive, red = negative). Owners' utility is linear in income, with marginal utility equal to 1.

Figure 8: Optimal allocation of road widening under each land-compensation rule.

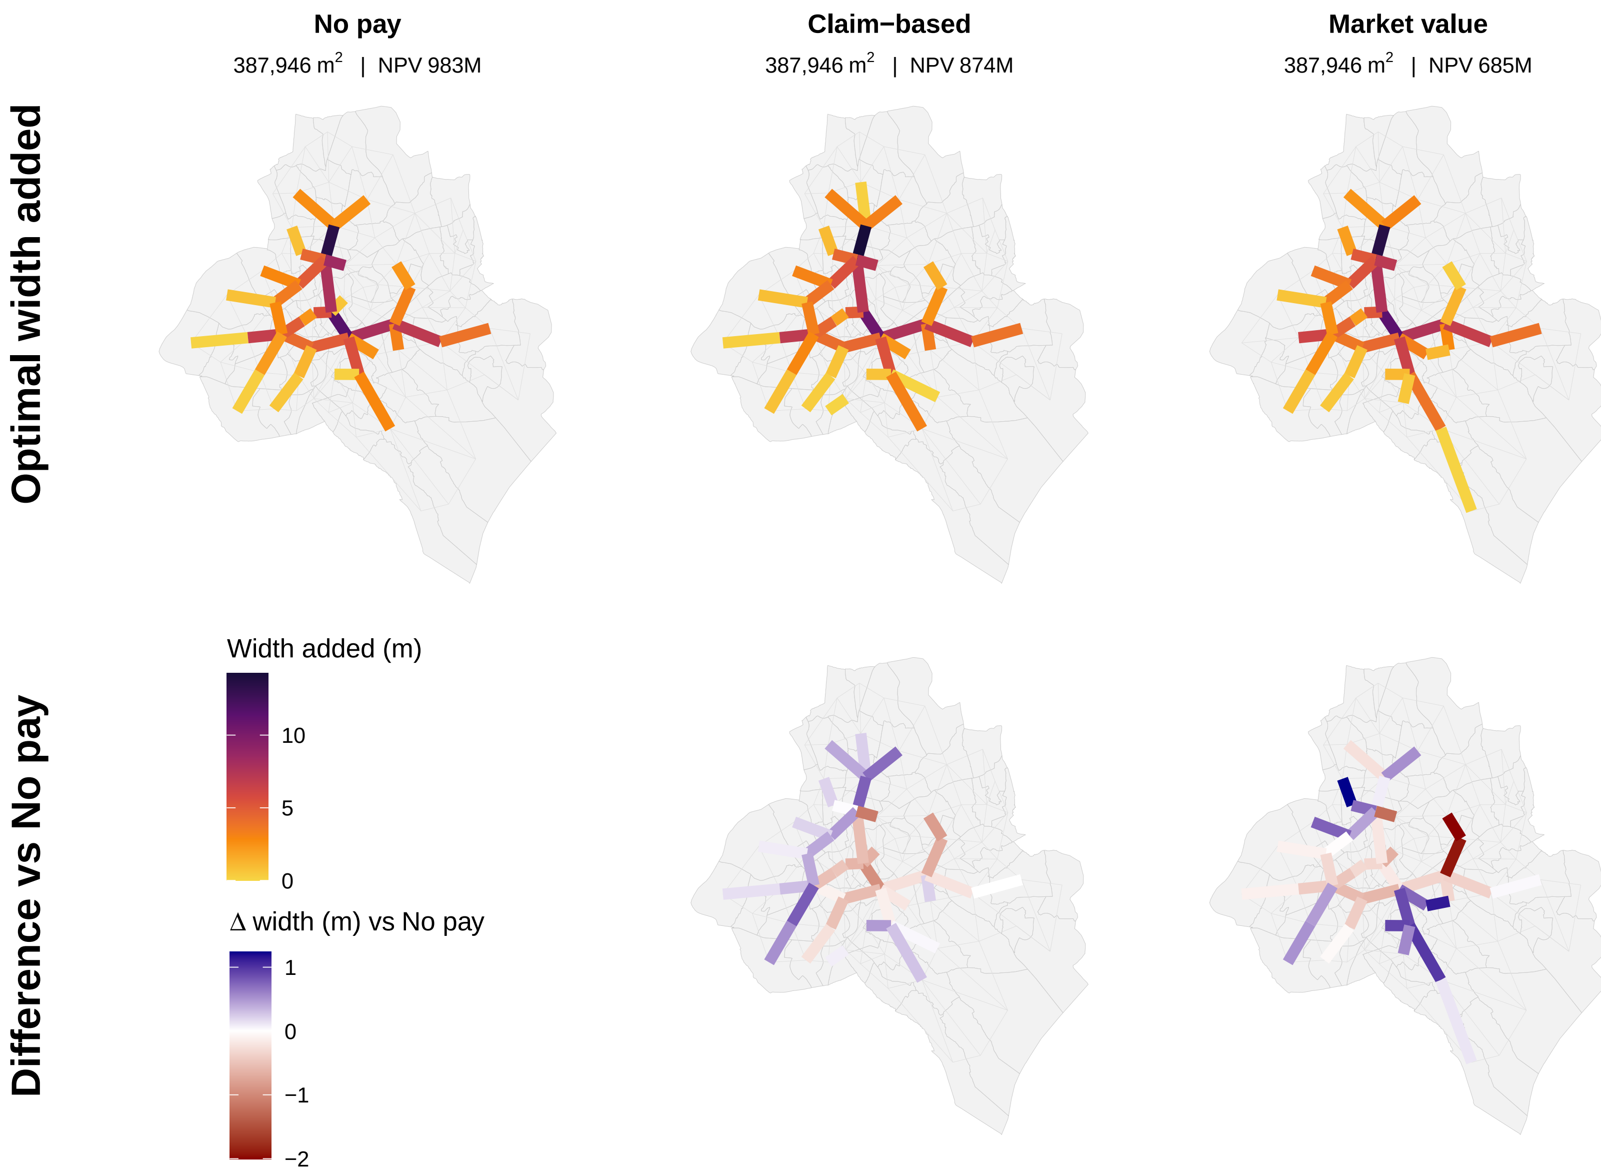

Notes: Top row: optimal width added to each link (metres) under no compensation, claim-based compensation, and market-value compensation, at $\eta = 4$ and $\delta = 0.06$ and financing land payments under the status-quo rule. *Bottom row:* the change in widening relative to the no-compensation allocation (red: less widening; blue: more), for the claim and market-value rules. Panel subtitles report the total widened area (m^2) and the net present value of the program (US$ millions). No compensation reproduces the first-best map, concentrating widening on the high-benefit central ring; market value is the most distorted, pulling investment out of the expensive centre toward cheaper peripheral links; claim-based compensation lies in between, closer to the no-compensation/first-best allocation than market value, since owners claim only on part of the affected land so the welfare wedge applies to a smaller share of the payment.

Figure 9: Optimal allocation under alternative financing rules by η.

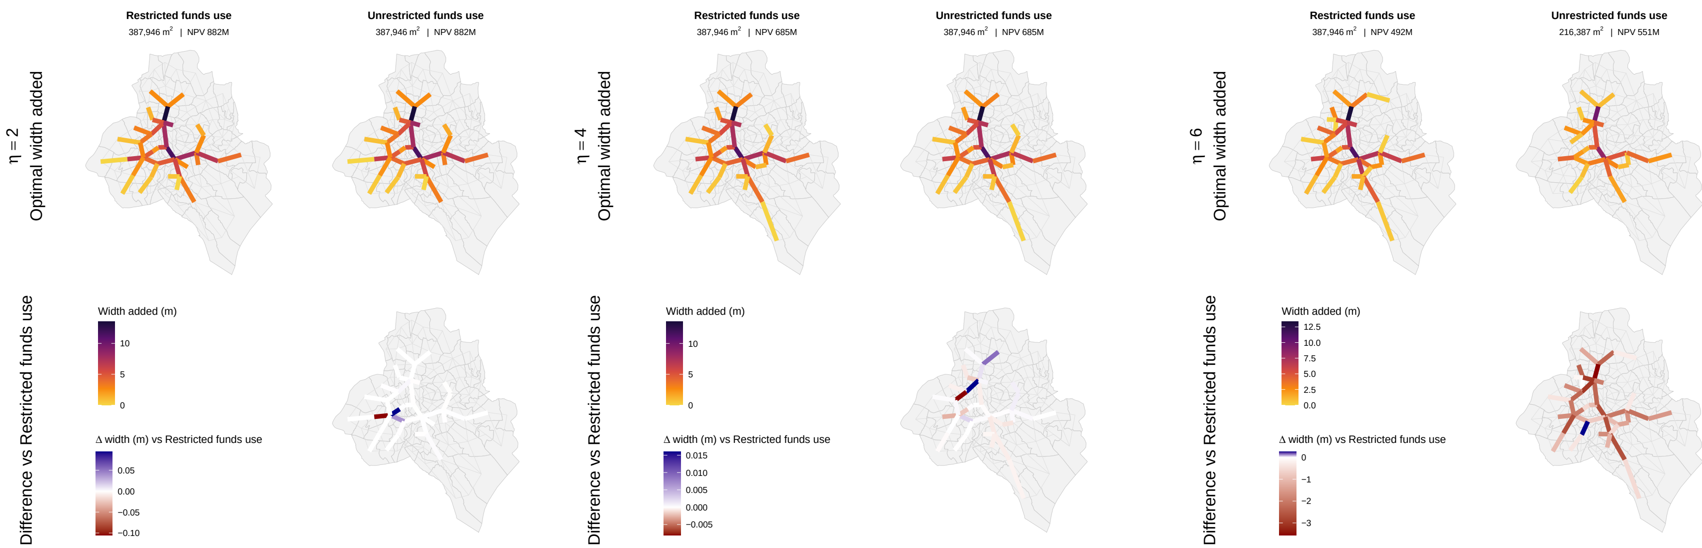

(a) $\eta = 2$ (b) $\eta = 4$ (c) $\eta = 6$

Notes: Top row: optimal width added to each link (metres) under market-value compensation, at $\delta = 0.06$, under two financing rules: *restricted funds use*, in which donor money is confined to construction and land payments must come from costly domestic revenue (the status quo), and *unrestricted funds use*, in which donor money may also finance land. *Bottom row:* the change in widening relative to the restricted rule (red: less widening; blue: more). Panel subtitles report the total widened area (m^2) and the net present value of the program (US$ millions). Lifting the restriction lets land be paid out of external funds rather than costly domestic revenue, neutralising the welfare wedge $(\eta - 1)\, p_k^{\text{land}}$ on land payments: the planner restores widening on the expensive central links it had withdrawn from, moving the allocation back toward the first best and raising net welfare. Panels report the optimal allocation of road widening for a low ($\eta = 2$, panel (a)), medium ($\eta = 4$, panel (b)), and high ($\eta = 6$, panel (c)) cost of public funds, under market-value compensation at $\delta = 0.06$. Because the restriction binds only through the wedge $(\eta - 1)\, p_k^{\text{land}}$, the value of lifting it grows with η: at $\eta = 2$ the restricted and unrestricted allocations differ only modestly, whereas at $\eta = 6$ unrestricting donor funds restores substantial widening on the expensive central links and yields a much larger net-welfare gain.

Table 1: Land Tenure and Compensation Outcomes

	Offered, comp. or claimed		Compensated (received)		Offered		Claimed	
	all	past	all	past	all	past	all	past
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Freehold	0.365	0.463	0.132	0.187	0.190	0.252	1.766	2.577
	(0.090)	(0.092)	(0.093)	(0.100)	(0.085)	(0.087)	(0.557)	(0.695)
Public land	0.288	0.342	0.255	0.244	0.287	0.293	0.681	1.270
	(0.060)	(0.068)	(0.059)	(0.066)	(0.058)	(0.065)	(0.408)	(0.533)
Land value (MV, std.)	0.043	0.071	0.006	0.030	0.020	0.052	0.243	0.458
	(0.038)	(0.038)	(0.039)	(0.043)	(0.038)	(0.039)	(0.218)	(0.327)
Estimator	OLS	OLS	OLS	OLS	OLS	OLS	Logit	Logit
Tenure + market value + wealth + parish controls				Yes				
Observations	311	257	310	257	311	257	212	178
R^2	0.15	0.20	0.15	0.18	0.19	0.24	—	—

Notes: Columns 1–6 report OLS estimates with heteroskedasticity-robust standard errors in parentheses. The reference tenure category is mailo. Each outcome is reported on the full sample ("all", odd columns) and on the subsample of already-completed (past) upgrades ("past", even columns). "Offered, comp. or claimed" (cols 1–2) equals one if the owner was offered, received, or claimed financial compensation. "Compensated" (cols 3–4) equals one if the owner received financial compensation. "Offered comp." (cols 5–6) equals one if the owner received an initial offer of financial compensation. "Claimed comp." (cols 7–8) is estimated by MLE logit on the subsample of owners not initially offered compensation and for whom the claiming decision was recorded; the dependent variable equals one if the owner filed a claim. "Market value" is the standardized broker-estimated market value of affected land (missing values imputed to zero with a missing indicator). "Wealth index" is a standardized mean of eight binary household indicators. Parish-level controls are parish population, share of owner-occupied households, and area. For the MLE logit columns, missing values in regressors are handled via mean replacement with a missing-indicator dummy; coefficients are in log-odds.

Table 2: Road Upgrades Increase Traffic Speed, Concentrated in Rush Hours

	Rush (morning + evening)				Off-peak			
	Trip-level		Road × month		Trip-level		Road × month	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Upgrade Completed (d)	4.683	4.628	4.235	4.295	5.994	3.038	2.781	0.998
	(0.795)	(1.034)	(1.020)	(1.534)	(1.160)	(0.935)	(0.997)	(0.855)
Upgrade In Progress (d)	0.742	1.195	1.100	-1.029	0.562	-0.140	0.093	-0.736
	(0.770)	(0.769)	(0.785)	(0.730)	(0.792)	(0.709)	(0.678)	(0.564)
Day FE	Yes	Yes			Yes	Yes		
Hour FE	Yes	Yes			Yes	Yes		
Month FE			Yes	Yes			Yes	Yes
Road FE		Yes	Yes	Yes		Yes	Yes	Yes
Observations	1,017	1,017	752	752	1,075	1,075	684	684
Mean dep. var.	21.25	21.25	21.31	21.31	24.01	24.01	24.26	24.26
Geo controls	Yes				Yes			
Road class	Yes				Yes			
Paved (baseline)	Yes				Yes			
Model	OLS	TWFE	TWFE	SA	OLS	TWFE	TWFE	SA
SA leave-one-cohort-out				[2.3, 4.9]				[0.2, 1.6]

Notes: Standard errors in parentheses. Standard errors are two-way clustered by road segment and day in the trip-level columns (1, 2, 5, 6) and by road segment in the panel columns (3, 4, 7, 8). The dependent variable is traffic speed (km/h) from Google Maps API queries for trips along policy roads. Columns (1)–(4) report estimates for trips during *rush hours* (morning and evening peaks pooled); columns (5)–(8) for *off-peak* hours (all remaining hours). Within each block, columns (1)/(5) and (2)/(6) are estimated at the trip level and columns (3)/(7) and (4)/(8) at the road × month panel level; the panel dependent variable is the road–month average of trip-level speed residualized on day and hour-of-day fixed effects. The reported mean dependent variable is the trip-level mean speed in the trip columns and the road–month mean speed in the panel columns. "Upgrade Completed" is a dummy equal to one after the upgrade is completed; "Upgrade In Progress" is a dummy equal to one during construction. Columns (1)/(5) are OLS with day and hour-of-day fixed effects, road distance, geographic controls (longitude, latitude, and their squares), road-class dummies, and baseline paving status. Columns (2)/(6) are trip-level TWFE with road-segment, day, and hour-of-day fixed effects. Columns (3)/(7) are panel TWFE with road-segment and month fixed effects. Columns (4)/(8) report the average treatment effect on the treated over the first twelve months following completion (relative periods 0–11) from the Sun and Abraham (2021) interaction-weighted estimator, with delta-method standard errors. The fixed-effect rows refer to the estimating regression; in the panel columns day-of-sample and hour-of-day variation is absorbed by residualizing the outcome prior to aggregation. The "SA leave-one-cohort-out" row reports the range of the windowed ATT as each completion-month cohort is dropped in turn from the aggregation (re-weighting with coefficients held fixed), summarizing sensitivity to individual cohorts.

Table 3: Calibrated parameters.

Parameter	Symbol	Value	Source
Share of expenditure on housing	$1-\beta$	0.235	Uganda National Panel Survey (2019)
Firms' expenditure share on labor	α	0.7	Own calc., manufacturing firm survey, Bassi et al. (2026)
Route-choice elasticity	ρ	50	Set (near least-cost routing)
Owner-utility curvature	ψ^o	0	Linear owner utility (main spec.)
Average owner parcel area (m^2)	$\bar{A}$	497	Broker survey (free-subdivision)
Discount rate	δ	$0.06, 0.12$	World Bank (2016)
Marginal cost of public funds	η	[1,8]	Range; disciplined in §7

Table 4: Realized upgrades vs. model-predicted benefits and costs

Land Cost Measure	Market Value		Upper Bound		Lower Bound		Claim		Tenure	
	(1) Level	(2) Per m	(3) Level	(4) Per m	(5) Level	(6) Per m	(7) Level	(8) Per m	(9) Level	(10) Per m
Marginal benefit (RHS)	0.061	0.030	0.072	0.042	0.062	0.031	0.086	0.057	0.122	0.086
	(0.042)	(0.051)	(0.041)	(0.050)	(0.043)	(0.052)	(0.039)	(0.049)	(0.038)	(0.054)
Land cost	-0.075	-0.067	-0.071	-0.073	-0.041	-0.036	-0.089	-0.101		
	(0.048)	(0.035)	(0.058)	(0.053)	(0.060)	(0.056)	(0.050)	(0.049)		
Share freehold									-0.099	-0.081
									(0.072)	(0.072)
Share public									0.083	0.077
									(0.058)	(0.058)
Link length	0.057	-0.001	0.034	-0.003	0.024	0.005	0.019	-0.009	-0.037	-0.010
	(0.042)	(0.046)	(0.044)	(0.045)	(0.048)	(0.046)	(0.040)	(0.045)	(0.041)	(0.043)
R^2	0.025	0.019	0.029	0.021	0.018	0.007	0.040	0.035	0.074	0.050
N	172	172	172	172	172	172	172	172	172	172

Notes: Dependent variable: indicator for whether any road mapped to a model link (between neighboring parishes) was upgraded (2018–2026), regressed on the model's standardized marginal benefit (RHS of Equation 24), a standardized land-acquisition cost, and link length; dyadic-robust standard errors. Columns are grouped by compensation rule (market value, upper-/lower-bound compensated, claim-based, and tenure shares). Within each rule, "Level" measures benefit and cost as totals per link; "Per m" divides by link length.

Appendix

A Context Appendix

A.1 Land Acquisition Timeline

The acquisition of private land for road widening was typically governed by the following timeline. First, the government decided which roads would be upgraded based on a collection of factors, including land availability; the program's appraisal document lists the availability of funds to compensate affected owners among its formal selection criteria (World Bank, 2014). An engineering study was then conducted to assess the optimal road design and affected land, under the supervision of the chief government valuer. Second, the government reached out to local leaders and affected owners to inform them about the road project, and either offered them financial compensation (*involuntary land take*) or asked for their consent without financial compensation (*voluntary land take*). Among the latter owners, those refusing to consent may go through formal proceedings to get compensation, as legally entitled to, conditional on satisfying all conditions, with the eminent domain imposing de facto an upper bound on the compensation that can be claimed by affected owners and limiting the scope for holdout. Finally, once all land parcels had been secured, the road construction could start. Eminent domain is a well-studied solution to holdout problems, where private sellers strategically wait and exert their monopoly power (increasing as they become the last holdout) to seek higher compensation, leading to project delays and inflated costs (Posner and Weyl 2017). I abstract from the time dimension in this analysis.

A.2 Land Tenure

Table A1: Validation of the Tenure Map Against Parcel-Level Registry Data

	Village-level shares derived from parcel-level data		
	(1) Mailo share	(2) Leasehold share	(3) Freehold share
KCCA freehold share	-0.124	0.089	0.035
	(0.018)	(0.017)	(0.010)
KCCA public share	-0.152	0.154	-0.001
	(0.023)	(0.022)	(0.008)
Observations	805	805	805
R^2	0.071	0.067	0.025
Mean dep. var.	0.727	0.227	0.043

Notes: The dependent variables are village-level tenure shares (mailo, leasehold, freehold) derived from parcel-level land-registry data from Alani et al. (2026), aggregated to the smallest administrative unit in Kampala (LC1, or "villages"), restricted to parcels in the database with a registered title. Each column regresses one parcel-derived village share on the freehold and public area shares from the project's tenure map (KCCA), with the mailo share omitted as the reference category. The KCCA tenure map is well aligned with the parcel-level map from Alani et al. (2026): a higher KCCA freehold share predicts a higher parcel-level freehold share, a higher KCCA public share predicts a higher leasehold share, and both predict a lower mailo share (i.e. KCCA mailo maps to parcel-level mailo). The modest R^2 and the remaining discrepancies between the two maps reflect the limited availability of parcel-level registered titles with tenure information (25% of total area) and the fact that the KCCA map is drawn from historical delineations rather than precise, parcel-level records. Robust standard errors in parentheses.

B Data Appendix

B.1 Broker survey

B.1.1 Sampling Strategy

To collect transaction-level price data, we surveyed real-estate brokers across Kampala. Sampling proceeded at the parish level: each enumerator was assigned a set of target parishes and instructed to complete, on average, two broker interviews per parish, with several enumerators sometimes working the same parish. Because most brokers are not formally registered, enumerators located them by asking local council leaders and residents where brokers could be found — typically congregating at busy public places such as markets and taxi stages. To be eligible, a respondent had to be at least 18 years old; have worked as a real-estate broker in the area for at least two years; have rented out at least three properties in Kampala in the previous three months and sold at least one property in Kampala in the previous year; and have transacted (sold or rented) in the assigned parish, among others. These activity thresholds ensured that respondents were active, experienced brokers with direct and recent knowledge of local rental and sale prices. Data collection was conducted in two waves: March-May 2024 and November-December 2025.

B.1.2 Comparison with Scrapped Online Property Ads

To increase confidence in the reliability of the broker survey data, I compare it with data scraped online from RealEstateDatabase.net, which includes posts of properties to rent or sale in Kampala and other parts of Uganda. This online database includes some geographical information (at the *neighborhood* level, which is less precise than parish and very noisy), as well as property characteristics. In addition, the date at which the property is posted can be recovered from the file name of the photos accompanying the post (but not the transaction date).

The distribution of online property prices is skewed to the right, with an average price list at $\$332,694$ (against $\$70,261$ in the survey data). Given that the average monthly net salary in Kampala is \$147 (KCCA 2019 Statistical Abstract) and 18.4% of Kampala households own a laptop (UBOS 2014), properties posted online are upscale properties only, not within reach of the average household, and not representative.

Reassuringly, however, while price levels are different across the databases, they are positively correlated over space. To check this, I run the following regression

$$\ln \mathrm{p}_{pt}^{b} = \alpha + \beta \times \ln \mathrm{p}_{pt}^{o} + \gamma_t + \gamma_p + e_{pt}$$

where p_{pt}^{b} is the average property price in parish p, year t in the broker data, and analogously for the online database $\ln \mathrm{p}_{pt}^{o}$. γ_t and γ_p are year and parish fixed-effects, respectively. The results are displayed on Table A2 and confirm the robust positive correlation between the online and broker survey databases at the parish and parish-year levels.

Figure A1: Number of Broker Interviews by Parish

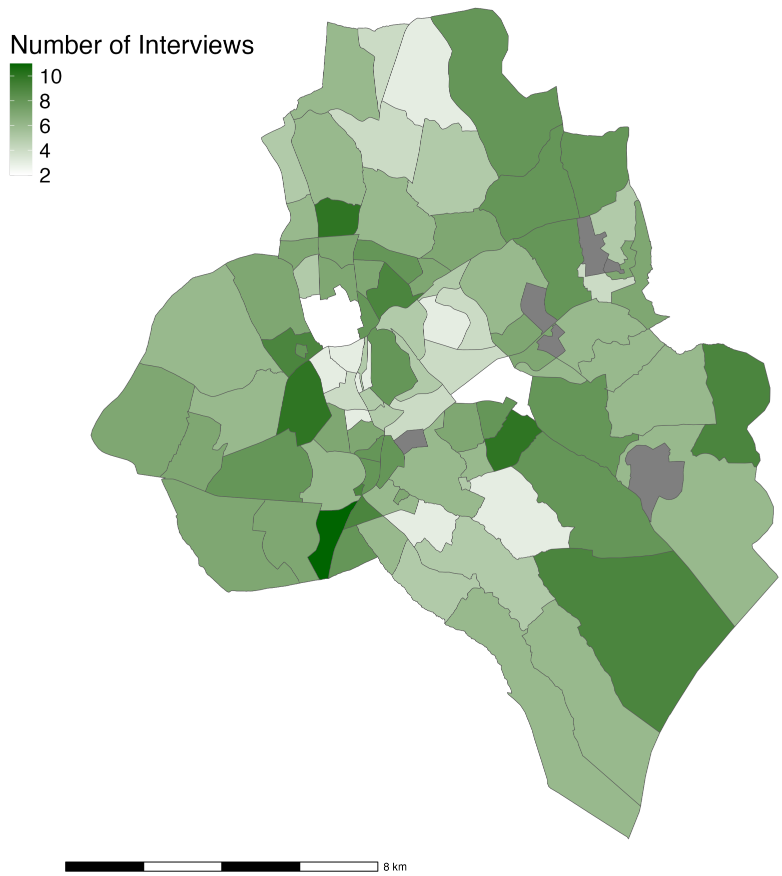

Notes: The map displays the number of broker interviews by parish.

Table A2: Comparison of Parish-Year Average Sales Price Between Online Data and Broker Data

Dependent Variable:	Log Parish-Year Mean Broker Survey Price (USD)			
Model:	(1)	(2)	(3)	(4)
Variables				
Log Parish-Year Mean Online Price (USD)	0.4340	0.4794	0.4649	0.2660
	(0.1178)	(0.1206)	(0.1699)	(0.1543)
Property type	Residential	Residential	Business	Business
Dep var trimmed top pct	1	5	1	5
Fit statistics				
Observations	114	111	72	67
R^2	0.10813	0.12659	0.09660	0.04370

Notes: This table assesses the consistency between broker-reported sale prices and an independent online price signal from RED, a Ugandan online real-estate marketplace. The unit of observation is a parish–year. In each column the dependent variable is the log of the average broker-survey *sale* price (USD) in a parish–year and the regressor is the log of the average online sale price (USD) in the same parish–year. Columns (1)–(2) cover residential properties and columns (3)–(4) business properties; within each property type the average price is computed after trimming the top and bottom 1% (columns 1 and 3) or 5% (columns 2 and 4) of the price distribution, applied separately to the survey and the online samples. Both samples are restricted to transactions and listings dated 2023 or later. Online listings are assigned to parishes through area-weighted matching of their self-reported location to administrative boundaries, with prices averaged across the listing's reported price fields and converted to USD; broker-survey prices come from brokers' reported past transactions. Only parish–years observed in both sources enter the regression. I.i.d. standard errors in parentheses.

Table A3: Broker Survey's Descriptive Statistics

	Mean	Sd	Median
Number of brokers interviewed	565		
Number of brokers interviewed (wave 1)	377		
Number of brokers interviewed (wave 2)	188		
Number of brokers interviewed per parish	5.88	2.42	6.00
Panel A: Broker characteristics			
Age	41.0	9.1	
Years operating as brokers	9.6	6.6	
Share of brokers operating in multiple parishes	0.67		
Share of brokers operating in multiple divisions	0.20		
Panel B: Activity			
Number of properties rented in the past 3 months in Kampala	9.22	10.86	
Number of properties sold in the past 3 months in Kampala	1.52	2.33	
Past Transactions (USD)			
Price - rented residential	148	123	108
Price - rented business	170	144	135
Price - sold residential	62,295	56,582	40,541
Price - sold business	72,335	64,782	56,757
Price - sold land	54,198	52,650	37,838
Price / m2 - sold residential	197	430	120
Price / m2 - sold business	312	483	153
Price / m2 - sold land	124	126	93
Appraisal Exercise - Standardized res property near road			
Relative residential sales price post upgrade	1.318	0.580	1.250

Notes: The descriptive statistics in this table come from the real estate broker survey. Panels A and B describe the characteristics of the respondents. Panel C describes the properties transacted by the respondents and panel D includes information about the impact of road upgrade on the hypothetical price of the standardized property. Wave 1 was conducted between March and May 2024, while wave 2 was conducted between November and December 2025.

B.2 Owner survey

B.2.1 Sampling Strategy

The sampling frame for the owner survey was at the road segment level. More precisely, roads within Kampala were split into categories based on their policy status: roads targeted by the policy with completed upgrade (pre-2020 or 2020-2024) and roads targeted by the policy but upgrade ongoing (2024 onwards). To ensure coverage of the different property right regimes in the city, the main tenure system within a 50 meters buffer of each road was identified. Then, for each triplet {division, land tenure, road category}, road segments were randomly ordered and the data collection team went down the list of valid road segments until data collection was completed.[39]

Enumerators conducted a listing of properties over a pre-specified road segment of approximately $\frac{1}{2}$ to 1 mile, between two intersections, starting from a randomly selected point over the road segment. In addition to the property listing, enumerators were instructed to interview all available potential respondents, moving from the (randomly selected) starting point to the end of the road segment. Selected road segments may include both affected and non-affected owners (for example as roads may have been widened on one side only) and the analysis in this paper is restricted to the subset of affected respondents. Figure A2 shows a map of the number of interviews of affected owners by parish, restricting to the main sample (owners who reported being officially contacted by the government about their land or property being affected by the road upgrade).

Eligible respondents were property (land or building) owners above 18 years old, their spouse or an individual with knowledge and decision power over land and property related issues. The targeted property must be on the side of one of the targeted roads. If the respondent was not available at that time, but willing to participate, the survey team offered to make an appointment and come back at a later time.

B.2.2 Overview of Survey Instruments

The survey instruments included the following sections:

1. Eligibility and consent
2. Road quality and road upgrade (name and characteristics of the closest road including a measure of road quality on a scale of 1 to 10, today and before the road upgrade ; road upgrade characteristics)
3. Land acquisition (amount of affected land ; compensation offer and received ; claim timeline ; reasoning ; perceptions including pivotality, potential returns from claiming, neighborhood interactions ; knowledge of the legal system ; suggestions)
4. Property value and property characteristics (self-assessed property value if sold or rented either (i) today, (ii) before the road upgrade, (iii) today if the road had not been upgraded ; ownership status ; property characteristics)

[39] Segments within approximately $1km$ from a previously selected segment were excluded, as well as segments where data collection was impossible (based on on-the-ground observations) or where road characteristics were different than in the database (for example a large paved road vs an unpaved road).

5. Household and neighborhood characteristics (religion ; ethnicity ; local leader popularity ; basic-socioeconomic characteristics on the model of the UNPS)

B.2.3 Comparing Survey Respondents to the Average Kampala Resident

In Table A4, I further compare owner survey respondents to the average Kampala resident as per the 2014 Population Census and the 2019 Uganda National Panel Survey (UNPS) along the lines of wealth proxies, education (panel A) and property characteristics (panel B). On average, survey respondents are wealthier and more educated than the average Kampala resident from the 2014 Population Census. Comparing the 2019 UNPS participants that are property owners and those that are renters (columns 6 and 7) reveals that property owners are wealthier than the average Kampala resident. In turn, the respondents to my owner survey tend to be wealthier than the average Kampala owner, which is consistent with property on the side of roads being wealthier than away from roads, which corresponds to the majority of properties in Kampala.

Figure A2: Spatial Coverage of the Owner Survey

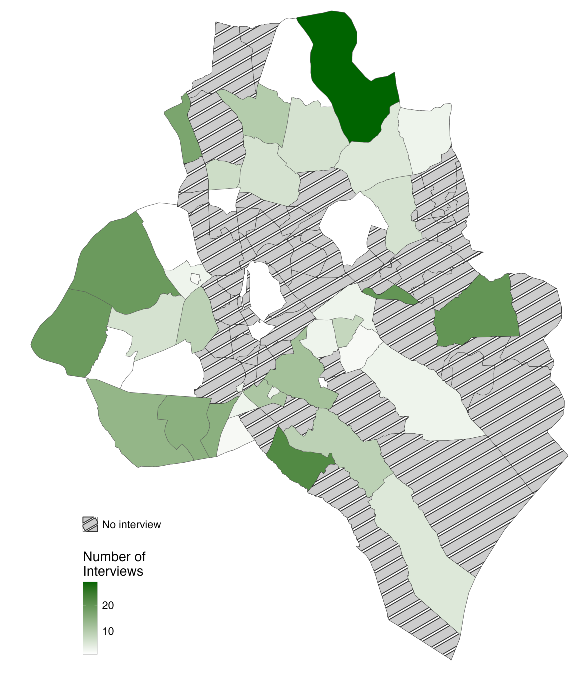

Notes: This map displays the number of owner interviews by parish. The sample is restricted to the owners in the main sample, i.e. who were officially contacted by the government about their land being affected by the road upgrade.

Table A4: Owner Survey: Descriptive Statistics

	Owner Survey				Census Data		
	Overall	Freehold	Public	Mailo	2014 Pop Census	2019 UNPS	2019 UNPS (owners)
Panel A: Sample Characteristics							
Number of owners	310	31	100	177			
Share on already upgraded roads	0.84	0.84	0.70	0.91			
Car ownership (share)	0.43	0.43	0.40	0.45		0.18	0.23
TV ownership (share)	0.97	0.97	0.98	0.97	0.68		
Bicycle ownership (share)	0.24	0.20	0.17	0.28	0.07	0.06	0.10
Computer ownership (share)	0.33	0.33	0.30	0.35	0.18		
Access to electricity (share)	0.97	1.00	0.99	0.96	0.84		
Bank account (share)	0.86	0.90	0.78	0.90	0.57		
Completed A-Level (share)	0.48	0.52	0.36	0.55	0.12		
Panel B: Property Characteristics							
Inherited property (share)	0.51	0.38	0.55	0.51			0.31
Owner since (year)	1989	1994	1987	1990			1994
Total land area (m2)	774.2	634.3	756.4	808.4			
Self-reported property value (USD)	154,796	204,626	141,539	154,744			
Cost Copy of Ownership Documents (USD)	1,718	1,764	1,616	1,778			
Floor: tiles (share)	0.50	0.45	0.51	0.50		0.13	0.15
Floor: cement screed (share)	0.46	0.52	0.46	0.45		0.78	0.77
Panel C: Land Acquisition							
Land Affected (m2)	77.7	48.3	60.6	92.1			
Initially Offered Compensation (share)	0.34	0.45	0.53	0.21			
Received Compensation by Time of Survey (share)	0.33	0.42	0.51	0.21			
Claimed Compensation (cond. no offer, share)	0.21	0.53	0.23	0.17			

Notes: Owner Survey columns cover the main analysis sample of affected owners the government officially reached out to (offered compensation or formally asked to consent), $N = 310$, split by robust land-tenure class (Freehold / Public / Mailo) assigned by spatial intersection of each interview's location; owners outside a mapped tenure polygon enter Overall only, so the tenure columns need not sum to Overall. *Census Data* columns report Kampala benchmarks: the 2014 Population Census (UBOS), the 2019 Uganda National Panel Survey (all households), and the 2019 UNPS restricted to owner-occupiers; blank cells indicate the indicator is unavailable in that source. "Claimed compensation" is conditional on not having been offered compensation. Area variables (land affected, total land area) and USD variables (property value, cost of an ownership-document copy) are trimmed at the top and bottom 5%; remaining rows are shares or means over the indicated (sub)sample.

B.3 Google Maps API Data

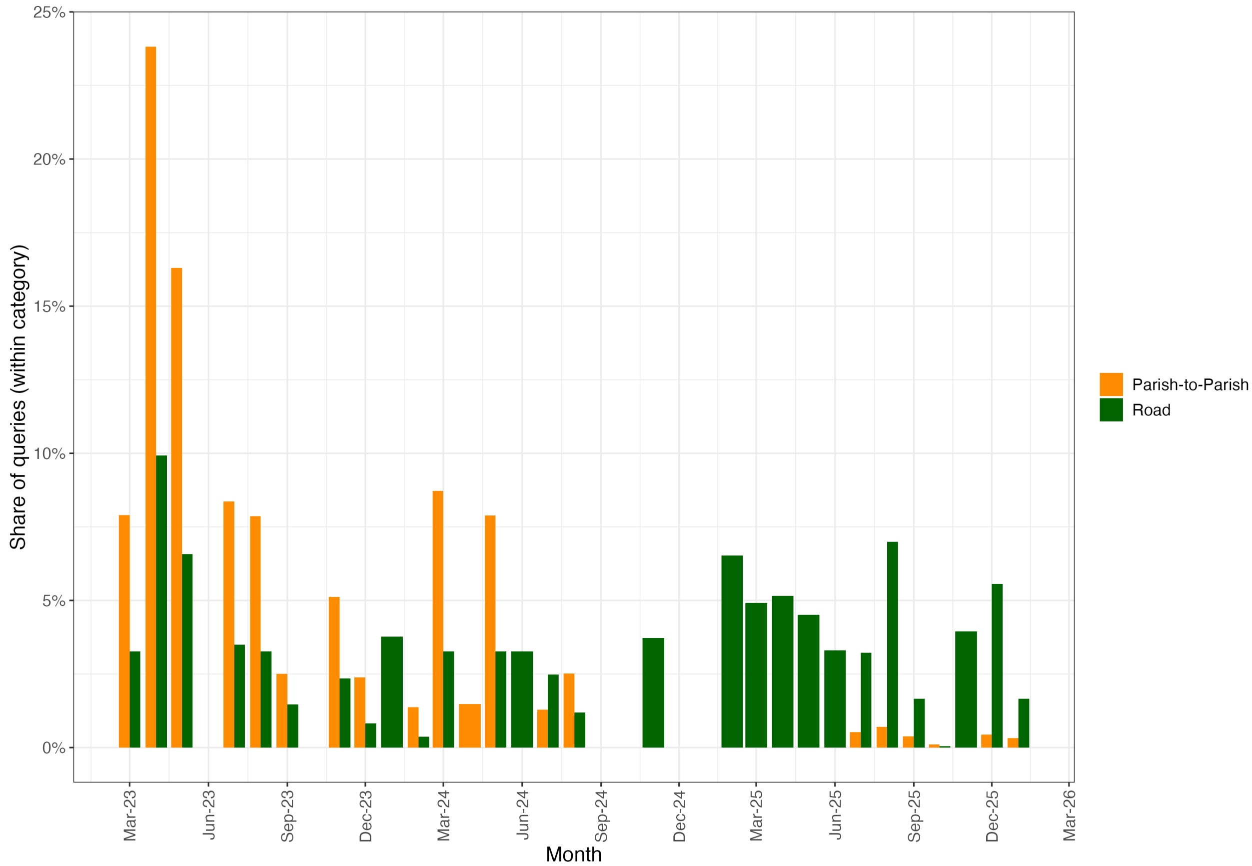

Figure A3: Google Maps API Queries Over Time 2023-2026

B.4 Ride-Hailing Data (2019-2024)

To study how workers will adjust their location choice in response to a change in commuting time, I partner with a local ride-hailing company running a popular app (Safeboda) with more than 17,000 rides per day.[40] I have access to the universe of motorcycle trips for a random sample of weeks between 2019 and 2024. As described in Appendix Table A5, my sample covers 119 days (5.9% of the period) and includes 2.2 million trips and more than 328,034 unique users (about 38.1% of the adult population in Kampala's 2014 Population Census). These users take an average of 6.9 trips on the app, but the standard deviation is very large (104.2), indicative of a subset of riders using the app daily, and consistent with 38.6% of users taking at least 5 trips on the app in my sample. The average trip lasts 20.7 minutes for a straight line distance of 5.1 kilometers. In addition, as displayed on Appendix Figure A4 for a randomly selected day in the data, most trips at morning commuting hours (6 am - 10 am) originate from the outskirts of the city to end at the city center. This commuting pattern towards the city center is consistent with the location of jobs from the Uganda Census of Business of Establishments displayed on Appendix Figure A14 Panel (b), but residents originating from the northern part of the city are overrepresented among morning commuters, compared to the universe of Kampala residents (Appendix Figure A14, Panel a).

Table A5: Ride-Hailing Descriptive Statistics

	Total		Per Month	
	Count / Mean	sd	Count / Mean	sd
Sample				
Number of Days Covered	119			
Share of Days Covered	0.06			
Trips				
Number of Trips	2,125,344		141,690	68,928
Length of Trip (min)	20.30	11.81	20.35	0.53
Straight Line distance (km)	5.08	2.82	5.18	0.40
Price (USD)	1.23	0.83	1.29	0.15
Users				
Number of Users	316,543		52,756	19,035
Avg Nb Trips per User	6.7	91.8	2.7	27.9
Share of users w/ more than 5 trips	0.38		0.15	

Notes: Descriptive statistics from the ride-hailing data. The sample covers a random sample of weeks between 2019 and 2024. Motorcycle rides only are included in the sample.

B.5 Travel Habit Surveys (2011 and 2016-2017)

The Travel Habit Survey (THS) was conducted between June 2016 and February 2017 as part of the Greater Kampala Metropolitan Area (GKMA) Multi-Modal Urban Transport Master Plan, commissioned by the Kampala Capital City Authority (KCCA) and financed by the World Bank. A representative sample of 844 households were interviewed. These interviews were combined with a pre-existing database

[40] Safeboda has been operating in Kampala since 2019, and dominates the app-based ride-hailing market for motorcycle rides.

Figure A4: Origin and Destination of Morning Ride-Hailing Trips

(a) Origin of Morning Trips (b) Destination of Morning Trips

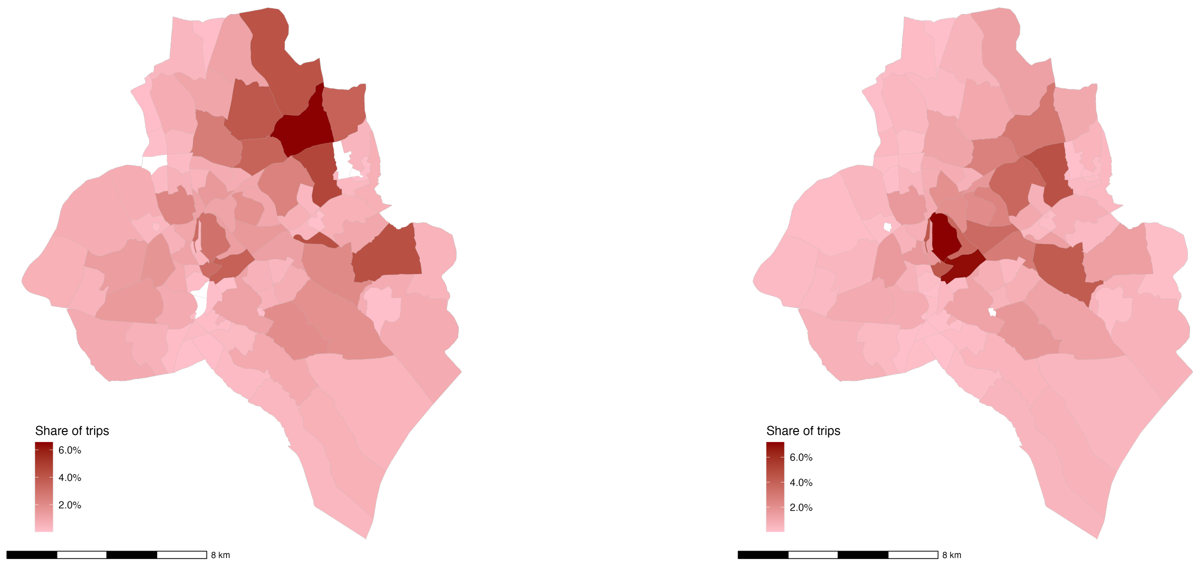

Notes: Parish-level shares of matched motorcycle (Safeboda) trip destinations on 2 July 2024 covering the morning rush (06:00–10:00). Panel (a) shows the share of these trips by parish of origin and panel (b) by parish of destination. The sample is restricted to trips with both origin and destination within Kampala. Darker shading indicates a larger share of trips starting or ending in a parish; parishes with no trips are left white.

of 2,654 households collected in 2011/2012 using the same methodology, for a total of 3,498 households and 28,323 recorded trips, including 5,130 home-work or work-home trips.[41]

For each trip, the survey recorded origin and destination (at the Traffic Analysis Zone level), departure times (both waves) and arrival times (2016/2017 wave only), transportation mode, trip purpose, and cost.

The main modes of transport for home to work trips are shared minibus *taxi* (56% of trips), private car (22%), motorcycle taxi or *boda-boda* (20%), and walking (3%).

[41]This includes 1,715 home-work or work-home trips cross-parishes (with positive distance) on the main commuting modes (walking, minibus taxis, motorcycle taxis, cars).

C Predicting Parish-Level Property Rates

I describe below how I recover pre-policy parish-level property sales rate and how I convert them into annual rental rates by leveraging transaction data from the broker survey.

Transaction-level preparation: From the broker survey, I recover the rate per m^2 for sold properties from information on price (USD) and parcel dimension (m^2). Raw rate per m^2 is only available for sold properties as rentals do not include information on parcel dimensions. These rates are trimmed at the top and bottom 2.5th percentile within each property class (residential, business) to remove outliers.

Recovering Pre-Upgrade Rates from Post-Upgrade Transactions: Sale transactions cover the 2018-2025 period. To recover pre-upgrade (2018) market rates – needed to compute the owner compensation and to calibrate the QSM –, I subtract the estimated treatment effect of road upgrade on parish-level property values from post-upgrade transactions. The pre-upgrade rate for properties transacted after T, corresponding to the parish's first road upgrade, is therefore $\log(\text{rate untreated}) = \log(\text{rate}_{t>T}) - \beta$, where β is the pooled OLS coefficient from the reduced-form impact regression (Table A14, col 3).[42]

Parish and Year Fixed Effects: I recover parish and year fixed-effects from a regression of transaction-level log rate per m^2 on parish α_i and year fixed effects δ_t. The predicted rate for parish i in year t is then defined as $rate_{i,t} = \exp(\alpha_i + \delta_t)$. Pre-policy parish-level rates are defined as $rate_{i,2018}$. For the 15 (out of 96) parishes without transaction data, I impute the parish fixed-effect from an inverse-distance weighting (IDW) approach from the $k \in \{3, 10\}$ nearest parishes with observed data (Shepard, 1968). The IDW power parameter is selected via leave-one-out cross-validation to minimize LOOCV RMSE.

From sales market rate to rental market rate: The QSM requires per-period rental rates paid by households and firms to landlords. As data on properties' size is only available for sales, to convert sales rate per m^2 into rental rate per m^2, I estimate the capitalization rate by pooling all sale and rental transactions and regressing log(price) on a rental dummy, controlling for property characteristics common to both sale and rental transactions (number of rooms, number of bathrooms (residential only), construction type (formal, temporary, floor type), distance to the nearest road), with parish and year fixed effects. The coefficient on the rental dummy captures the log monthly rent-to-price ratio; annualizing gives the capitalization rate for residential (6.1%, P/E = 16.4) and business (9.7%, P/E = 10.3) properties.

Integrating over the implemented road upgrades: To recover the market value of affected land from the realized road upgrades, I multiplied market rates in each parish by the length of road upgrades in each parish, assuming that each road upgrade comes with land acquisition equivalent to one additional lane (3.4 meters).

[42] I use the pooled OLS rather than the TWFE coefficients because only a subset of parishes exhibit variation in treatment status during the period covered by the transaction data. I further use a single β across residential and business properties.

D Appendix and Robustness on Land Costs Facts

Figure A5: Owners' Reasons to Consent to Voluntary Land Take Without Compensation

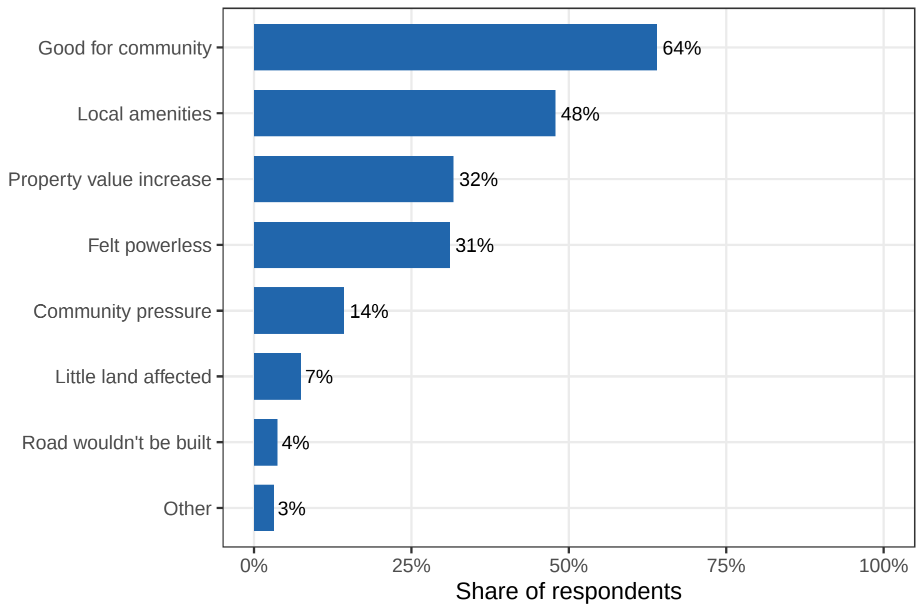

Notes: Reasons (multiple-choice) given by the 161 owners in the survey who were not offered compensation and decided not to claim it.

Figure A6: Parish-Level Predicted Probability of Claiming Compensation Across Specifications

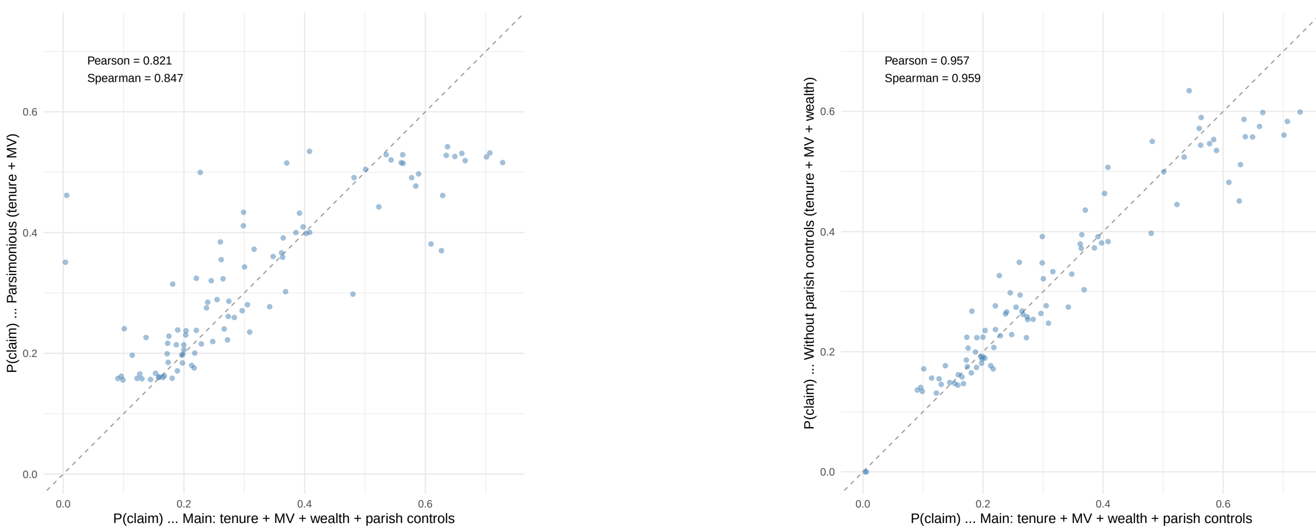

(a) Main vs. parsimonious (tenure + MV)

(b) Main vs. without parish controls

Notes: Each dot is a parish. The horizontal axis is the predicted probability of claiming compensation $\mathbb{P}_i(\text{claim})$ from the *main* specification used in the model (tenure shares + market value + wealth + parish controls); the vertical axis is an alternative specification — the parsimonious tenure + market value spec (panel a) and the main spec without parish controls (panel b). Estimates are the MLE logit estimates from the claim columns of Table 1. Parish-level predictions evaluate the estimated logit at parish-level tenure shares (KCCA land tenure map), parish-level wealth (UBOS census), and parish controls, with remaining owner-level variables held at their sample means. The dashed line is the 45-degree line; Pearson and Spearman correlations are shown on each panel. I assume 37 m^2 of land (the survey median) is taken from each owner.

Table A6: Mechanisms Driving Claiming Decisions

	(1)	(2)	(3)	(4)
	Claimed compensation (dummy)			
Wealth index (std)	0.07			0.41
	(0.17)			(0.21)
Above-median doc. cost (d)		-0.79		-0.37
		(0.37)		(0.39)
Nb. friends in neighborhood (std)			-1.00	-1.17
			(0.50)	(0.54)
Estimator	Logit	Logit	Logit	Logit
Observations	212	212	212	212
Mean dep. var.	0.226	0.226	0.226	0.226

Notes: Each column reports maximum likelihood estimates of a logit model for the probability of claiming compensation among owners not initially offered financial compensation. The dependent variable equals one if the owner filed a claim. "Above-median doc. cost" is a dummy for self-reported costs of obtaining ownership documents exceeding the sample median. "Land security" is a standardized 4-point scale response to: "Would you say that landowners are secure from eviction or expropriation by the government?" "Wealth index" is a standardized mean of eight binary household indicators (completed A-level education, electricity, permanent floor, bank account, television, computer, bicycle, ≥2 meals/day). Coefficients are in log-odds; standard errors in parentheses. Missing values handled via mean replacement with a missing-indicator dummy.

Table A7: Land Tenure Mechanisms

	(1)	(2)	(3)
	Wealth index (std)	Above-med doc. cost (d)	Nb. friends (std)
Freehold	-0.227	-0.192	-0.280
	(0.111)	(0.082)	(0.081)
Public	-0.552	-0.067	-0.296
	(0.096)	(0.056)	(0.077)
Estimator	OLS	OLS	OLS
Observations	478	415	458
Mean dep. var.	0.002	0.520	-0.023

Notes: Each column reports OLS estimates with heteroskedasticity-robust standard errors in parentheses. The reference tenure category is mailo. The sample is survey respondents who were not initially offered financial compensation "Above-med doc. cost" is a dummy for self-reported costs of obtaining ownership documents exceeding the sample median. "Land security" is a standardized 4-point scale of perceived security from government expropriation. "Wealth" is a standardized mean of eight binary household indicators. Even-numbered columns add individual-level SES controls (education, language, bank account, formal credit, employment, household head, nationality).

Figure A7: Robustness: Tenure Coefficient for the Probability of Being Offered or Claim Compensation

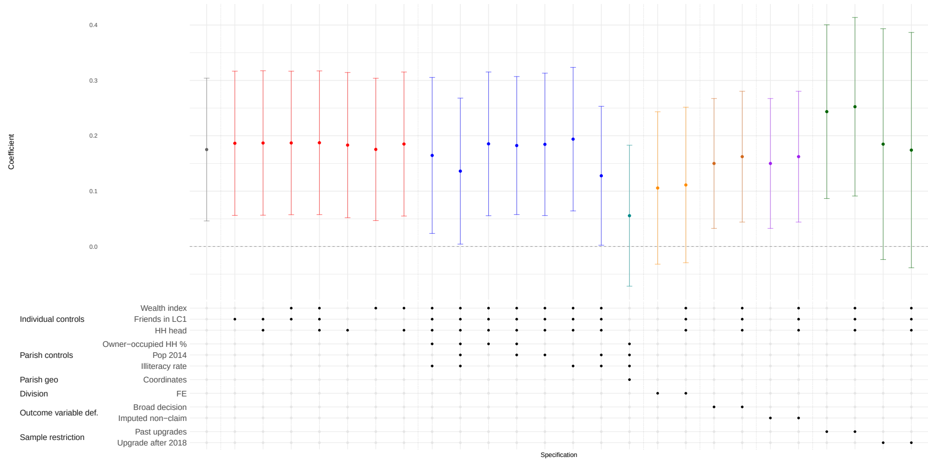

(a) Freehold Coef (Proposed, OLS)

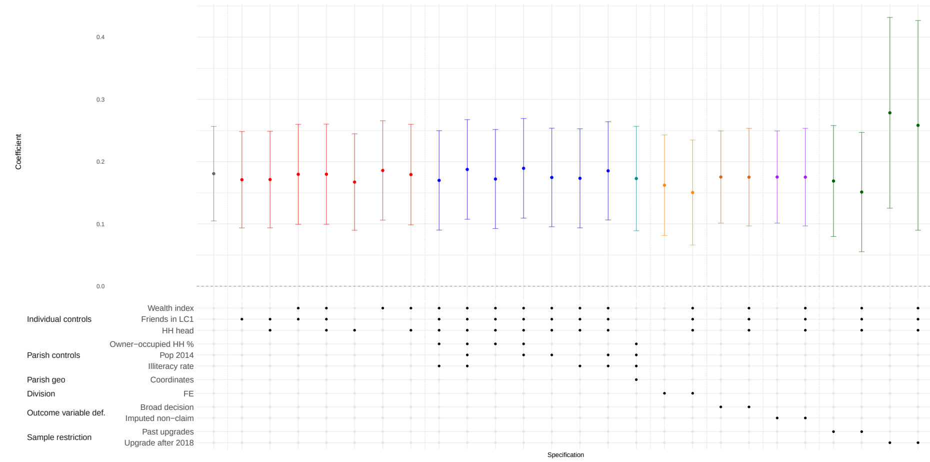

(b) Public Coef (Proposed, OLS)

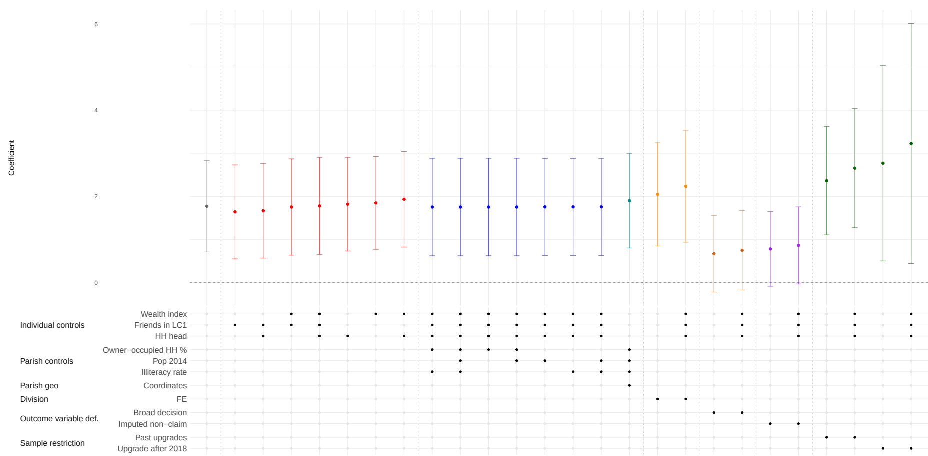

(c) Freehold Coef (Claimed, MLE)

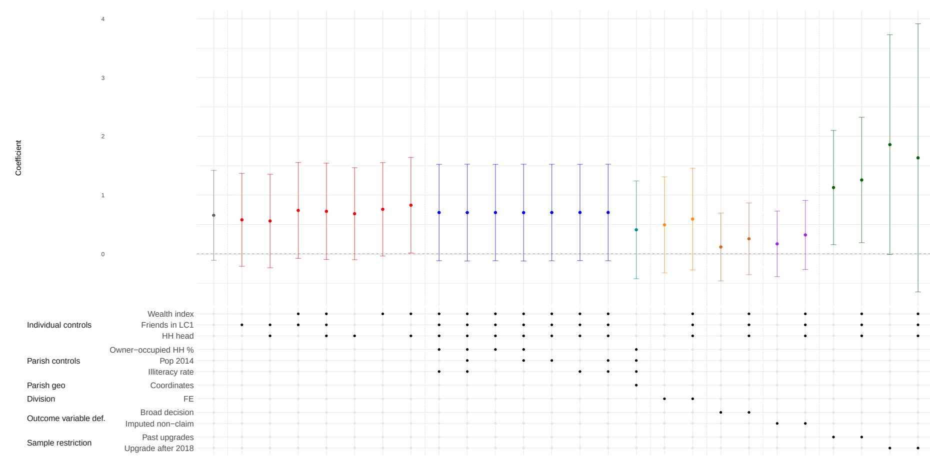

(d) Public Coef (Claimed, MLE)

Notes: Each panel shows specification curves for a tenure coefficient across alternative specifications of the OLS regression of receiving compensation (top panels) and the MLE logit model for the probability of claiming compensation among owners not initially offered one (bottom panels). Point estimates are shown with 95% confidence intervals and specifications are ordered similarly across panels. In red, I add combinations of individual-level controls to show that the tenure estimates are not driven by correlated observables. In blue, I control for parish-level socioeconomic characteristics and in cyan I add flexible controls for parish coordinates. In orange, I include division fixed effects. In brown, I expand the sample to owners who reported not being officially approached by government officials about their land being affected by the road upgrade. In purple, I assume that owners with no recorded claim decision — a subset of respondents from the first wave of data collection who reported not being officially approached by the government despite their land being affected — did not claim and were not offered compensation. In green, I restrict the sample to respondents who reported that the road upgrade was either already completed or completed after 2018.

E Appendix and Robustness on Road Upgrade Benefits

Road scarcity (context for ξ): Table A8 places Kampala's road network in comparative perspective: its exceptionally low road-area share and narrow streets are what make a high speed–capacity elasticity a sensible response to relieving an unusually binding capacity constraint, as discussed in the main text.

Table A8: Kampala's roads are scarce and narrow, even by Sub-Saharan African standards

City	Roads' share of built-up area (%)		Avg. street width (m)		Roads <4 m wide (%)		Roads >16 m wide (%)	
	Pre 1990	1990–2014	Pre 1990	1990–2014	Pre 1990	1990–2014	Pre 1990	1990–2014
Kampala	**13**	**12**	**6.7**	**4.5**	**20**	**42**	**4**	**1**
Other Sub-Saharan African cities								
Lagos	17	16	10.1	7.1	6	20	10	3
Kinshasa	14	13	9.5	5.2	28	37	4	3
Kigali	17	14	7.9	5.5	18	32	7	1
Nakuru	24	21	10.8	5.5	14	31	20	2
Johannesburg	25	18	13.2	7.4	6	22	32	7
High-income comparison cities								
Santiago	25	18	12.6	7.9	5	16	27	10
Chicago	27	25	11.6	10.0	8	27	42	30
Paris	21	15	9.2	6.2	10	28	11	5
London	19	10	9.5	7.5	9	18	9	4

Notes: Road-layout metrics from the Atlas of Urban Expansion (Angel et al., 2016), comparing areas built before 1990 with those built 1990–2014. Across the Atlas's global sample of 200 cities, roads occupy on average 21% of built-up area worldwide and about 16% in Sub-Saharan Africa. Kampala combines an unusually low road-area share, the narrowest post-1990 streets in this comparison, a rising share of sub-4 m roads, and almost no roads wide enough (>16 m) to carry fast through-traffic—a profile of exceptionally scarce road capacity.

Speed: controls, horizon, and estimator. The estimated impact of road upgrade on local speed is robust to a wide range of specifications, including to adding geographic controls, splitting rush from off-peak, and replacing the twelve-month window with the un-windowed Sun and Abraham (2021) ATT averaged over all post-completion periods, as shown in Table A9. Table A10 reports the pooled all-hours effect across the same set of estimators, of which the main table's rush/off-peak split (Table 2) is a decomposition.

Speed: selection into timing.

Speed: spillovers onto controls (SUTVA). The estimated impact of road upgrade on local speed would be biased if speed on control roads was affected through traffic diversion either onto nearby unimproved roads, or off them onto the faster route. I show that this channel is unlikely to drive the estimates by building two measures of each control road's exposure to nearby completed upgrade: an indicator and a count of completed policy roads within 1–5 km. I then implement two exercises, presented in Appendix Table A12.

Table A9: Road Upgrades and Traffic Speed: Robustness to Controls and Estimation Horizon

	Rush (morning + evening)			Off-peak		
	(1)	(2)	(3)	(4)	(5)	(6)
Upgrade Completed (d)	2.762	4.526	3.376	3.115	6.109	1.328
	(1.286)	(1.152)	(1.286)	(1.655)	(1.555)	(1.214)
Upgrade In Progress (d)	0.335	0.891	-1.029	-0.130	0.726	-0.736
	(1.035)	(0.938)	(0.730)	(1.153)	(0.906)	(0.564)
Day FE	Yes	Yes		Yes	Yes	
Hour FE	Yes	Yes		Yes	Yes	
Month FE			Yes			Yes
Road FE			Yes			Yes
Observations	1,017	1,017	752	1,075	1,075	684
Geo controls		Yes			Yes	
Model	OLS	OLS	SA	OLS	OLS	SA
SA leave-one-cohort-out			[1.6, 4.9]			[0.6, 1.7]

Notes: Standard errors in parentheses. Standard errors are two-way clustered by road segment and day in the OLS columns (1, 2, 4, 5) and by road segment in the Sun–Abraham columns (3, 6). The dependent variable is traffic speed (km/h) from Google Maps API queries for trips along policy roads. Columns (1)–(3) report estimates for rush hours (morning and evening peaks pooled); columns (4)–(6) for off-peak hours. Columns (1)/(4) and (2)/(5) are trip-level OLS with day and hour-of-day fixed effects: columns (1)/(4) include no controls, and columns (2)/(5) add road distance and geographic controls (longitude, latitude, and their squares). Columns (3)/(6) report the average treatment effect on the treated from the Sun and Abraham (2021) interaction-weighted estimator, estimated on the road × month panel with road-segment and month fixed effects (the outcome residualized on day and hour-of-day fixed effects before aggregation) and averaged over *all* post-completion periods — the un-windowed analogue of the main-table specification. The "SA leave-one-cohort-out" row reports the range of this full-horizon ATT as each completion-month cohort is dropped in turn from the aggregation.

Table A10: Road Upgrades and Average Traffic Speed (All Hours)

	Trip-level				Road × Month Panel	
	(1)	(2)	(3)	(4)	(5)	(6)
Upgrade Completed (d)	2.904	5.238	5.215	3.613	2.912	1.387
	(1.413)	(1.300)	(0.884)	(0.910)	(0.920)	(1.073)
Upgrade In Progress (d)	0.051	0.732	0.540	0.451	0.300	-1.266
	(1.066)	(0.897)	(0.754)	(0.706)	(0.634)	(0.634)
Day FE	Yes	Yes	Yes	Yes		
Hour FE	Yes	Yes	Yes	Yes		
Month FE					Yes	Yes
Road FE				Yes	Yes	Yes
Observations	2,092	2,092	2,092	2,092	1,044	1,044
Mean dep. var.	22.67	22.67	22.67	22.67	22.71	22.71
Geo controls		Yes	Yes			
Road class			Yes			
Paved (baseline)			Yes			
Model	OLS	OLS	OLS	TWFE	TWFE	SA
SA leave-one-cohort-out						[0.1, 2.4]

Notes: Standard errors in parentheses. Standard errors are two-way clustered by road segment and day in columns (1)–(4) and by road segment in columns (5)–(6). The dependent variable is traffic speed (km/h) from Google Maps API queries for trips along policy roads, pooling *all* hours of the day; this table therefore reports the average effect across rush and off-peak periods (decomposed in Table 2). Columns (1)–(4) are estimated at the trip level and columns (5)–(6) at the road × month panel level, where the dependent variable is the road–month average of trip-level speed residualized on day and hour-of-day fixed effects; the reported mean dependent variable is the trip-level mean in columns (1)–(4) and the road–month mean in columns (5)–(6). Columns (1)–(3) are OLS with day and hour-of-day fixed effects, adding controls progressively: column (1) none, column (2) road distance and geographic controls (longitude, latitude, and their squares), and column (3) further road-class dummies and baseline paving status. Column (4) is a trip-level TWFE specification with road-segment, day, and hour-of-day fixed effects. Column (5) is a panel TWFE with road-segment and month fixed effects. Column (6) reports the average treatment effect on the treated over the first twelve months following completion (relative periods 0–11) from the Sun and Abraham (2021) interaction-weighted estimator, with delta-method standard errors; the "SA leave-one-cohort-out" row reports its range as each completion-month cohort is dropped in turn.

First, I conduct a placebo on the controls: on road-months not yet treated, I regress residualized rush-hour speed on exposure, with road and month fixed effects. I show that exposure does not consistently predict

Table A11: Road Upgrade Timing is Not Predicted by Observable Characteristics

	Start Construction (month)		Completion (month)	
	(1)	(2)	(3)	(4)
Log Road Length	0.0721	0.7248	10.40	8.859
	(6.993)	(8.133)	(4.598)	(4.965)
Paved baseline	20.24	16.03	19.52	14.30
	(19.92)	(21.02)	(13.40)	(13.34)
Road Class: Primary	26.34	28.41	1.207	-3.518
	(32.81)	(32.50)	(17.93)	(18.29)
Road Class: Secondary	29.69	18.98	11.52	0.7436
	(17.43)	(18.38)	(12.89)	(13.11)
Road Class: Tertiary	35.88	17.09	20.72	9.541
	(17.80)	(20.54)	(14.27)	(15.03)
Longitude		-329,647.9		-298,092.9
		(457,898.9)		(344,834.0)
Longitude sqrd		5,059.8		4,579.2
		(7,027.4)		(5,292.6)
Latitude		903.0		1,780.8
		(3,715.3)		(2,407.1)
Latitude sqrd		-2,036.1		-3,243.9
		(5,688.8)		(3,686.4)
Observations	40	40	51	51
R^2	0.19959	0.32869	0.21812	0.34110
F-test	1.6956	1.6321	2.5107	2.3584

Notes: Each column reports an OLS regression of upgrade timing (in months) on road-level characteristics. The dependent variable in columns (1)–(2) is the month construction started; in columns (3)–(4), the month construction was completed. Columns (2) and (4) add geographic controls (longitude, latitude, and their squares). The sample in columns (1)–(2) is restricted to roads with observed (non-imputed) start dates. Road class dummies are relative to non-major roads.

control speeds, consistent with no contamination. At 3 km an additional completed neighbor is worth -0.352 km/h (s.e. 0.544), or -8.3% of the treatment effect and insignificant, and the estimate is small at 1–3 km and flips sign across different measures of exposure. The estimated coefficient, grows with distance and is marginally significant at 5 km (-0.72, s.e. 0.42). Yet, this pattern is not suggestive of credible spillovers, which would be local and fade with distance. Instead, it is likely the sign of the neighbor count saturating and absorbing citywide trend rather than local diversion as, by 5 km nearly every control already has a completed neighbor, leaving the binary measure no within-road variation.

Second, and most importantly, adding exposure to the main specification barely moves the effect: the completion coefficient goes from 4.235 km/h to 4.067 km/h (-4%), and is stable at every radius.

Commuting: take-up of upgraded roads. In Table A13 I show that upgrades also draw commuters onto the improved segments. I estimate this take-up effect on both the number of trips passing through a segment and its share of all trips, and on the post-April-2023 subsample for which both pre- and post-upgrade months are observed; the effect is stable in sign and significance across the two outcomes and samples.

Property values: stated-preference and amenity evidence. Figure A8 shows that the distribution of brokers' hypothetical post-upgrade price ratios underlying the appraisal is right-tailed for both residential and business premises. Figure A9 documents that upgrades also improved local amenities, so that its

Table A12: Spillovers onto Control Roads: Direct Placebo and Exposure Control

	Reg A: $\hat{\zeta}$ (spillover on controls)		Reg B: $\hat{\beta}$ (main effect, w/ exposure)	
	Binary	Count	Binary	Count
1 km	-0.225	0.319	4.226	3.954
	(0.459)	(0.716)	(1.043)	(1.023)
2 km	-0.225	0.081	4.147	3.876
	(0.459)	(0.707)	(1.000)	(1.052)
3 km	–	-0.352	4.235	4.067
		(0.544)	(1.020)	(1.037)
5 km	–	-0.718	4.235	4.168
		(0.415)	(1.020)	(1.006)

Notes: Evidence on the no-spillover (SUTVA) assumption for the rush-hour speed DiD. *Exposure* is the number of policy-road upgrades completed within the stated radius of a road by month t (Count), or an indicator for at least one (Binary). Reg A regresses residualized rush-hour speed on exposure, restricting to road-months not yet treated (own completion and own in-progress indicators both zero), with road and month fixed effects; $\hat{\zeta}$ is the effect of a nearby completion on a still-control road's speed. Reg B adds exposure to the main two-way fixed-effects specification on the full sample; $\hat{\beta}$ is the coefficient on own completion, to be compared with the no-exposure main estimate of 4.235 km/h. Standard errors clustered by road in parentheses. Binary exposure is collinear with the fixed effects at 3–5 km and those cells are empty; inverse-distance weighting is omitted because touching road geometries produce zero pairwise distances. Reg A includes still-control road-months only. Binary exposure is collinear with the FE at 3-5 km as all control roads have a treated road within this radius at some point.

Table A13: Upgraded Roads are More Likely to Be Used by Commuters

	Number of Trips		Share of Trips	
	(1)	(2)	(3)	(4)
Upgrade Completed (d)	109.5	221.0	0.0059	0.0058
	(66.10)	(137.5)	(0.0018)	(0.0036)
Upgrade in progress (d)	9.975	62.11	-0.0005	-0.0012
	(43.43)	(75.90)	(0.0018)	(0.0027)
Month	Yes	Yes	Yes	Yes
Road-segment	Yes	Yes	Yes	Yes
Observations	10,808	3,360	10,808	3,360
R^2	0.53226	0.55577	0.69088	0.68959
Mean Dep Var	173	175	0.0199	0.0213
Trips Sample	Parish-to-Parish	Parish-to-Parish	Parish-to-Parish	Parish-to-Parish
Sample	All Policy Roads	Upgraded Post 04-23	All Policy Roads	Upgraded Post 04-23
Unit	Road-Segment x Month	Road-Segment x Month	Road-Segment x Month	Road-Segment x Month

Notes: Each column reports a TWFE regression at the road-segment × month level. The dependent variable in columns (1)–(2) is the number of parish-to-parish Google Maps trips passing through the road segment in a given month; in columns (3)–(4), the share of all parish-to-parish trips. Columns (1) and (3) use all policy roads in the Google Maps sample while columns (2) and (4) restrict the sample to roads upgraded after April 2023, e.g. for which we have pre and post-upgrade observations. All specifications include road-segment and month fixed effects. Standard errors are clustered by road segment and month.

accessibility-only welfare gains understate the full local capitalization by the amenity component.

Figure A8: Distribution of brokers' hypothetical post-upgrade price ratio.

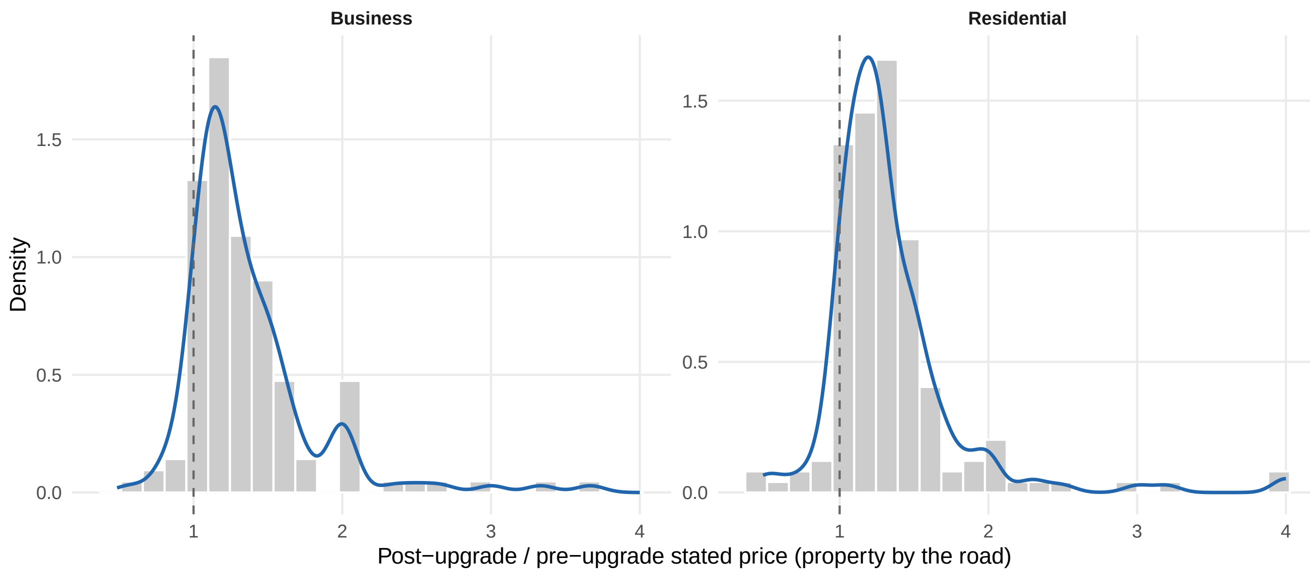

Notes: Across brokers, the ratio of the stated sale price of a standard property *after* a hypothetical road upgrade to its price *before*, for a property next to the road. Left: residential appraisals; right: business. The dashed line marks a ratio of one (no change); mass to the right indicates an expected price gain. The median post-pre upgrade price ratio is 1.25 for residential properties and 1.22 for business properties, consistent with a right-tailed distribution of answers.

Figure A9: Self-Assessed Road Quality Before and After Upgrade

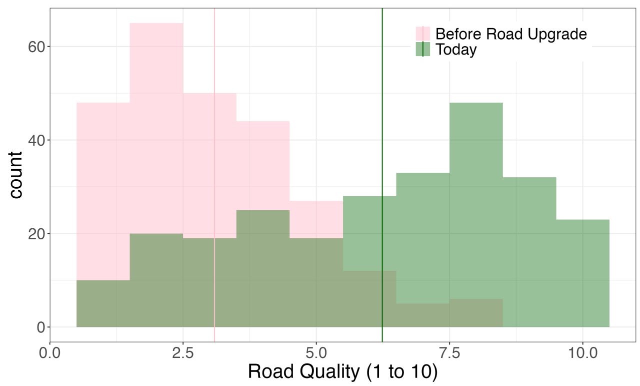

(a) Overall

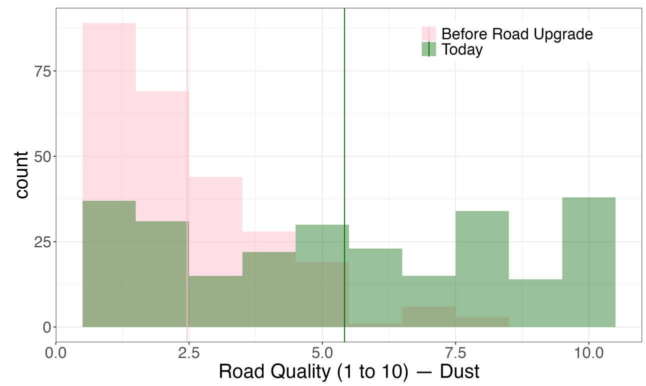

(b) Dust

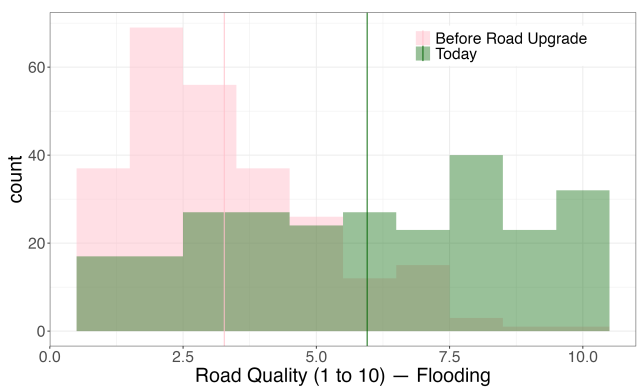

(c) Flooding

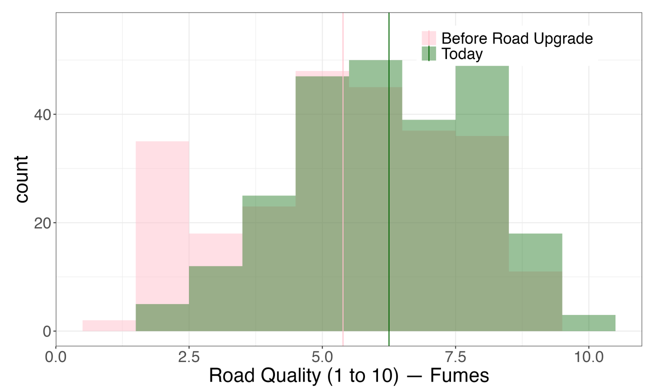

(d) Fumes

Notes: Self-assessed road quality on a 1–10 Likert scale, reported by affected property owners who experienced a past road upgrade ($N = 259$). Pink: recalled quality before the upgrade; green: quality today. Vertical lines indicate group means. Mean improvement: overall 3.1, dust 3, flooding −4.6, fumes 8.2 points.

Table A14: Road Upgrades Increased Local Property Values (Transaction-Based Evidence)

	(1)	(2)	(3)	(4)	(5)	(6)
Road upgrade (ATT)	0.107	0.187	0.095	0.024	0.145	0.284
	(0.056)	(0.108)	(0.085)	(0.104)	(0.302)	(0.118)
Outcome	Rent	Rent	Sale	Sale	Sale	Sale
Property class	All	All	All	All	All	All
Parish sample	All	All	All	All	CS-eligible	CS-eligible
Geo FE	Sub-county	Parish	Sub-county	Parish	Parish	Parish
Time FE	Year	Year	Year	Year	Year	Year
Estimator	Pooled OLS	TWFE	Pooled OLS	TWFE	CS	SA
Observations	2,044	2,044	1,619	1,619	1,547	1,547

Notes: Each column reports the estimated effect of a parish-level road upgrade on log property prices. The dependent variable is the residual from a first-stage regression of log prices on property characteristics (floor type, number of rooms, building type, land dimensions for sales) and month fixed effects. Columns (1)–(2) use rental transactions while columns (3)–(6) use sale transactions. Columns (1) and (3) include sub-county fixed effects and parish-level socio-economic controls; columns (2) and (4) include parish fixed effects (TWFE). Columns (5)–(6) implement the heterogeneity-robust staggered DiD estimators of Callaway and Sant'Anna (2021) and Sun and Abraham (2021), with parish fixed effects. All prices are trimmed at the top and bottom 5th percentiles. Standard errors are two-way clustered by broker and parish for the OLS specifications (columns 1–4); the Callaway and Sant'Anna estimator clusters by parish.

F Model Appendix

F.1 Optimal Road Improvement: Lagrange Problem

The government problem of optimal road widening $\{\Delta R_{kl}\}$ in general equilibrium can be rewritten as the following optimization problem with shadow constraints for all the endogenous variables of the model and corresponding Lagrange multipliers:[43]

$$
\begin{aligned}
LL =& \frac{1}{\delta}\left[\bar{L}W^w + \sum_i \left(N_i^{of}W_i^{of} + N_i^{or}W_i^{or}\right) - \eta \sum_i \sum_m p_i^{land}\frac{l_{im}}{2}\Delta R_{im}\right] - \bar{p}\sum_i \sum_{m \in C_i} \frac{l_{im}}{2}\Delta R_{im} \\
& [1] - \gamma^W \left[W^w - f_W(q^r_{\forall i}, \tau_{ij \forall i,j}, w_{i \forall i})\right] \\
& [N] - Z^{W_i^{of}}\left[W_i^{of} - f_{W_i^{of}}(l_i^{of})\right] \\
& [N] - Z^{W_i^{or}}\left[W_i^{or} - f_{W_i^{or}}(l_i^{or})\right] \\
& [N \times N] - \sum_{ij} \lambda_{ij}^{L_{ij}}\left[L_{ij} - f_{L_{ij}}\left(W^w, q_i^r, \tau_{ij}, w_j\right)\right] \\
& [N] - \sum_i \eta_i^{q_i^r}\left[q_i^r - f_{q_i^r}\left(w_{j \forall j}, L_{ij,\forall j}, \tau_{ij \forall j}, H_i^r\right)\right] \\
& [N] - \sum_i \eta_i^{q_i^f}\left[q_i^f - f_{q_i^f}\left(L_{ij \forall i}, H_j^r, w_j\right)\right] \\
& [N] - \sum_i \eta_i^{w_i}\left[w_i - f_{w_i}\left(H_j^f, L_{ij \forall i}\right)\right] \\
& [N \times N] - \sum_{ij} \mu_{ij}^{\tau_{ij}}\left[\tau_{ij} - f_{\tau_{ij}}\left(d_{kl,\forall kl}\right)\right] \\
& [E] - \sum_{kl} \psi_{kl}\left[d_{kl} - f_{d_{kl}}\left(R_{im}\right)\right] \\
& [N] - \sum_i \phi_i^{H_i^r}\left[H_i^r - f_{H_i^r}\left(R_{im,\forall m}\right)\right] \\
& [N] - \sum_i \phi_i^{H_i^r}\left[H_i^r - f_{H_i^r}\left(R_{im,\forall m}\right)\right] \\
& [N] - \sum_i \zeta_i^{l_i^{of}}\left[l_i^{of} - f_{l_i^{of}}\left(q_i^r, q_i^r\right)\right] \\
& [N] - \sum_i \zeta_i^{l_i^{or}}\left[l_i^{or} - f_{l_i^{or}}\left(q_i^r, q_i^r, H_i^r, H_i^r, R_{im,\forall m}, p_i^{\text{land}}\right)\right] \\
& [N] - \sum_i \Phi_i^{p_i^{land}}\left[p_i^{\text{land}} - f_{p_i^{land}}\left(R_{im,\forall m}\right)\right] \\
& [1] - \lambda\left[\bar{M} - \bar{p}\sum_i \sum_{m \in C_i} \frac{l_{im}}{2}\Delta R_{im}\right].
\end{aligned}
$$

Taking the First Order Conditions w.r.t. the optimal policy $\{\Delta R_{kl}\}$ and all endogenous variables gives a

[43] Throughout the appendix, owners' utility is logarithmic ($\psi^o = 1$ in equation (8)), so $\partial W_i^o / \partial l_i^o = \bar{C}^o / l_i^o$. Both welfare and the land payment are per-period flows — compensation is the rent-equivalent of the taken land and enters owners' rental income each period — so they enter the objective in present value ($\frac{1}{\delta}$ times the flow); only construction is paid upfront.

system of $(2 \times E + 2 \times N^2 + 10 \times N + 1)$ equations:

$$
\begin{aligned}
0 =& -\frac{l_{kl}}{2}\left(\frac{\eta}{\delta}\left(p_k^{land} + p_l^{land}\right) + (1+\lambda)\bar{p}\right) + \psi_{kl}\frac{\partial d_{kl}}{\partial \Delta R_{kl}} + \psi_{lk}\frac{\partial d_{lk}}{\partial \Delta R_{kl}} + \phi_k^{H_k^r}\frac{\partial H_k^r}{\partial \Delta R_{kl}} + \phi_l^{H_l^r}\frac{\partial H_l^r}{\partial \Delta R_{kl}} \\
& + \phi_k^{H_k^f}\frac{\partial H_k^f}{\partial \Delta R_{kl}} + \phi_l^{H_l^f}\frac{\partial H_l^f}{\partial \Delta R_{kl}} + \zeta_k^{l_k^{or}}\frac{\partial l_k^{or}}{\partial \Delta R_{kl}} + \zeta_l^{l_l^{or}}\frac{\partial l_l^{or}}{\partial \Delta R_{kl}} + \Phi_k^{p_k^{land}}\frac{\partial p_k^{land}}{\partial \Delta R_{kl}} + \Phi_l^{p_l^{land}}\frac{\partial p_l^{land}}{\partial \Delta R_{kl}} && \Delta R_{kl}[E] \\
0 =& \frac{1}{\delta}\bar{L} - \gamma^W - \theta\sum_{ij}\lambda_{ij}^{L_{ij}}\frac{L_{ij}}{W^w} && \Delta W^w[1] \\
0 =& \frac{1}{\delta}N_i^{of} - Z_i^{W_i^{of}} && W_i^{of}[N] \\
0 =& \frac{1}{\delta}N_i^{or} - Z_i^{W_i^{or}} && W_i^{or}[N] \\
0 =& -\phi_i^{H_i^r} + \eta_i^{q_i^r}\frac{\partial q_i^r}{\partial H_i^r} + \zeta_i^{l_i^{or}}\frac{\partial l_i^{or}}{\partial H_i^r} && H_i^r[N] \\
0 =& -\phi_i^{H_i^f} + \eta_i^{q_i^f}\frac{\partial q_i^f}{\partial H_i^f} + \eta_i^{w_i}\frac{\partial w_i}{\partial H_i^f} + \zeta_i^{l_i^{or}}\frac{\partial l_i^{or}}{\partial H_i^f} && H_i^f[N] \\
0 =& -\lambda_{ij}^{L_{ij}} + \sum_i \eta_i^{q_i^r}\frac{\partial q_i^r}{\partial L_{ij}} + \sum_i \eta_i^{q_i^r}\frac{\partial q_i^r}{\partial L_{ij}} + \sum_i \eta_i^{w_i}\frac{\partial w_i}{\partial L_{ij}} && L_{ij}[N^2] \\
0 =& -\eta_i^{q_i^r} + \gamma^W\frac{\partial W}{\partial q_i^r} + \sum \lambda_{ij}^{L_{ij}}\frac{\partial L_{ij}}{\partial q_i^r} + \zeta_i^{l_i^{of}}\frac{\partial l_i^{of}}{\partial q_i^r} + \zeta_i^{l_i^{or}}\frac{\partial l_i^{or}}{\partial q_i^r} && q_i^r[N] \\
0 =& -\eta_i^{q_i^f} + \zeta_i^{l_i^{of}}\frac{\partial l_i^{of}}{\partial q_i^f} + \zeta_i^{l_i^{or}}\frac{\partial l_i^{or}}{\partial q_i^f} && q_i^f[N] \\
0 =& -\eta_j^{w_j} + \gamma^W\frac{\partial W}{\partial w_i} + \sum_m \lambda_{mi}^{L_{mi}}\frac{\partial L_{mi}}{\partial w_i} + \sum_m \eta_m^{q_m^r}\frac{\partial q_m^r}{\partial w_i} + \eta_i^{q_i^r}\frac{\partial q_i^r}{\partial w_i} && w_i[N] \\
0 =& -\zeta_i^{l_i^{of}} + Z^{W_i^{of}}\frac{\partial W_i^{of}}{\partial l_i^{of}} && l_i^{of}[N] \\
0 =& -\zeta_i^{l_i^{or}} + Z^{W_i^{or}}\frac{\partial W_i^{or}}{\partial l_i^{or}} && l_i^{or}[N] \\
0 =& -\Phi_i^{p_i^{land}} - \eta\sum_{m\in C_i}\frac{l_{im}}{2}\Delta R_{im} + \zeta_i^{l_i^{or}}\frac{\partial l_i^{or}}{\partial p_i^{land}} && p_i^{\text{land}}[N] \\
0 =& -\mu_{ij}^{\tau_{ij}} + \gamma^W\frac{\partial W}{\partial \tau_{ij}} + \lambda_{ij}^{L_{ij}}\frac{\partial L_{ij}}{\partial \tau_{ij}} + \eta_i^{q_i^r}\frac{\partial q_i^r}{\partial \tau_{ij}} && \tau_{ij}[N^2] \\
0 =& -\psi_{kl} + \sum_{ij}\mu_{ij}^{\tau_{ij}}\frac{\partial \tau_{ij}}{\partial d_{kl}} && d_{kl}[E]
\end{aligned}
$$

With some algebra, we can then further simplify the FOC for ΔR_{kl} as

$$
\begin{aligned}
&\underbrace{\frac{l_{kl}}{2}\left(\eta\left(p_k^{land} + p_l^{land}\right) + \delta(1+\lambda)\bar{p}\right)}_{\text{financial cost}} - \underbrace{\phi_k^{H_k^r}\frac{\partial H_k^r}{\partial \Delta R_{kl}} + \phi_l^{H_l^r}\frac{\partial H_l^r}{\partial \Delta R_{kl}}}_{\text{loss of res land}} - \underbrace{\phi_k^{H_k^f}\frac{\partial H_k^f}{\partial \Delta R_{kl}} + \phi_l^{H_l^f}\frac{\partial H_l^f}{\partial \Delta R_{kl}}}_{\text{loss of bus land}} \\
&- \underbrace{\zeta_k^{l_k^{or}}\frac{\partial l_k^{or}}{\partial \Delta R_{kl}} + \zeta_l^{l_l^{or}}\frac{\partial l_l^{or}}{\partial \Delta R_{kl}}}_{\text{extra burden on affected owners}} - \underbrace{\Phi_k^{p_k^{land}}\frac{\partial p_k^{land}}{\partial \Delta R_{kl}} + \Phi_l^{p_l^{land}}\frac{\partial p_l^{land}}{\partial \Delta R_{kl}}}_{\text{higher share of people paid}} = \underbrace{(\psi_{kl} + \psi_{lk})\left(\xi\kappa\bar{t}_{kl}d_{kl}\right)\frac{1}{R_{kl}^{1+\xi}}}_{\text{speed benefit}}
\end{aligned}
$$

Each term of this first-order condition has a direct economic interpretation, and the partial derivatives appearing in it have closed forms. The right-hand side is the per-period benefit of widening link kl: the speed gain, valued by the shadow price of lower commuting costs on the link, $(\psi_{kl} + \psi_{lk})$. The left-hand side collects the four marginal costs.

1. The financial costs: the opportunity cost of the construction $\frac{l_{kl}}{2}\bar{p}(1+\lambda)$ and the fiscal burden of land payments $\frac{l_{kl}}{2}\eta(p_k^{\text{land}}+p_l^{\text{land}})$.

2. The loss of residential and business floorspace used by residents and firms: the shadow costs of residential and business floorspace lost $\phi_m^{H^r}, \phi_m^{H^f}$, scaled by the floor-to-area ratio in these locations. For each endpoint $m \in \{k,l\}$, these floorspace losses are $\frac{\partial H_m^r}{\partial \Delta R_{kl}} = -\frac{l_{kl}}{2}\gamma_m \mu_m^r$ and $\frac{\partial H_m^f}{\partial \Delta R_{kl}} = -\frac{l_{kl}}{2}(1-\gamma_m)\mu_m^f$.

3. The change in roadside owners' income: decline in their rental income weighted by the shadow cost of lower owners' income $\zeta_k^{l_i^{or}}$ and $\zeta_l^{l_i^{or}}$, with $\frac{\partial l_k^{or}}{\partial \Delta R_{kl}} = \frac{\frac{l_{kl}}{2}}{N_k^{or}}\left(p_k^{\text{land}} - \left(q_k^r\gamma_k\mu_k^r + q_k^f(1-\gamma_k)\mu_k^f\right)\right)$ and analogously for l;

4. The change in the compensation rate: under the claim-based compensation regime, $\mathbf{P}_m(\text{claim})$ is increasing in the potential compensation and therefore in the amount of land taken, as derived in Appendix F. This is weighted by the shadow costs of higher land payments in k and l, $\Phi_k^{p_k^{\text{land}}}$ and $\Phi_l^{p_l^{\text{land}}}$, respectively.

with $\frac{\partial p_i^{land}}{\partial \Delta R_{mi}}$ depending on the compensation rule:

1. No compensation: $\frac{\partial p_i^{land}}{\partial \Delta R_{mi}} = 0$

2. Baseline market value compensation: $\frac{\partial p_i^{land}}{\partial \Delta R_{mi}} = 0$

3. Land-tenure-based compensation:

$$\frac{\partial p_i^{land}}{\partial \Delta R_{mi}} = a(\bar{q}_i^0)^2 \frac{l_{mi}}{2N_i^{or}} \frac{\exp\left(\frac{a\bar{q}_i^0 \sum_{m\in C_i} \Delta R_{mi}\frac{l_{mi}}{2}}{N_i^{or}} - \bar{Z}_i\right)}{\left(1+\exp\left(\frac{a\bar{q}_i^0 \sum_{m\in C_i} \Delta R_{mi}\frac{l_{mi}}{2}}{N_i^{or}} - \bar{Z}_i\right)\right)^2}$$

where $\bar{q}_i^0 \equiv \gamma_m \mu_m^r q_m^{r,0} + (1-\gamma_m)\mu_m^f q_m^{f,0}$.

F.2 Solving for the Optimal Policy

I solve the planner's problem directly as a single large-scale nonlinear program. The choice variables are the per-link widenings $\{\Delta R_{kl}\}$, bounded below at zero (no road destruction) and above at a width cap $\overline{\Delta R}$. The spatial equilibrium is imposed as equality constraints: for any candidate set of widths, the commuting-cost matrix, the labor-mobility shares, wages, residential and business rents, and the welfare aggregates must jointly satisfy the equilibrium conditions of Section 5. The objective is the net present value of workers' and owners' welfare, capitalized at the discount rate δ, net of construction and land-compensation costs financed at the relevant shadow price of public funds. I implement the program in `JuMP` (Dunning et al., 2017) and solve it with the interior-point solver `Ipopt` (Wächter and Biegler, 2006).

Idiosyncratic route choice (Allen and Arkolakis, 2022) makes the equilibrium objects smooth in the road widths, so the program is differentiable but non-convex. Unlike the optimal-network problem of Fajgelbaum and Schaal (2020), which is globally convex when congestion dominates the returns to infrastructure, the analogous convexifying force here—land scarcity, which raises the marginal cost of widening as land is taken—is partly a transfer to owners and co-moves with the access benefit it creates, so convexity is not guaranteed; the claim rule, under which compensation rate is endogenous to the policy, adds a further non-convexity. Ipopt therefore locates only a local optimum whose basin can depend on the starting point. I address this with multistart and warm-started continuation across the grid of scenarios I solve—every combination of compensation rule, funding regime, cost of public funds η, and discount rate δ: I solve each cell from several seeds and keep the highest-welfare solution, sweep η downward warm-starting each solve from the adjacent one, and initialize the donor-financed-land solves from their restricted-funding counterparts. The claim rule, the most sensitive to initialization, is solved as a single warm-start chain across η and funding regimes.

F.3 Proofs

F.3.1 Assumptions

Assumption 1: *Assume market value compensation at rate* $p_k^{land} = \gamma_k \mu_k^r q_k^{r,0} + (1-\gamma_k)\mu_k^f q_k^{f,0}$*, independent of* ΔR_{kl}*.*

Assumption 2: *Assume that owners' utility is linear in income, s.t.* $W_i^{or} = l_i^{or}$ *and* $W_i^{of} = l_i^{of}$*.*[44]

F.3.2 Proposition 1

Proposition 1: *Under Assumption 1, the welfare-maximizing road widening policy* ΔR_{kl} *is decreasing in the shadow value of residential* $\phi_k^{H^r}, \phi_l^{H^r}$ *and business* $\phi_k^{H^f}, \phi_l^{H^f}$ *floorspaces at the endpoints* k *and* l*. A planner who ignores the opportunity cost of land use* ($\phi_H^r = \phi_H^f = 0$) *overinvests in locations with high opportunity cost of land use.*

Proof: The FOCs for ΔR_{kl} give us

$$\frac{l_{kl}}{2}\left(\eta(p_k^{\text{land}} + p_l^{\text{land}}) + \delta(1+\lambda)\bar{p}\right) - \phi_k^{H^r}\frac{\partial H_k^r}{\partial \Delta R_{kl}} - \phi_l^{H^r}\frac{\partial H_l^r}{\partial \Delta R_{kl}} - \phi_k^{H^f}\frac{\partial H_k^f}{\partial \Delta R_{kl}} - \phi_l^{H^f}\frac{\partial H_l^f}{\partial \Delta R_{kl}} - \zeta_k^{l_k^{or}}\frac{\partial l_k^{or}}{\partial \Delta R_{kl}} - \zeta_l^{l_l^{or}}\frac{\partial l_l^{or}}{\partial \Delta R_{kl}}$$
$$- \Phi_k^{p_k^{\text{land}}}\frac{\partial p_k^{\text{land}}}{\partial \Delta R_{kl}} - \Phi_l^{p_l^{\text{land}}}\frac{\partial p_l^{\text{land}}}{\partial \Delta R_{kl}} = \frac{\kappa\xi\bar{t}_{kl}d_{kl}}{(R_{kl}^0 + \Delta R_{kl})^{1+\xi}}(\psi_{kl} + \psi_{lk}) \quad \forall kl \quad (27)$$

where the LHS is increasing in $\phi_m^{H^v}, \quad m \in k, l \quad v \in r, f$ as the shadow value of lower commuting costs on kl is weakly positive and the RHS is decreasing in ΔR_{kl} as it follows from the concavity of owners' utility function that the shadow value of higher owner income $\zeta_m^{l_m^{or}}, \quad m \in k, l$ is positive.[45]

[44] Note that $\bar{C}^o \neq 1$ would imply that 1 dollar owned by the government would not have the same welfare impact as 1 dollar owned by owners.

[45] In the quantitative model, the shadow values $\phi_m^{H^v}, \quad m \in k, l \quad v \in r, f$ are strongly positively correlated with equilibrium rental rates q_m^r and $q_m^f, \quad \forall m$, as displayed on Appendix Figure A13, so the planner who ignores land use overinvests disproportionally in high-rent locations.

F.3.3 Proposition 2

Proposition 2: *Under Assumption 1, and when the shadow price of public funds equals one ($\eta = 1$), the optimal road widening policy $\{\Delta R_{kl}\}$ is independent of the compensation rule $\{p_k^{land}\}$. When the shadow price of public funds is greater than one, the optimal road widening policy $\{\Delta R_{kl}\}$ is decreasing in the land compensation rates at the endpoint locations k and l.*

Proof: when owner utility is linear in income, and after multiplying the first-order condition through by δ (so the welfare shadow values absorb the $\frac{1}{\delta}$), we can simplify the following:

- $\zeta_m^{l^{or}} = N_m^{or} \quad m \in k, l$
- $\frac{\partial l_m^{or}}{\partial \Delta R_{kl}} = \frac{l_{kl}}{2} \frac{1}{N_m^{or}} \left(p_m^{\text{land}} - \left(\gamma_m \mu_m^r q_m^r + (1 - \gamma_m) \mu_m^f q_m^f \right) \right) \quad m \in k, l$

Plugging these in equation (24) and rearranging gives us:

$$
\begin{aligned}
&\frac{l_{kl}}{2}(\eta - 1)\left(p_k^{\text{land}} + p_l^{\text{land}}\right) \\
+\frac{l_{kl}}{2}\delta(1+\lambda)\bar{p} - \phi_k^{H^r} \frac{\partial H_k^r}{\partial \Delta R_{kl}} - \phi_l^{H^r} \frac{\partial H_l^r}{\partial \Delta R_{kl}} &- \phi_k^{H^f} \frac{\partial H_k^f}{\partial \Delta R_{kl}} - \phi_l^{H^f} \frac{\partial H_l^f}{\partial \Delta R_{kl}} \\
+\frac{l_{kl}}{2}\left(\gamma_k \mu_k^r q_k^r + (1-\gamma_k)\mu_k^f q_k^f\right) &+ \frac{l_{kl}}{2}\left(\gamma_l \mu_l^r q_l^r + (1-\gamma_l)\mu_l^f q_l^f\right) \\
-\Phi_k^{p_k^{\text{land}}} \frac{\partial p_k^{\text{land}}}{\partial \Delta R_{kl}} - \Phi_l^{p_l^{\text{land}}} \frac{\partial p_l^{\text{land}}}{\partial \Delta R_{kl}} &= \frac{\kappa \xi \bar{t}_{kl} d_{kl}}{(R_{kl}^0 + \Delta R_{kl})^{1+\xi}}(\psi_{kl} + \psi_{lk}) \quad \forall kl,
\end{aligned}
$$

where only the terms on the first row depend on $p_m^{\text{land}} \quad m \in k, l$.

When $\eta = 1$, the first row equals to 0 and the optimal policy $\{\Delta R_{kl}\}$ does not depend on the land compensation $\{p_k^{\text{land}}\}$.

When $\eta > 1$, the LHS is increasing in p_k^{land} and p_l^{land}, while the RHS is decreasing in ΔR_{kl}, so ΔR_{kl} is decreasing in p_k^{land} and p_l^{land}.

F.4 Simulation: A Simple Economy on a Grid

To illustrate *Proposition 1* and *Proposition 2* and highlight the mechanisms of the model, I present simulations for a simple economy of a monocentric city on a 5×5 grid (25 locations), represented in Appendix Figure A10. Productivity is high in the central location and low everywhere else (panel a). All locations have similar residential amenities (panel b), total land supply, floor-to-area ratio, and road infrastructure at baseline. The distributions of equilibrium residential population (panel d), employment (panel e), residential rental rate (panel f) and business rental rate (panel g) are intuitive: most residents want to work in the central location to access the highest productivity and wages, and, as a result, their demand for land is decreasing in distance to the center, as they must incur a longer commute. This geography is a simplification of Kampala's, where most jobs are located at the center (Figure A10).

I then solve for the optimal road widening policy under different assumptions to illustrate how the level and location of road investments depends on the joint distribution of benefits and land-related costs. Summary statistics for the optimal policies are summarized in Appendix Table A15.

Illustration of Proposition 1: A planner who ignores the opportunity cost of land use over-invests in locations with high opportunity cost of land use: in Figure A12 I solve for the optimal policy assuming that roads do not use land ($\phi_H^r = \phi_H^f = 0$) (panel a) and accounting for the opportunity cost of land use (panel b). The former is the dominant approach in the literature and is equivalent to optimizing on spatially heterogeneous benefits alone. For simplicity, I assume no financial compensation to affected owners and I assume that owners' utility is linear in their rental income and therefore only partially captures the distributional channel of compensation for land. The roads connecting the central location host the well-paying jobs and thus have the largest marginal benefits (panel a). Yet, these central locations also have the largest opportunity cost of land use, for both firms who use this land for production and workers who would like to live as close as possible to well-paying jobs. Accounting for this opportunity cost of land use moves investments away from these central locations (panel b) towards more peripheral areas. Therefore, failure to account for land scarcity would lead the planner to overinvest in locations where land pressure is the highest, which, in monocentric cities like the simulated one, are also where the marginal benefits of better transportation infrastructure are the largest.

Illustration of Proposition 2: the optimal policy depends on the cost of public funds η: in panel (c) and (d) of Figure A12 , I solve for the optimal policy (with land use) with market value compensation to affected owners (*Assumption 1*). To isolate the mechanism of interest, I assume that owners' utility is linear in their rental income (*Assumption 2*). In panel (c), I set $\eta = 1$ and obtain the same investment allocation as in the absence of compensation (panel b). In panel (d), I set $\eta > 1$, so that effectively a fiscal loss is incurred on every transfer to affected owners. Investment declines everywhere, including in the more central locations where the benefits of road widening are the largest. The mechanism holds when assumptions (1) and (2) are relaxed, but the overall magnitude may be mitigated by the distributional channel of compensation and the correlation between compensation probability and the marginal benefits of road widening, as discussed below.

Figure A10: Simulated Economy - Baseline Equilibrium

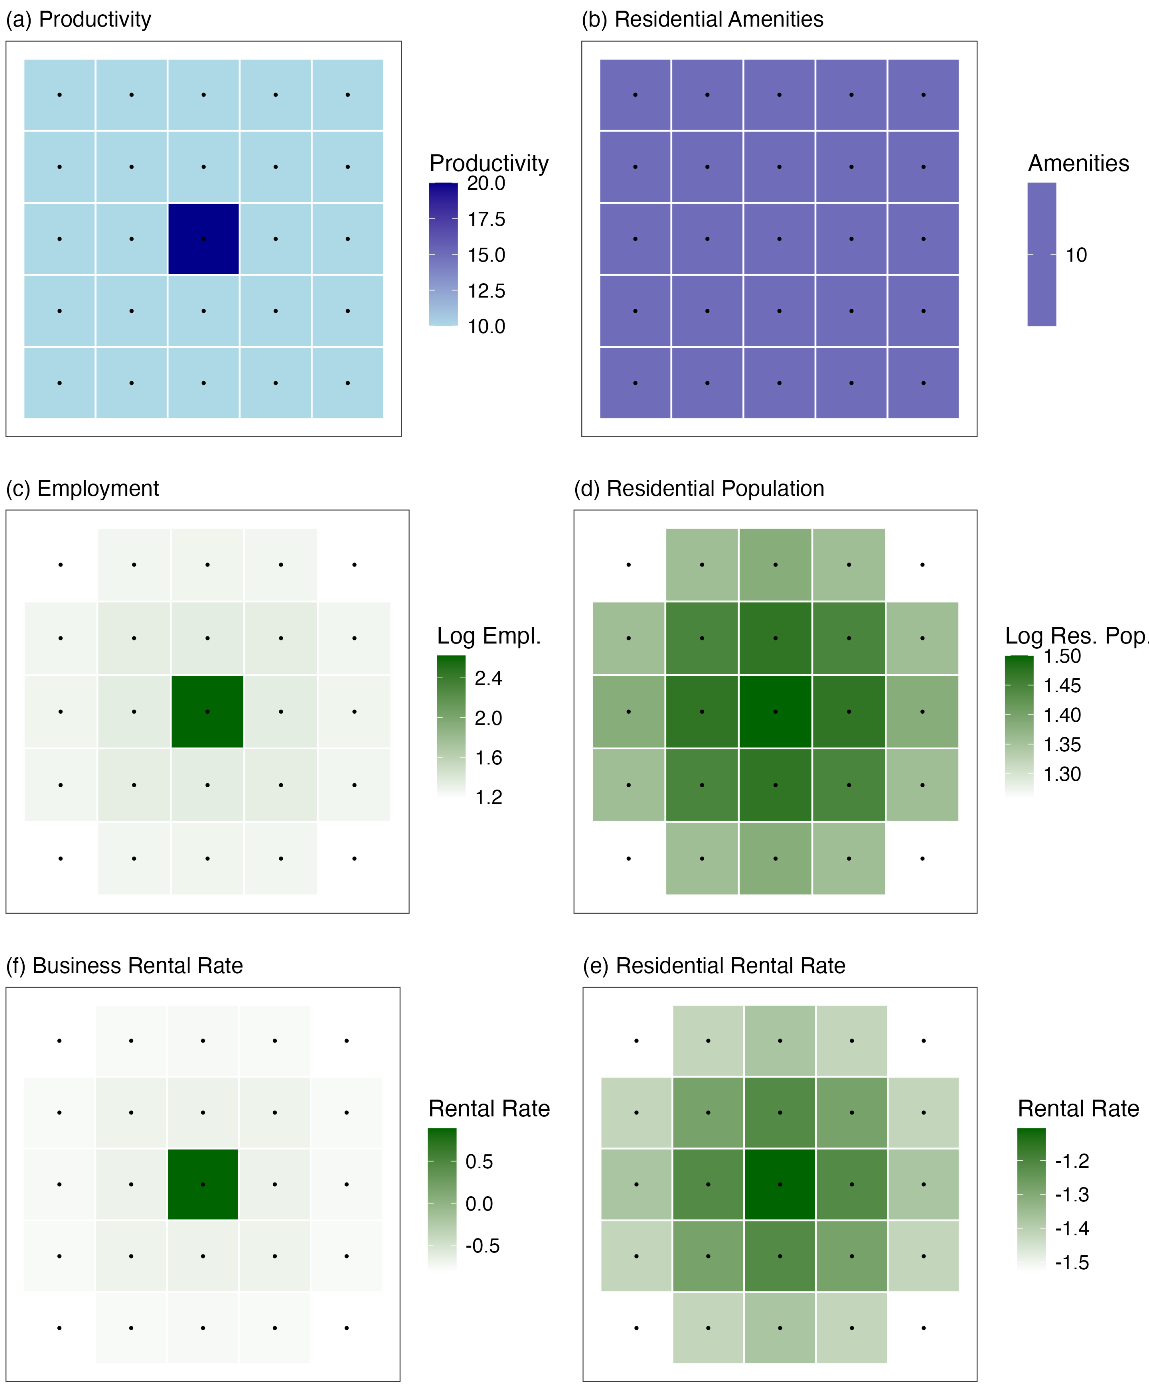

Figure A11: Simulated Economy - Compensation Probabilities for Figure A12, Panels (e) and (f)

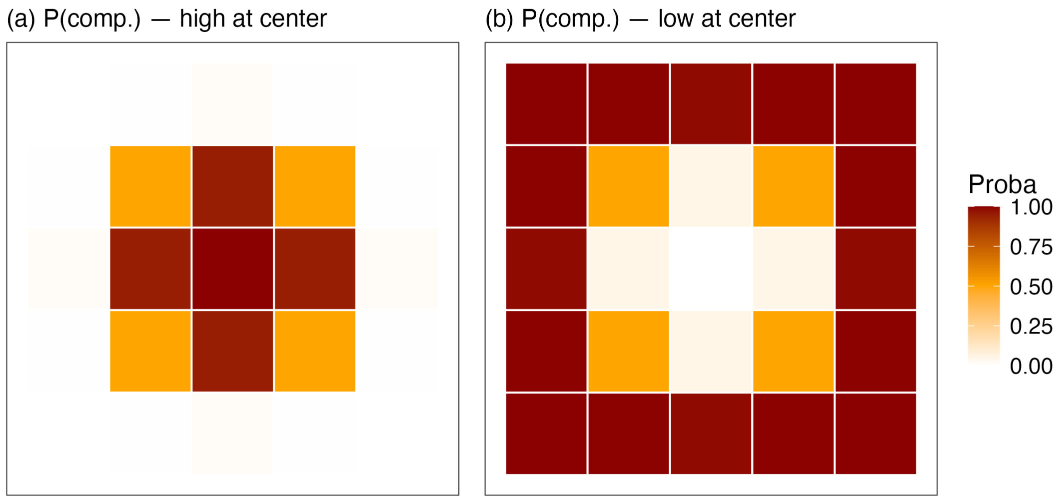

Remark: A generous compensation rule shapes the optimal policy through two channels going in potentially opposite directions: With $\eta > 1$, a compensation rule with $0 < \mathbf{P}_i < 1$ for some i affects the optimal policy through two effects, compared to a scenario where $\mathbf{P}_i = 1 \quad \forall i$ (*Assumption 1*): (i) a price effect and (ii) an income effect. The price effect leads to reallocation of investment towards locations with relatively lower compensation probabilities. The income effect lowers the overall fiscal burden, potentially increasing investment in high land market value locations and aggregate investment. Under *Assumption 2*, the income effect always has a positive impact on welfare, but the effect of the price effect is ambiguous and depends on the correlation between the compensation probability and the marginal benefits of road widening. I illustrate this channel on Figure A12, panel (e) and (f). On panel (e) $\mathbf{P}_i$ is positively correlated with baseline rental rates, while they are negatively correlated on panel (f), as displayed on Appendix Figure A11. When $0 < \mathbf{P}_i(\text{compensated}) < 1$ is positively (negatively) correlated with baseline rental rates, investments are shifted away from (toward) central locations, associated with a relatively higher (lower) – although nominally lower than in the market value compensation case – compensation than peripheral locations. The net welfare effects depend on whether the income or price effect dominates.

The compensation probabilities for the optimal policies with spatially heterogeneous probabilities of getting compensated at market value (Figure A12, panels (e) and (f)) are displayed on Appendix Figure A11.

Table A15: Simulated Economy - Optimal Policy: Summary Statistics

	(a) No land use No compensation	(b) Land use No compensation	(c) Land use Compensation MV $\eta = 1$	(d) Land use Compensation MV $\eta > 1$	(e) Land use Compensation high center $\eta > 1$	(f) Land use Compensation low center $\eta > 1$	NA NA NA
Total road widening ($\sum \Delta R_{kl}$)	10.00	10.00	10.00	7.00	10.00	10.00	6.67
% Δ commuting cost (fixed pop.)	-1.10%	-0.98%	-0.98%	-0.72%	-0.97%	-0.98%	-0.69%
% Δ commuting cost (new pop.)	-0.45%	-0.33%	-0.33%	-0.23%	-0.34%	-0.32%	-0.22%
Total compensation for land	0	0	9.93	6.99	6.108	2.71	6.66
% Δ net welfare	+0.862%	+0.236%	+0.236%	+0.022%	+0.085%	+0.170%	+0.023%

Notes: Summary statistics for the optimal policy as displayed in Figure A12, with column headers corresponding to the figure's panel. *% Δ commuting costs (fixed pop.)* displays the percentage change in average commuting costs from the policy, weighted by baseline population while *% Δ commuting costs (new pop.)* weights by the new population.

Figure A12: Simulated Economy - Optimal Policy

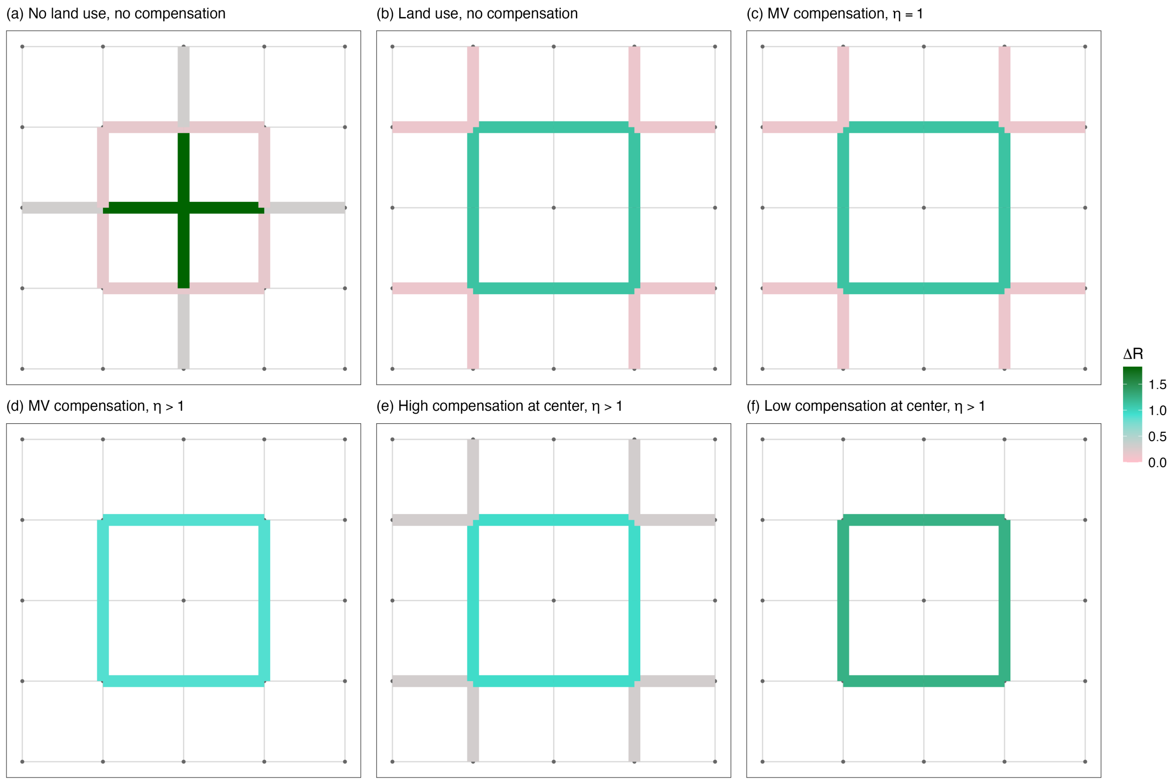

Figure A13: Simulated Economy - Correlations of Rental Rates and Land Shadow Values

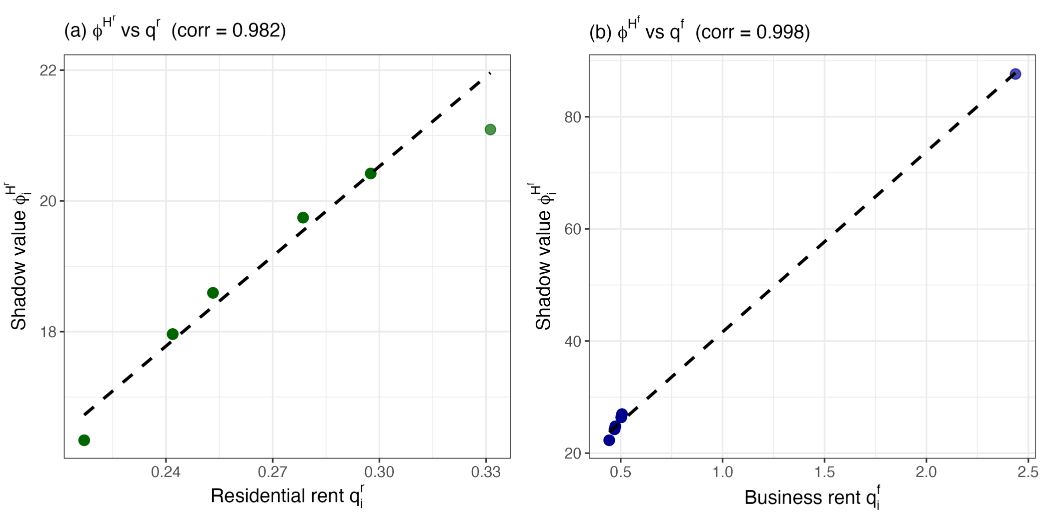

F.5 Model Quantification Appendix

F.5.1 Gravity Estimation of $\theta\kappa$

I estimate the composite semi-elasticity $\theta\kappa$ from the gravity equation of commuting flows on commuting times. The gravity Equation (11) implies that

$$L_{ij} \propto \exp\big(-\theta\kappa\, t_{ij}\big), \tag{28}$$

so that $\theta\kappa$ can be recovered from a Poisson pseudo-maximum-likelihood (PPML) regression of flows on travel times, controlling for origin and destination fixed effects (Silva and Tenreyro, 2006).

I estimate $\theta\kappa$ using the ride-hailing data aggregated to parish-group pair (i, j) in year-month t in two steps. Kampala's 96 parishes are consolidated into 18 parish groups to limit the sparsity of the commuting-flow matrix. First, I predict travel time (in minutes) by parish group pair (i, j) in year-month t by a trip-level regression of log duration on log straight-line distance, hour-of-day indicators, and year-month fixed effects, weighted by the observed distribution of trip request hours during weekday morning peak hours (6–11 am) in the data. Second, I regress commuting flows — defined as the number of distinct riders traveling between each parish group pair during weekday morning peak hours in a given month — on predicted travel times, including origin × year-month $\mu_{i,t}$ and destination × year-month $\nu_{j,t}$ fixed effects:

$$L_{ij,t} = \exp\big(-\theta\kappa\, \hat{t}_{ij,t} + \mu_{i,t} + \nu_{j,t}\big) + \varepsilon_{ij,t}, \tag{29}$$

with standard errors clustered at the parish group-pair level. The identification assumption for $\theta\kappa$ is that, conditional on origin × year-month and destination × year-month fixed effects, bilateral travel times $t_{ij,ym}$ are uncorrelated with the error term $\epsilon_{ij,ym}$. These fixed effects absorb all time-varying origin- and destination-specific determinants of flows, so identification relies on variation in travel times across pairs within the same origin-month and destination-month. This condition would be violated if unobserved bilateral factors affecting commuting flows — such as historical linkages between specific neighborhoods — were also correlated with bilateral travel times. My preferred specification uses predicted travel times from the first-stage speed technology regression, which mitigates reverse causality from transport investment responding to demand on high-flow corridors.

My preferred estimate is $\hat{\theta}\hat{\kappa} = 0.0319$ (Table A18, column 1). The estimate is robust across alternative specifications, including using measured (rather than predicted) travel times (column 2), using non-aggregated parishes (column 3), and estimating the gravity equation on the 2016/2017 Travel Habit Survey work-related commuting flows instead of ride-hailing data (columns 4-5).[46] While the Travel Habit Survey covers a representative sample of commuters, its small sample size leads to less robust estimates and requires high levels of aggregation. My preferred estimate using predicted times from the ride-hailing data (column 1) is stable across alternative specifications and falls between the THS estimates using predicted and measured times (columns 4-5). Self-reported travel times in the survey are subject to recall error, which, under classical measurement error, attenuates the coefficient toward zero and provides a lower bound on $\theta\kappa$. Conversely, predicted times based on distance alone attenuate measurement error but may

[46] Since only the 2016/2017 wave includes travel times, the THS gravity uses this wave only.

capture bilateral variation in commuting times beyond distance, providing an upper bound.

While very stable across specifications, it is worth acknowledging two potential issues with the ride-hailing data. First, there is a potential selection issue as users of this motorcycle taxi company may be different than the average commuter in Kampala. Second, using data on actual trips (rather than predicted) prevents me from including an observation (pair $\times$ period) for which no trip was recorded in a given time period. My estimate $\hat{\theta\kappa}$ is within the range but at the lower end of estimates in other contexts, including Bogota (Tsivanidis (2026), $\hat{\theta\kappa} \in [0.028, 0.046]$). As $\theta\kappa$ governs workers' response to a change in commuting times, a larger $\theta\kappa$, as estimated for Santiago (Bordeu (2024), $\hat{\theta\kappa} = 0.0656$) and Berlin (Ahlfeldt et al. (2015), $\hat{\theta\kappa} = 0.0683$), would imply larger benefits from road improvements.

F.5.2 Separate Identification of κ

To isolate θ and κ separately from $\hat{\theta\kappa} = 0.0319$, I exploit mode choice variation in the Travel Habit Survey to pin down κ. As the different modes of transportation (walking, shared minibus taxis, private cars, motorcycle taxis) have distinct speed technologies, commuters choosing among modes for the same origin-destination pair face different travel times. κ is identified from this within-OD, across-mode variation in travel times in a conditional logit framework.

I proceed in two stages. First, I estimate mode-specific speed technologies by regressing log travel time on log distance interacted with mode indicators and departure-hour fixed effects using all trips with observed travel times in the four key modes ($N = 1,880$ trips; Table A19). Data on travel times is only available for the 2016/2017 wave. Second, I use the first-stage estimates to predict mode-specific travel times for each commute trip, creating a choice set of four alternatives per trip, and estimate a conditional logit model (McFadden, 1974):

$$P(\text{mode } m \mid i) \propto \exp\bigl(-\kappa\, \hat{t}_{m,i} + \delta_m\bigr), \tag{30}$$

where $\hat{t}_{m,i}$ is the predicted travel time for trip i under mode m and δ_m are alternative-specific constants. The preferred sample comprises 142 home-work and work-home trips for the 2016/2017 wave, across 76 parish-group OD pairs, of which 27 are served by two or more modes. The results are presented in Appendix Tables A19 and A20. I estimate $\hat{\kappa} = 0.0089$ (robust SE $= 0.0078$, $p = 0.251$), which, combined with $\theta\kappa = 0.0319$ from the gravity specification implies $\theta = 4.36$.

Several caveats apply. First, the predicted travel times are generated regressors, so the reported standard errors may be understated. Second, I assume that all four modes are available to every commuter at every OD pair, which likely overstates the effective choice set for some trips. Third, the limited number of work commuting trips prevents me from precisely estimating the coefficient.

As a robustness check, I re-run the second stage pooling commute trips from both the 2011 and 2016/2017 THS waves, which increases the sample size but assumes that the transportation technology was stable between the two survey rounds — a strong assumption given the fast pace of Kampala's growth. I obtain $\hat{\kappa} = 0.0184$ (SE $= 0.0054$), confirming that the estimate is robust to pooling both survey waves (Table A20, column 2). In the Appendix, I re-solve the model for this alternative value of κ and the implied θ.

Despite these limitations, the point estimate is reassuringly close to Tsivanidis (2026)'s conditional-logit estimates from Bogota ($\kappa_H = 0.013$, $\theta_H = 3.10$ for high-skill workers; $\kappa_L = 0.011$, $\theta_L = 2.47$ for low-skill workers) and implies a Frechet shape parameter θ that falls within the range estimated in other developing-country urban settings.

F.5.3 Infrastructure-Free Travel Time

The baseline travel time t^0_{kl} is available for 320 out of 344 links. To recover $\bar{t}_{kl}$ for every edge, I regress log $\bar{t}_{kl}$ on the log of kl's length, log distance to the CBD and log total road length within the edge's catchment area, and use the regression-predicted values for all edges, ensuring that $\bar{t}_{kl}$ varies smoothly across the network rather than mixing observed and imputed values. The correlation between predicted $\bar{t}_{kl}$ and $\bar{t}_{kl}$ recovered from t^0_{kl} and road width is 0.47 for the 93% of edges matched to Google Maps data.

F.5.4 Floorspace Productivity

I recover the productivity term μ_i residually from floorspace market clearing. For residential floorspace, equating floorspace supply to demand at observed rents gives floorspace

$$H_i^R = (1-\beta)\,\mathbb{E}[w]_i\, L_i^r / q_i^R, \tag{31}$$

and hence

$$\mu_i^R = \frac{H_i^R}{\gamma_i\, H_i^{pr0}} = \frac{(1-\beta)\,\mathbb{E}[w]_i\, L_i^r}{q_i^R\,\gamma_i\, H_i^{pr0}}. \tag{32}$$

Business productivities μ_j^F are recovered analogously from firms' floorspace market clearing. By construction, the model reproduces observed broker rents at baseline.

F.5.5 Number of Landowners

The number of landowners is recovered from baseline rents under a free-subdivision assumption: every owner holds a parcel of the average area $\bar{A} = 497$ m^2 (broker survey) earning the average land rental rate $\bar{r}$, for a common baseline income $\bar{y} = \bar{A}\,\bar{r}$. The number of owners in parish i is parish rental income divided by this per-owner income,

$$\bar{N}_i \;=\; \frac{q_i^{r,0} H_i^r + q_i^{f,0} H_i^f}{\bar{A}\,\bar{r}}, \tag{33}$$

so owner counts scale with parish land value. Each owner holds a square parcel of side $\ell_i = \sqrt{A_i/\bar{N}_i}$, where A_i is the parish's private ground area, and I split owners by road frontage — the roadside count is the one-sided road length divided by the parcel side:

$$N_i^{or} \;=\; \frac{\sum_{m\in\mathcal{C}_i} l_{mi}/2}{\ell_i}, \qquad N_i^{of} \;=\; \bar{N}_i - N_i^{or}, \tag{34}$$

capped so both groups stay non-empty.[47] This implies a common baseline owner income of 359 USD/-month, about 2.44× the average worker wage, with roadside owners accounting for 6.9% of all owners.[48]

[47]Road expansion is assumed to occur on one side, and N_i^{or} is capped at $\bar{N}_i$.

[48]This income gap — the source of the redistributive cost of compensation in the case of non-linear owner utility — is not an artifact of how the number of owners is recovered. An alternative headcount based on census owner-occupancy

Table A16: Estimation of ξ

	Rush (AM+PM)		Non-rush		All hours	
	(1)	(2)	(3)	(4)	(5)	(6)
Log Infrastructure (width m)	0.505	0.367	0.306	0.142	0.344	0.126
	(0.122)	(0.158)	(0.107)	(0.132)	(0.103)	(0.136)
Estimator	TWFE	SA	TWFE	SA	TWFE	SA
Road + Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	752	752	684	684	1,044	1,044
Roads (total)	50	50	50	50	51	51
Roads (switchers)	23	23	22	22	23	23

Notes: Standard errors are displayed in parentheses. Standard errors are two-way clustered by road segment and day (columns 1–4) or by road segment (columns 5–6). The dependent variable is log traffic speed (km/h). Road infrastructure is measured as total road width in meters: number of lanes × 3.4m per lane, plus one additional lane upon completion of the upgrade. Columns (1)–(3) are trip-level OLS with day and hour-of-day fixed effects; column (1) includes no controls, column (2) adds road distance and geographic controls (longitude, latitude, and their squares), and column (3) further adds road-class dummies and baseline paving status. Column (4) is a trip-level TWFE specification with road-segment and day fixed effects. Columns (5)–(6) use a road × month panel: log speed is residualized on day and hour-of-day fixed effects at the trip level, then collapsed to road–month means. Column (5) estimates TWFE with road-segment and month fixed effects. Column (6) reports the implied elasticity from the Sun and Abraham (2021) interaction-weighted estimator: $\hat{\xi}_{\text{SA}} = \widehat{ATT}_{\text{log(speed)}} \,/\, \overline{\Delta \log(\text{infra})}$, where $\overline{\Delta \log(\text{infra})}$ is the mean change in log infrastructure at completion among treated roads. "Switchers" are roads observed both before and after upgrade completion.

rates, rather than the land-value/free-subdivision construction above, yields an owner-to-worker income ratio of comparable magnitude (about 2.5, versus 2.44 here).

Table A17: Estimation of ξ - Different Times of Day

	Non-rush hours			Evening rush		Morning rush		Rush (M+E)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Infrastructure	0.687	0.306	0.142	0.621	0.581	0.416	0.202	0.367
	(0.119)	(0.107)	(0.132)	(0.075)	(0.127)	(0.160)	(0.297)	(0.158)
Level	Trip	Panel	Panel	Trip	Panel	Trip	Panel	Panel
Model	OLS	TWFE	SA	OLS	TWFE	OLS	TWFE	SA
Day + Hour FE	Yes			Yes		Yes		
Road + Month FE		Yes	Yes		Yes		Yes	Yes
Full controls	Yes			Yes		Yes		
Observations	1,075	684	684	831	618	186	161	752
Roads (total)	50	50	50	50	50	45	45	50
Roads (switchers)		22	22		23		5	23

Notes: Standard errors are displayed in parentheses. Standard errors are two-way clustered by road segment and day (trip-level columns) or by road segment (panel columns). Same specifications as Table A16, varying the time-of-day subsample. Non-rush hours exclude both morning (6:00–10:59am) and evening (3:00–8:59pm) rush. Trip-level OLS columns include day and hour-of-day fixed effects with the full set of controls (geographic controls, road class, baseline paving status). Panel columns use the road × month panel with road-segment and month fixed effects. Column (8) pools morning and evening rush for the Sun-Abraham estimator. SA columns report the implied $\xi = \widehat{ATT}_{\log(\text{speed})} \,/\, \overline{\Delta \log(\text{infra})}$.

Table A18: Gravity estimation of $\theta \times \kappa$

Dependent Variable:			Flows		
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Travel time (minutes)	-0.0319	-0.0352	-0.0317	-0.0876	-0.0117
	(0.0039)	(0.0036)	(0.0015)	(0.0134)	(0.0053)
Fixed-effects					
Origin $\times$ Month	Yes	Yes	Yes		
Dest. $\times$ Month	Yes	Yes	Yes		
Origin				Yes	Yes
Dest.				Yes	Yes
Fit statistics					
Standard-Errors		OD Pair		Origin & Dest.	
Observations	4,536	4,513	128,448	324	95
Pseudo R^2	0.89758	0.90050	0.61861	0.61198	0.53604
Dataset	Ride-hailing	Ride-hailing	Ride-hailing	Travel survey	Travel survey
Travel time	Predicted	Measured	Predicted	Predicted	Measured
Spatial units	Parish group	Parish group	Parish	Parish group	Parish group

Notes: Poisson pseudo-maximum likelihood (PPML) estimates of equation (29). The dependent variable is the flow of trips between origin–destination pairs. Columns (1) and (3) use predicted travel times from the first-stage speed technology regression; column (2) uses measured travel times from the ride-hailing data. Column (3) uses parish-level spatial units; columns (1)–(2) and (4)–(5) use aggregated parish groups matching the travel survey zone crosswalk. Columns (1)–(3) include origin $\times$ year-month and destination $\times$ year-month fixed effects; columns (4) and (5) include origin and destination fixed effects. Standard errors are clustered at the origin–destination pair level in columns with bilateral fixed effects, and two-way clustered at the origin and destination level otherwise.

Table A19: First stage: mode-specific speed technology

	Dependent variable:
	log(time)
log(distance)	0.894
	(0.149)
Taxi	4.311
	(0.875)
Car	3.984
	(0.882)
Boda-boda	2.632
	(0.911)
log(distance) × Taxi	−0.631
	(0.154)
log(distance) × Car	−0.644
	(0.156)
log(distance) × Boda-boda	−0.433
	(0.161)
Constant (Walking)	−2.660
	(0.882)
Departure hour FE	Yes
F-test mode effects	30.98
Observations	1,880
R^2	0.253
Adjusted R^2	0.241
F Statistic	22.340 (df = 28; 1851)

Notes: OLS regression of log travel time on log distance interacted with transport mode, with departure hour fixed effects. The base category is walking. The sample includes all inter-parish trips from both waves of the GKMA Household Travel Survey. Predicted travel times from this regression serve as inputs to the second-stage mode choice model and to the gravity estimation. Heteroskedasticity-robust standard errors in parentheses.

Table A20: Second stage: conditional logit on mode choice

	Dep. variable: mode chosen	
	2016/2017 wave	Both waves
Travel time (min.)	-0.0089	-0.0184
	(0.0078)	(0.0054)
Taxi (ASC)	2.7994	3.5036
	(0.5724)	(0.2209)
Car (ASC)	1.9198	1.8465
	(0.6998)	(0.2917)
Boda-boda (ASC)	2.0175	0.8700
	(0.6073)	(0.2698)
Trip strata	Yes	Yes
Base mode	Walking	Walking
Observations (trips × modes)	568	6,860
N trips	142	1,715
N OD pairs	76	171

Notes: Conditional logit (McFadden) estimates. The dependent variable is an indicator for the chosen transport mode among four alternatives: walking (base), minibus taxis, car, and motorcycles. Each trip defines a choice occasion. Travel times are predicted from the first-stage speed technology regression. The sample includes home–work and work–home inter-parish trips from the 2016-2017 wave (column 1) and from both waves (column 2) of the GKMA Household Travel Survey, weighted by household survey weights. Robust standard errors in parentheses.

Figure A14: Kampala at Baseline

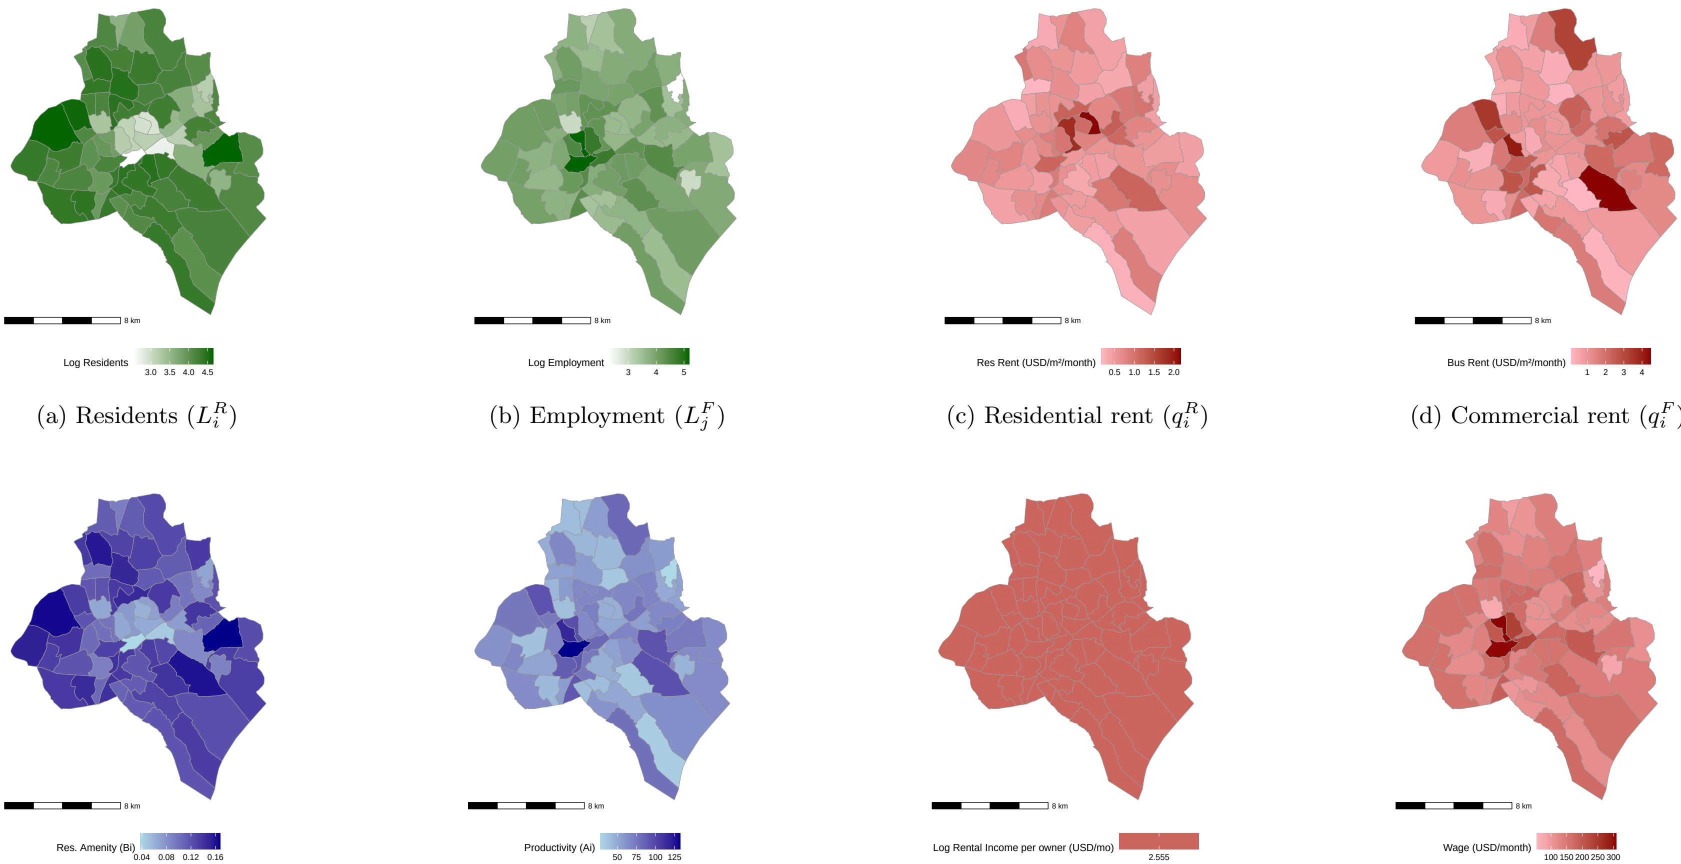

(a) Residents (L_i^R) (b) Employment (L_j^F) (c) Residential rent (q_i^R) (d) Commercial rent (q_i^F)

(e) Residential amenity (B_i) (f) Productivity (A_j) (g) Per Owner Rental Income (l_{io}) (h) Baseline wage (w_j)

Notes: Baseline distribution of economic activity across Kampala's 90 locations. Panels (a) and (b) show observed residential population and employment. Panels (c) and (d) show observed rents. Panels (e) and (f) show recovered residential amenities and productivity. Panel (g) shows owner rental income. Panel (f) shows local equilibrium wages.

Figure A15: Predicted Land-Acquisition Cost per m^2 by Parish, Across Compensation Definitions

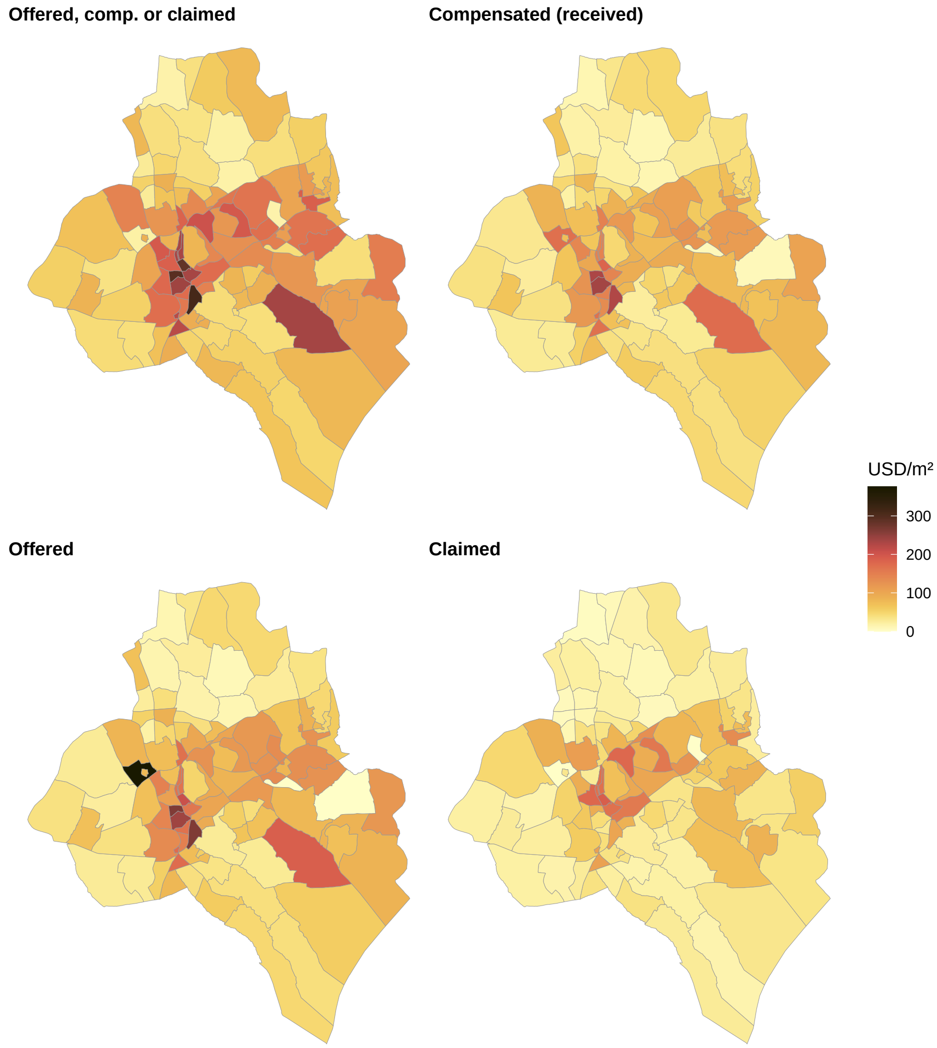

Notes: Each panel maps the expected land-acquisition cost per m^2 of land taken, by parish: the blended pre-policy market rate (USD/m^2; $0.5\times$(residential + business), 2018) times the predicted probability of compensation under one definition — offered, compensated, or claimed (union); compensated (received); offered; and claimed. Probabilities are parish-level predictions from the main specification (tenure + market value + wealth + parish controls; Table 1), full sample, evaluated at the survey-median taking. Common color scale across panels; the market rate alone (full-compliance / eminent-domain benchmark, $\mathbb{P} = 1$) is the upper envelope.

G Realized and Optimal Road Improvements Appendix

Figure A16: Realized Externally Funded Road Upgrades in Kampala: Additional Width Per Link

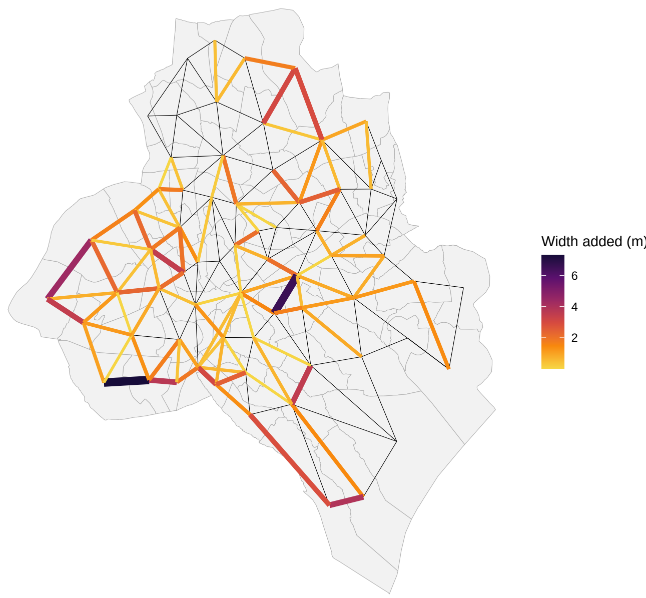

Notes: Each link connects two neighboring parishes. Links with externally funded road upgrades are highlighted, with color and thickness proportional to the effective width added (in meters). Grey links have no upgrade. Road projects are mapped to the closest link in the model's aggregated network.

Figure A17: Change in Total Rental Income by Parish Under the Realized Road Upgrades

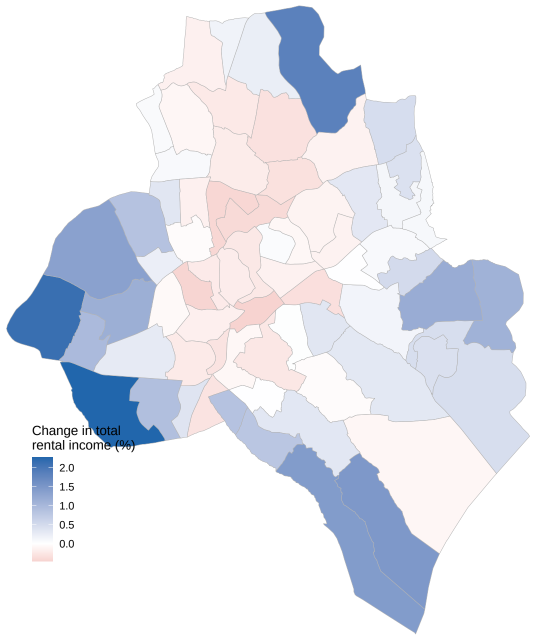

Notes: Each parish is shaded by the percentage change in its total private rental income — the sum of residential and business rents, i.e. rate times floorspace ($q_i H_i$) — between the pre-policy baseline and the new spatial equilibrium induced by the realized externally funded upgrades. Blue parishes gain and red parishes lose, with white near zero. The heterogeneity reflects general-equilibrium reallocation: rents rise where improved accessibility attracts residents and firms, and fall in locations workers move away from. Because it measures the change in market rents, the map is independent of how owners are compensated.

Figure A18: Heterogeneous Welfare Impacts on Affected Owners Depend on the Compensation Rule

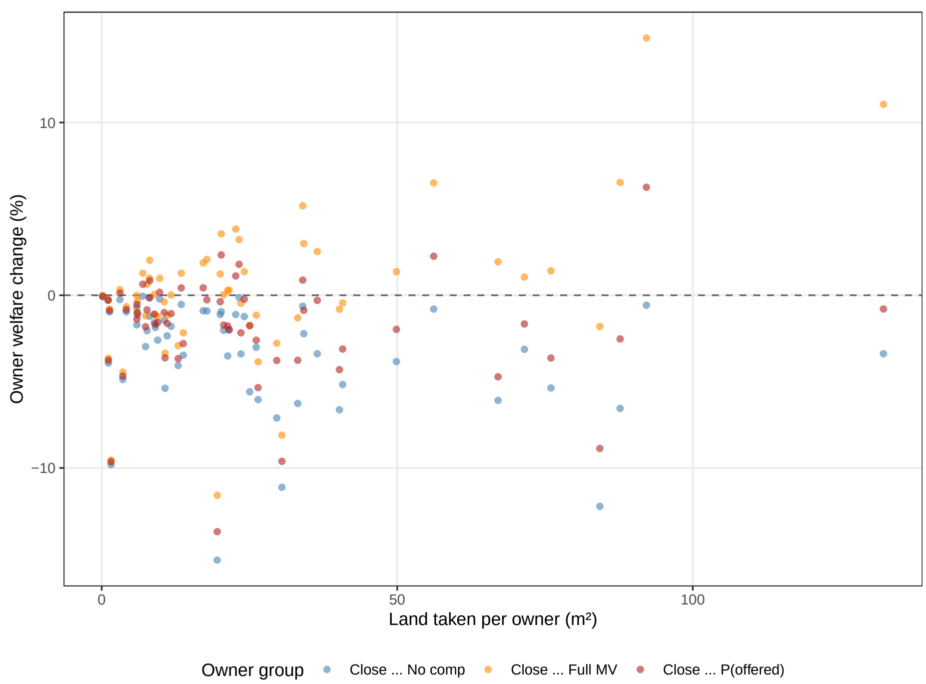

Notes: Each point is a parish. x-axis: land taken per close owner (m^2); y-axis: per-owner welfare change (%). Colors indicate compensation regime: no compensation (blue), full market-value compensation (orange), and claim-based compensation (red).

Table A21: Realized upgrades vs. model-predicted benefits (welfare) and costs

Land Cost Measure	Market Value		Upper Bound		Lower Bound		Claim		Tenure	
	(1) Level	(2) Per m	(3) Level	(4) Per m	(5) Level	(6) Per m	(7) Level	(8) Per m	(9) Level	(10) Per m
Benefit (welfare)	0.080	0.058	0.085	0.066	0.080	0.059	0.096	0.077	0.125	0.104
	(0.042)	(0.054)	(0.042)	(0.053)	(0.042)	(0.053)	(0.043)	(0.055)	(0.041)	(0.056)
Land cost	-0.070	-0.066	-0.062	-0.072	-0.033	-0.036	-0.080	-0.099		
	(0.048)	(0.035)	(0.056)	(0.053)	(0.060)	(0.056)	(0.047)	(0.048)		
Share freehold									-0.086	-0.079
									(0.068)	(0.069)
Share public									0.080	0.076
									(0.060)	(0.059)
Link length	0.055	0.009	0.033	0.007	0.023	0.016	0.020	0.001	-0.030	0.001
	(0.042)	(0.046)	(0.044)	(0.045)	(0.048)	(0.046)	(0.040)	(0.045)	(0.041)	(0.043)
R^2	0.036	0.028	0.038	0.030	0.028	0.016	0.048	0.045	0.082	0.062
N	172	172	172	172	172	172	172	172	172	172

Notes: Dependent variable: indicator for whether any road mapped to a model link (between neighboring parishes) was upgraded (2018–2026), regressed on the model's standardized marginal benefit as defined by the change in residents' welfare, a standardized land-acquisition cost, and link length; dyadic-robust standard errors. Columns are grouped by compensation rule (market value, upper-/lower-bound compensated, claim-based, and tenure shares). Within each rule, "Level" measures benefit and cost as totals per link; "Per m" divides by link length.

Figure A19: Marginal benefit and land cost of widening are positively correlated across links

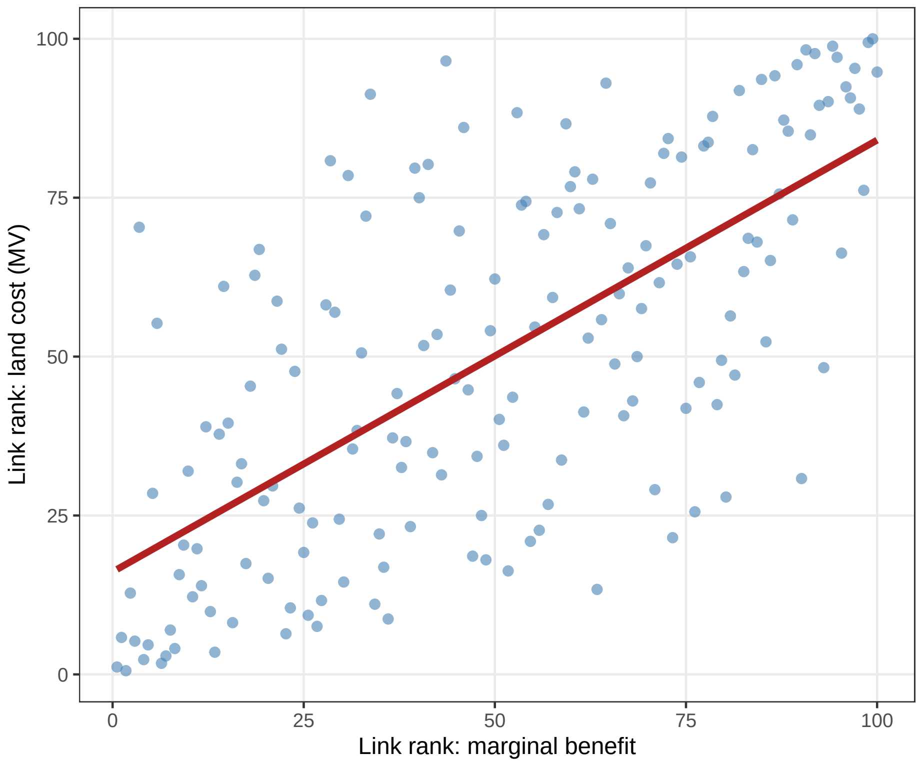

Notes: Each point is one undirected link in the commuting network. The horizontal axis is the within-network percentile of the model-implied marginal benefit of widening the link by one lane; the vertical axis is the percentile of its land cost under market-value compensation, measured from the broker survey. The positive relationship (rank correlation 0.68) is the empirical counterpart of the covariance $\mathrm{Cov}(p^{\mathrm{land}}, MB)$ that Proposition 3 identifies as the driver of the allocative distortion: land is dearest where widening helps most, so the fiscal wedge distorts investment most where returns are highest.

Figure A20: Optimal allocation as the cost of public funds η rises (market-value compensation, $\delta = 0.06$)

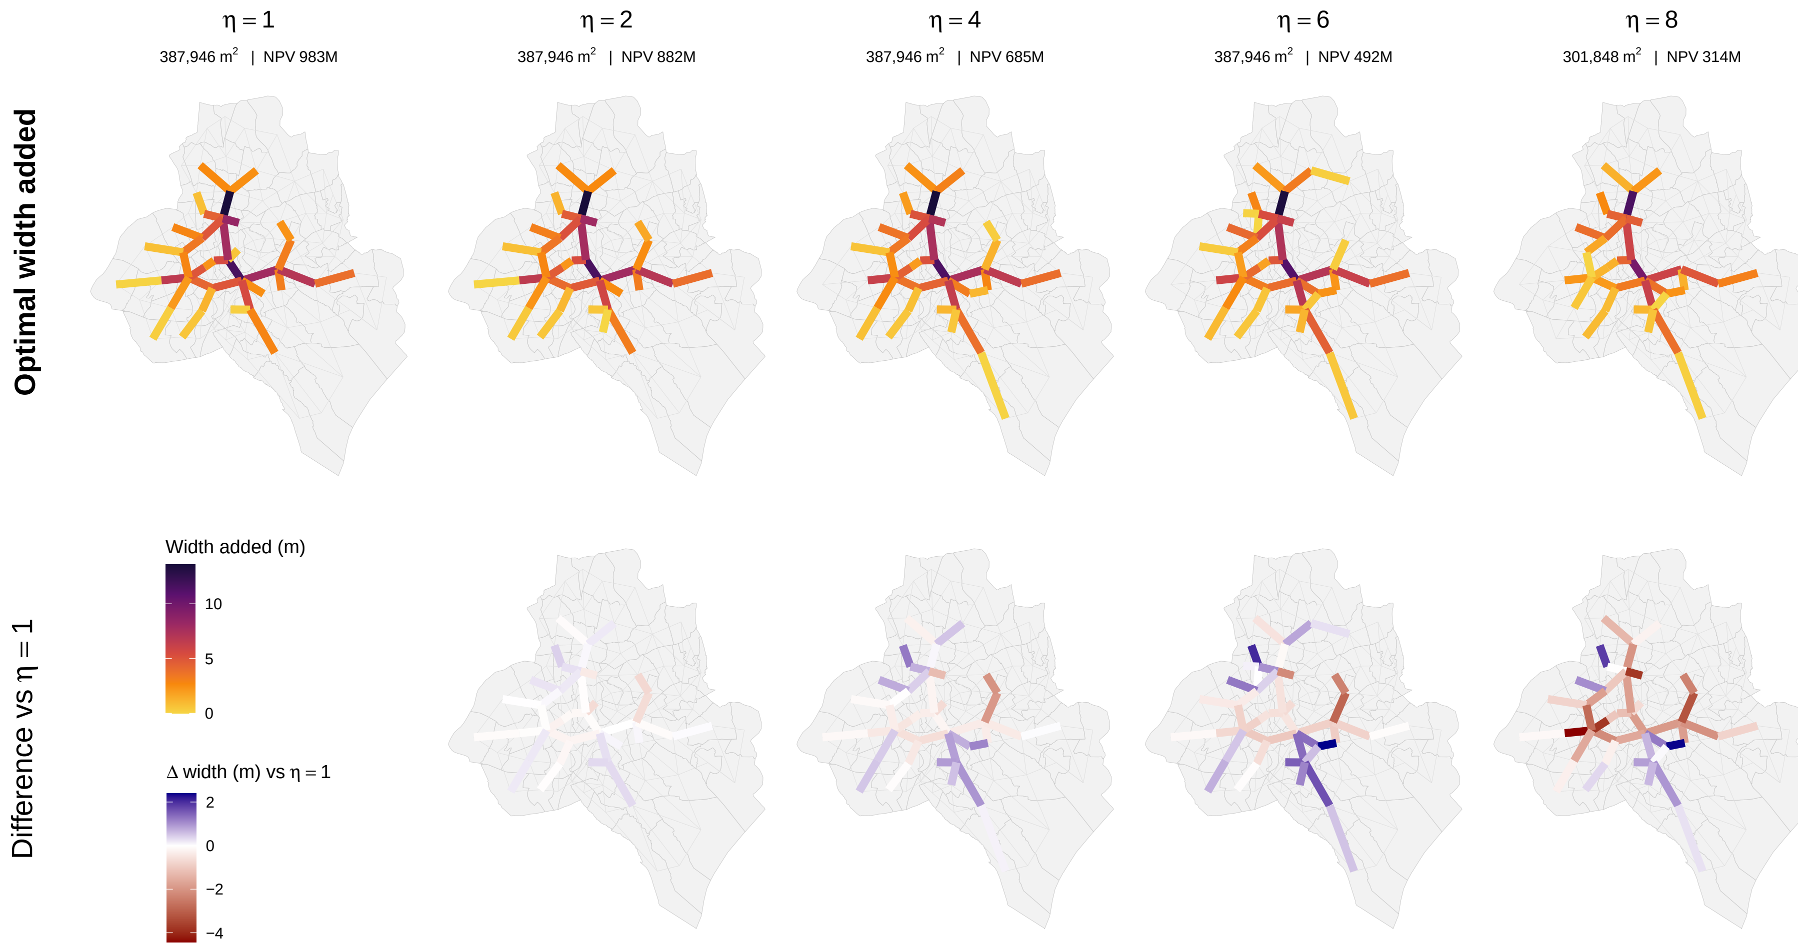

Notes: Top row: optimal width added to each link (metres) at $\eta \in \{1, 2, 4, 6, 8\}$, holding the discount rate at $\delta = 0.06$ and financing land payments under the status-quo rule. *Bottom row:* the change in widening relative to the no-wedge benchmark $\eta = 1$ (red: less widening; blue: more). Panel subtitles report the total widened area (m^2) and the net present value of the program (US$ millions). At $\eta = 1$ the planner implements the first-best map, concentrating widening on the high-benefit central ring; as η rises the welfare wedge $(\eta - 1)\, p_k^{\text{land}}$ pulls investment out of the expensive centre and what remains tilts toward cheaper peripheral links.

Figure A21: Optimal allocation across the discount rate δ.

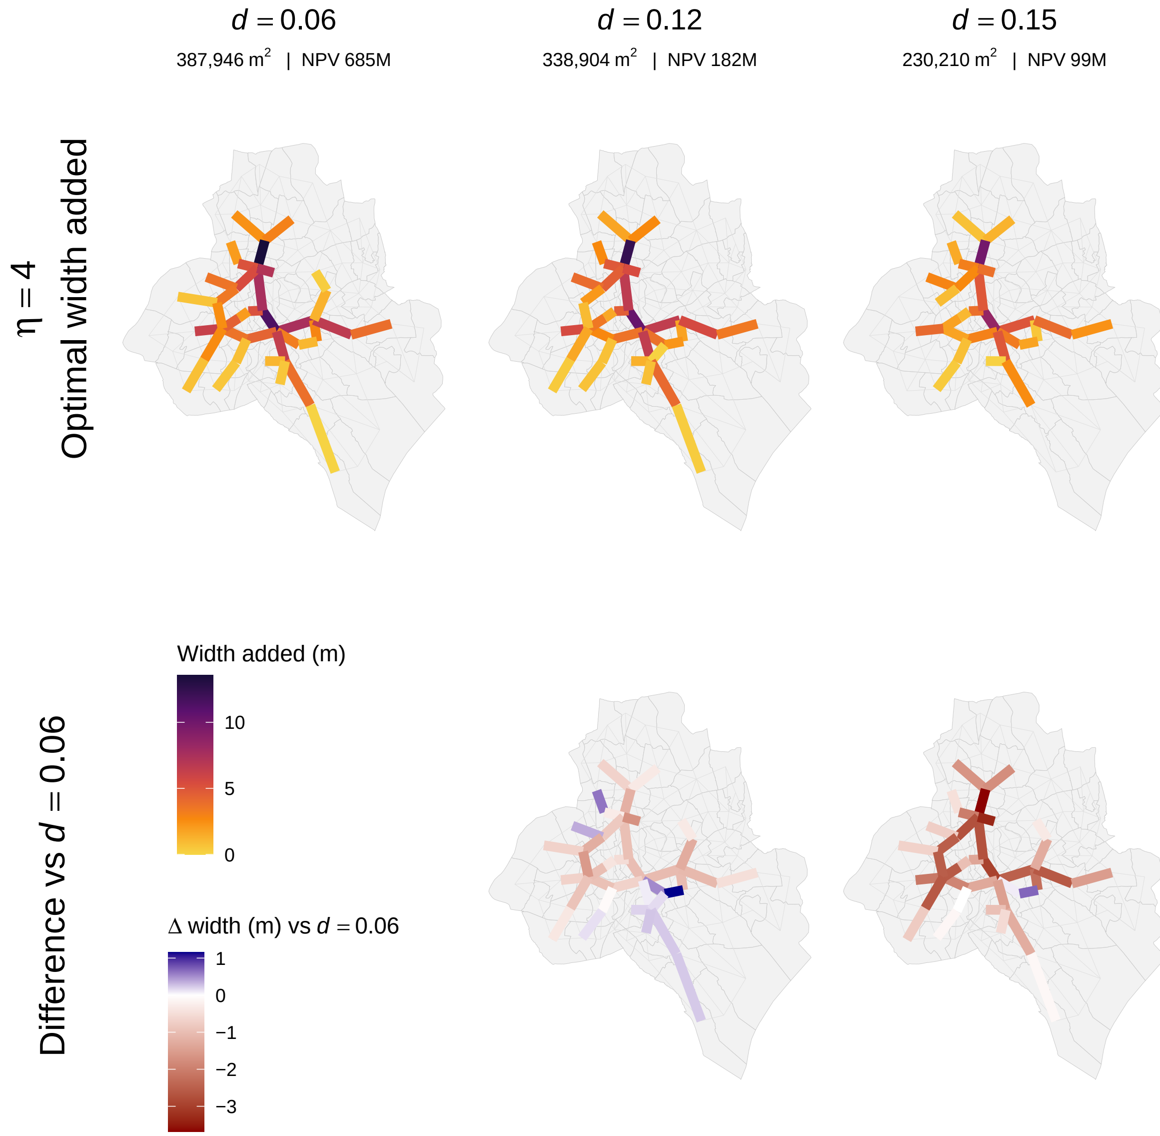

Notes: Top row: optimal width added to each link (metres) at $\delta = 0.06$, $\delta = 0.12$, and $\delta = 0.15$. η is fixed to 4.0 across all maps. *Bottom row:* the change relative to $\delta = 0.06$. A higher discount rate leaves the land-cost wedge—a stock—unchanged but lowers the present value of the commuting benefits—a flow—so the wedge binds sooner and investment contracts. Panel subtitles report the total widened area (m^2) and the net present value of the program (US$ millions).

Figure A22: Marginal Income and Substitution Effects of Land Compensation Rules

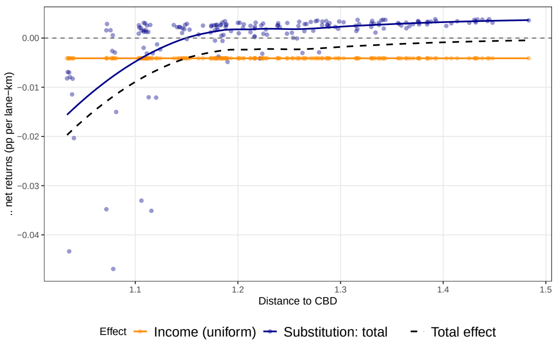

(a) No compensation to market value compensation

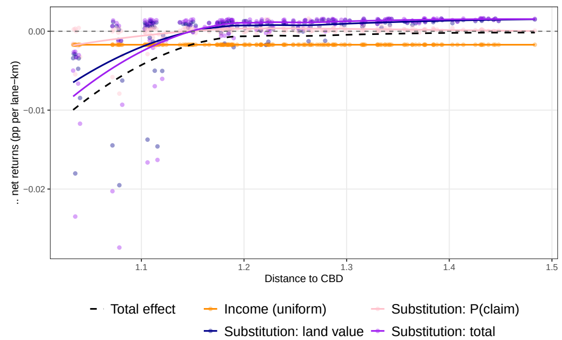

(b) No compensation to claim-based compensation

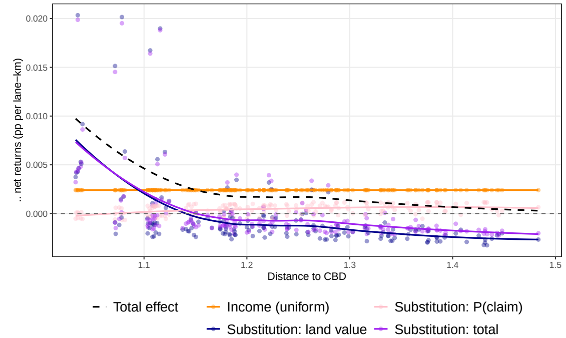

(c) Market value to claim-based compensation

Notes: Each panel decomposes the marginal welfare effect of switching land-compensation rule into a uniform *income* (level) effect and link-specific *substitution* (reallocation) effects. Each point is a candidate road-widening link, plotted against its distance to the CBD; solid curves are local-polynomial smooths of the components, the black dashed line is their total, and the grey dashed line marks zero. The vertical axis is the change in the net marginal returns associated to widening this link (welfare net of compensation costs), in percentage points of baseline welfare per added meter of road. Panel (a) compares the market value compensation against no compensation. Panel (b) compares the claim-based compensation against no compensation. Panel (c) compares the claim-based compensation against the market value compensation. With η the shadow price of public funds ($\eta = 4$), F_l the per-meter fiscal cost of market-value compensation, $\bar{F}$ its mean, and $\bar{p}$ the land-payment–weighted mean claim probability (42%), the wedge $(\eta - 1)/\eta$ is the share of each compensation dollar lost to distortionary financing. The total effects are $-\frac{\eta-1}{\eta}F_l$ (a, no comp $\to$ MV), $-\frac{\eta-1}{\eta}p_l F_l$ (b, no comp $\to$ claim), and $\frac{\eta-1}{\eta}F_l(1-p_l)$ (c, MV $\to$ claim); the income term evaluates these at $\bar{F}, \bar{p}$, and the substitution terms collect the link-specific deviations in land value ($F_l - \bar{F}$) and, in (b) and (c), claim probability ($p_l - \bar{p}$). Because owners value transfers below their fiscal cost as they are wealthier than workers, compensation lowers net returns (a, b negative on average) while moving from blanket market value to claim-based raises them (c positive).

Table A22: Net returns to the realized upgrades under alternative parameters

	Realized GE outcome		Net welfare (USD m), by compensation			
	Δ commuting time (%)	Δ rental income (%)	No payment	Payment (lower)	Payment (upper)	Market value
Main ($\eta = 4,\ \delta = 0.06$)	-3.6	0.7	172	118	94	35
Cost of public funds η						
$\eta = 1$	-3.6	0.7	172	172	172	172
$\eta = 2$	-3.6	0.7	172	154	146	126
$\eta = 6$	-3.6	0.7	172	82	42	-56
$\eta = 8$	-3.6	0.7	172	45	-10	-147
Discount rate δ						
$\delta = 0.03$	-3.6	0.7	401	347	323	264
$\delta = 0.12$	-3.6	0.7	57	3	-21	-80
Speed–capacity elasticity ξ						
$\xi = \frac{1}{2}\hat{\xi}$	-1.7	0.4	13	-51	-80	-150
Commuting split ($\theta\kappa$ held at estimate)						
$\kappa = 0.01$ (Berlin), θ implied	-3.4	0.7	183	118	89	19
Owner utility						
Log	-3.6	0.7	157	92	63	-7

Notes: Net returns to the realized road upgrades under alternative parameter values. Each row changes a single parameter from the main specification (shadow price of public funds $\eta = 4$, discount rate $\delta = 0.06$, speed–capacity elasticity $\xi = \hat{\xi}$, estimated θ/κ split, linear owner utility, status-quo financing); all other parameters are held at their main values. The first two columns report general-equilibrium outcomes: the percentage change in average commuting time and in citywide rental income, baseline to post-policy. The remaining columns report the net present value of the upgrades—the present value of resident and owner welfare gains, less construction cost and η times land compensation, in millions of USD—under four compensation assumptions: no payment ($\mathbf{P} = 0$); the lower and upper empirical bounds on the share of owners compensated (lower = $\mathbf{P}$(received), upper = $\mathbf{P}$(ever compensated)), which bracket the realized policy; and full market value ($\mathbf{P} = 1$). $\xi = \frac{1}{2}\hat{\xi}$ halves the estimated speed–capacity elasticity; the κ row holds the estimated product $\theta\kappa$ fixed and sets κ to the Berlin value (Ahlfeldt et al., 2015), implying $\theta = \theta\kappa/\kappa$; the log row replaces linear owner utility with logarithmic. Negative entries indicate the policy is net welfare-reducing under that compensation assumption.

G.1 Non-Linear Owner Utility

In the main specification, owners' utility is linear in income. Here, I solve the model for another parameterization of the CRRA utility. I set $\psi^o = 1$ so owners' utility is equal to $W_i^o \equiv \bar{C}^o \ln l_i^o$. Additionally, I set the reference income to the average wage $\bar{C}^o \equiv \bar{w}^{\psi^o}$, so that $\frac{\partial W_i^o}{\partial l_i^o} = \frac{\bar{w}}{l_i^o}$ and $= 1$ at $l_i^o = \bar{w}$. For rental income l_i^o greater than the average worker wage $\bar{w}$, owners' marginal utility of income is less than 1.

The NPV of the net returns of the realized road upgrades under this alternative owner utility parameterization is displayed in Appendix Figure A23. The patterns are similar as under linear owner utility, with market-value compensation leading to negative net returns for a high enough $\{\eta, \delta\}$ combination. Noticeably, while under linear owner utility, $\eta = 1$ led to similar net returns no matter the compensation rule because it removed all wedges, with log owner utility, even under $\eta = 1$, market-value compensation is associated with lower net returns of the policy because compensation is a regressive policy: for every dollar taken to society as a whole to give to wealthy owners, the marginal utility gains are only $\frac{\bar{w}}{l_i^o} < 1$ dollar equivalent.

Figure A23: Net present value of the realized upgrades under *log* owner utility.

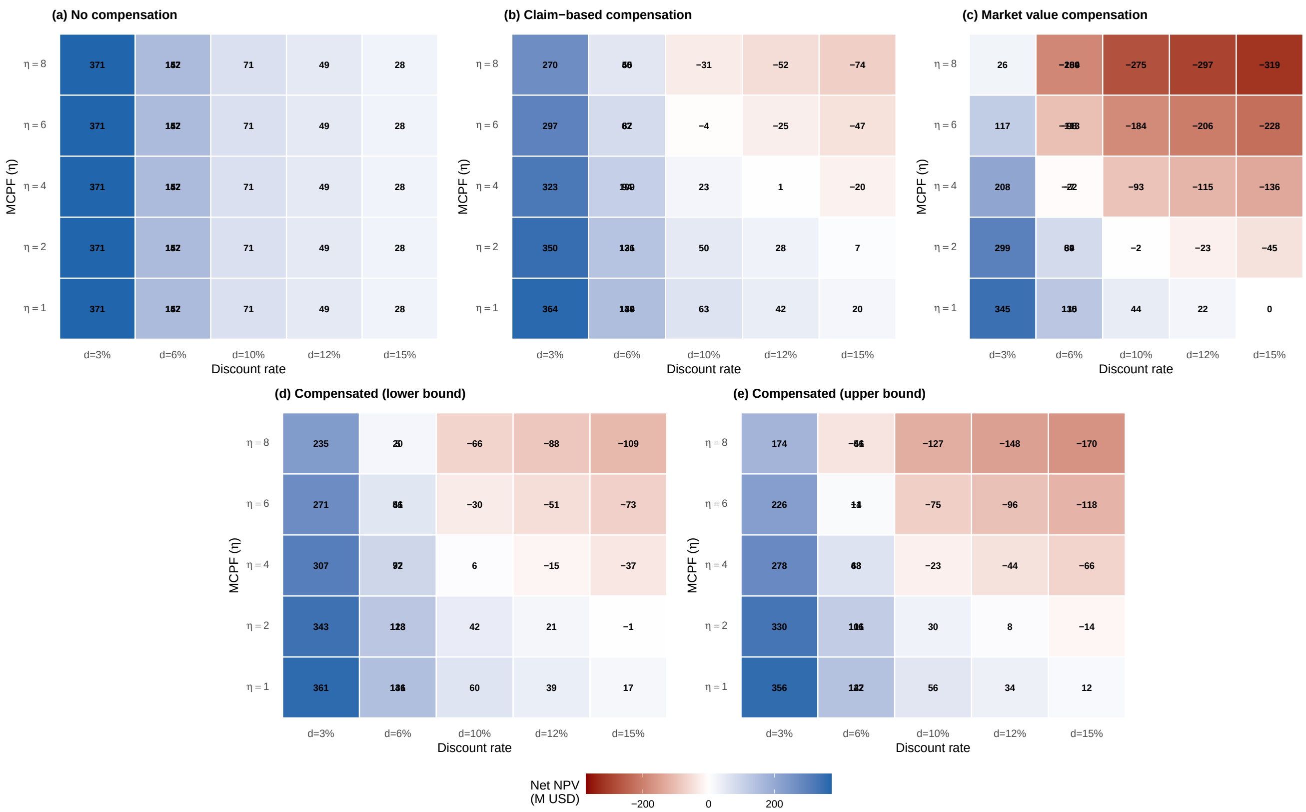

Notes: As Figure 6, but with logarithmic owner utility (diminishing marginal utility of owner income) in place of the linear, money-metric benchmark used in the main text. Each cell reports the net NPV (million USD) of the realized external upgrades for a given cost of public funds η (rows) and discount rate δ (columns). Top row: the three compensation rules—no compensation, claim-based, and market value; bottom row: the two empirical bounds on the share of owners compensated (lower $= \mathbf{P}$(received), upper $= \mathbf{P}$(comp)). Owner welfare is anchored to be money-metric at the mean wage, so $\partial W^o/\partial l = 1$ at income equal to the average worker's; relative to the linear benchmark, the log specification additionally penalizes transfers to owners, who are richer than average.

Table A23: Net returns to the optimal allocation under alternative parameters

	Net welfare (USD m)	Road area built (% env.)	% of first best
$\eta = 1$ (first best)	983	100	100
Main — market value ($\eta = 4,\ \delta = 0.06$)	685	100	70
Cost of public funds η			
$\eta = 2$	882	100	90
$\eta = 6$	492	100	50
$\eta = 8$	314	78	32
Discount rate δ			
$\delta = 0.03$	1,701	100	85
$\delta = 0.12$	182	87	38
Speed–capacity elasticity ξ			
$\xi = \frac{1}{2}\hat{\xi}$	68	40	24
Commuting split ($\theta\kappa$ held at estimate)			
$\kappa = 0.01$ (Berlin), θ implied	869	100	75
Owner utility			
Log	506	100	60
Compensation rule			
Claim rule, endogenous P(claim)	874	100	89
Claim rule, exogenous P(claim)	863	100	88
Funding rule			
Construction & land at η ($\eta = 4$)	506	100	52
Construction & land at η ($\eta = 6$)	244	59	25

Notes: Net returns to the welfare-maximizing (optimal) road-widening allocation under alternative parameter values, re-solving the planner's problem for each row under market-value compensation and status-quo financing. Each row changes a single parameter from the main specification ($\eta = 4$, $\delta = 0.06$, $\xi = \hat{\xi}$, estimated θ/κ split, linear owner utility); all else held at its main value. Columns report the net present value of the optimal program (million USD), net welfare per working-age resident (USD), the road area built as a share of the realized construction envelope, the share of the first-best gains captured, and the percentage change in average commuting time. The first best is the no-compensation allocation at $\eta = 1$ (defined in Section 8); "% of first best" is each row's net welfare relative to that benchmark, computed at the row's own δ, ξ, κ, and owner utility, so under linear owner utility the $\eta = 1$ row equals 100%. $\xi = \frac{1}{2}\hat{\xi}$, the κ row, and the log row are defined as in Table A22. The two claim-rule rows re-solve the optimum under claim-based compensation: "endogenous **P**" lets the claim probability respond to the taking, as in the main model, while "exogenous **P**" freezes it at the realized one-lane predicted level; their near-equality isolates the policy-feedback channel (which depends on the imprecise α) from the spatial level of claiming.

G.2 Robustness to the Speed-Capacity Elasticity ξ

ξ scales the present value of accessibility benefits against the fixed dollar-denominated cost of land and construction. To assess the sensitivity of the results to ξ, I re-solve the model at half the estimated elasticity, $\xi = \frac{1}{2}\hat{\xi}$ holding all else fixed (Tables A22 and A23).

The paper's *comparative* conclusions hold and even sharpen. More generous compensation still lowers net returns monotonically; market-value compensation turns them sharply negative even at $\eta = 4$, $\delta = 0.06$

(against a positive return at the estimated ξ), so the reversal sets in at a lower cost of public funds; and in the optimal allocation the welfare-maximizing program under market value captures a *smaller* share of the first best (24% against 70%), because the dollar-denominated fiscal wedge looms larger relative to the smaller benefits. The misallocation from land payments is thus worse, not better, at a lower ξ.

The *level* of the realized net gain is sensitive to ξ. Pricing only the closed-city accessibility channel, the realized upgrades' net return at $\frac{1}{2}\hat{\xi}$ is positive when owners are uncompensated (USD 13 million) but turns negative once the realized compensation is paid (USD -51 to -80 million) and sharply negative under market value (USD -150 million; Table A22). These figures reflect benefits of road upgrades coming only through this low ξ, while open-city, out-of-city and local-amenity benefits are omitted from the model, so these numbers should be interpreted as lower bounds.