

University of Kent

School of Economics Discussion Papers

Highway Improvements and Local Economic Development: Evidence from Tanzania's Central Corridor

Amrit Amirapu

Mumba Ngulube

Xinxin Lyu

August 2026

KDPE 2605



Highway Improvements and Local Economic Development: Evidence from Tanzania's Central Corridor*

Amrit Amirapu[†]

Mumba Ngulube[‡]

Xinxin Lyu[§]

July 2026

Abstract

Transport corridor investments absorb enormous capital in settings where fiscal space is scarce, yet rigorous evidence on their economic returns remains thin—particularly in Sub-Saharan Africa. We estimate the local economic effects of upgrading two segments of Tanzania's Central Corridor, using a staggered difference-in-differences strategy on a panel of 10 km × 10 km grid cells covering 2003–2018. We take satellite-derived per capita consumption (Huber and Mayoral, 2024) as our primary outcome and define control units as those grid cells proximate to other major highways that were not exposed to upgrades. We find that the two upgrades have considerably different effects: relative to the control-group trend, consumption near the Singida–Shelui segment (completed 2009) increases steadily until it is 14 percent higher after eight years, while consumption near the Lushunga–Kagera segment (completed 2013) falls around 3 percent over the five years we observe post upgrade. Our central finding is that neither effect is homogeneous within its catchment. Initially larger and richer cells capture the bulk of the gains along Singida–Shelui, while smaller and poorer cells account for all of the negative effects along the Lushunga–Kagera improvement. Better roads may therefore not merely distribute gains unequally: where the local economy is thin, improved connectivity can leave an area absolutely worse off. Secondary findings include the fact that population falls persistently near both segments and in rich and poor areas alike, which is consistent with out-migration toward newly accessible urban centres. The results hold against an alternative control group drawn from cells near planned-but-uncompleted projects, while a synthetic difference-in-differences specification preserves the Singida–Shelui gains while rendering the Lushunga–Kagera decline statistically insignificant.

*We thank the UK Foreign Commonwealth Development Office in Tanzania for providing financial support to the World Bank to undertake this research, and the Government of Tanzania and its Agencies for their collaboration in providing data and feedback on results. Special thanks also to Nana Soetantri, Bahar El Hifnawi and Melvin Daniel Mndeme from the World Bank for support and assistance – and to Muneeza Mehmood Alam, Nino Pkhikidze, Michele Valsecchi, Diego Marino Fages, Bansi Malde, and workshop participants at DEW Durham and the University of Kent for very helpful feedback and suggestions. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily represent the views of the World Bank, their Executive Directors, Development Partners or the countries they represent. The authors may be contacted at a.amirapu@kent.ac.uk, mngulube@worldbank.org and xlyu@worldbank.org.

[†]University of Kent, Department of Economics.

[‡]World Bank Group.

[§]World Bank Group.

1 Introduction

Investing in transport corridors requires immense capital, creating challenging trade-offs for developing nations already navigating constrained budgets and competing demands from health, education and social protection. The rationale for prioritizing these investments extends beyond their direct user benefits (e.g., improved transit times, greater reliability, reduced costs of moving goods and people). A growing body of evidence suggests that strategic corridor development can unlock broader economic and social gains: stimulating trade, accelerating structural transformation, creating employment, and improving welfare across large geographic areas.

Sub-Saharan Africa (SSA) may present the most compelling case for investing in transport corridors. The region faces high poverty rates alongside a widening infrastructure deficit. The United Nations projects the region's population to nearly double to 2.5 billion by 2050, driving rapid growth in demand for mobility, trade, and economic services. Meanwhile, the transport sector is characterized by fragmented transport networks, weak intermodal connectivity, and logistics costs that limit market integration, suppress trade, and leave rural populations with inadequate access to jobs, markets, and essential services (Foster and Briceño-Garmendia, 2010; Beuran et al., 2015). Given these pressures, it is not surprising that national governments, philanthropy organizations, bilateral donors, and multilateral institutions continue to prioritize investments in transport. However, evidence on how transport corridors shape economic outcomes remains surprisingly thin for SSA, leaving policy makers without the empirical foundation needed to design, sequence, and evaluate investments of this scale.

What credible evidence does exist comes overwhelmingly from large national programmes evaluated as a whole. Donaldson (2018) shows that access to colonial India's railway network lowered trade costs and raised real incomes by around 16 percent in the districts that received it. Faber (2014) finds that connecting peripheral Chinese counties to the National Trunk Highway System *slowed* their growth relative to comparable counties left unconnected. Asher and Novosad (2020) evaluate India's rural road programme and find that new roads move workers out of agriculture without measurably raising village income, assets, or consumption. Like us, Gertler et al. (2024) study improvements in road quality rather than new construction, and find shifts from informal into formal employment in Indonesia.

For Sub-Saharan Africa the evidence is thinner and more aggregated still. Jedwab and Moradi (2016) show that colonial railways reshaped the distribution of economic activity across the continent and that the pattern persisted long after the railways themselves fell into disuse. Jedwab and Storeygard (2022)

provide the most comprehensive contemporary estimates, using road construction in distant locations as a source of variation in market access across 39 countries over five decades, and report a thirty-year elasticity of city population with respect to market access of roughly 0.08 to 0.13. Storeygard (2016) uses oil price fluctuations interacted with distance to the primary port to show that transport costs constrain the growth of African cities, especially those far inland. Closest to our own setting, Aggarwal et al. (2024) study northern Tanzania and find that improvements to main roads raised agricultural technology adoption by more than improvements to rural roads did.¹

That transport improvements have uneven effects is not a new idea. Economic geography has long held that falling transport costs between a core and a periphery may concentrate activity in the core rather than disperse it (Krugman, 1991; Redding and Venables, 2004), and some empirical work bears this out: Faber (2014) finds losses in the connected periphery, while Jedwab and Storeygard (2022) report that market-access elasticities are larger for small, remote cities (and smaller in politically favoured and agriculturally productive areas). What this literature characterises, though, is heterogeneity across regions—different cities, different districts, different countries—recovered by comparing many treated regions with many untreated ones.

Two other dimensions have received much less attention, largely because the settings that support credible identification seldom permit them. The first is heterogeneity *across segments of the same network*. Testing this requires detailed project-level information (e.g., when each segment was completed), which programme-level evaluations typically neither observe nor exploit. The second is heterogeneity across narrowly defined units within the same region. Testing this requires outcome data at a spatial resolution finer than the country-, district- and city-level data on which most of this literature rests. A practical constraint compounds the second problem: household surveys are rarely available at fine resolution over long panels, so work in low-income settings has leaned on satellite-measured nighttime luminosity as a proxy for local activity. However, this proxy has been shown to be limited – especially in the context of SSA, where large numbers of populated cells emit no detectable light at all – and subject to potential biases (Huber and Mayoral, 2024).

This paper seeks to fill these gaps by examining the impact of upgrading two sections of Tanzania’s Central Corridor on local economic activity, consumption and migration. Drawing on high-resolution geospatial data and satellite-derived welfare estimates covering the period 2003 to 2018, we analyse how impacts evolve over time, offering new evidence on both the magnitude and the temporal dynamics of corridor investment returns

¹This is a deliberately narrow selection. For broader surveys of what transport investment does to firms, trade, agriculture, and human capital, see Redding and Turner (2015) and Foster et al. (2023); we return to the specific channels in Section 3. Prior work on the Central Corridor itself is descriptive rather than causal (Okonge and Nyaronga, 2025).

in a low-income African setting. We use the Callaway and Sant'Anna (2021) Differences in Differences estimator, which recovers group-time average treatment effects *in the presence of treatment heterogeneity* and allows us to trace the full dynamic path of impacts. Treatment is assigned to grid cells within 25 km of a project segment;² and control cells are those that are within 25 km of non-Central Corridor highway segments and at least 50 km from any completed transport project in mainland Tanzania.

To address potential endogeneity concerns related to the non-random selection of road segments, we employ a four-pronged identification strategy. First, we exclude projects in Dodoma and Dar es Salaam to prevent these high-growth urban hubs from biasing our estimates. Second, we use event-study models to verify the absence of differential pre-trends between treatment and control groups.³ Third, we check robustness against an 'alternative control group' using grid cells near "planned-but-uncompleted" infrastructure projects as a counterfactual for the growth potential of areas selected for treatment – in the absence of treatment. Finally, we re-estimate the headline effects using the synthetic difference-in-differences estimator of Arkhangelsky et al. (2021), which reweights the control group so that its pre-treatment trajectory is parallel to that of the treated group by construction.

Our main findings, focusing on (logged) per capita consumption as our primary outcome, are that the two upgrades had sharply different local effects, and that pooling them obscures rather than summarises what happened. Estimated over both projects together, consumption in treated grid cells appears to decline slightly before rising steadily to reach roughly 14 percent above the control-group trend eight years after completion—suggesting a large but delayed payoff. Separating the projects reveals that this delay is an artefact of composition rather than a feature of the dynamics. The Singida–Shelui upgrade (completed 2009) generates gains almost immediately, rising from around 3 percent in its second year to 14 percent by its eighth, with no sustained initial dip. The Lushunga–Kagera upgrade (completed 2013) instead shows negative effects of around 3 percent across the five years for which we can observe it.

We also look at nighttime lights and population as secondary outcomes. Nighttime light effects are more muted and sensitive to the control group specification, partly reflecting well-documented measurement concerns with satellite luminosity data in Sub-Saharan Africa – including the fact that the analysis relies on a very small number of lit cells. Population levels near treated roads fall persistently for both projects, consistent with net out-migration toward newly accessible urban centers, though these results should also

²We also use 50 km thresholds for treatment and control assignment in secondary analyses.

³It is well understood that the absence of differential pre-trends is neither necessary nor sufficient for the main DiD identifying assumption to hold, but we take the presence of approximately parallel trends in the pre-period as strengthening our confidence that the identifying assumption holds (Kahn-Lang and Lang, 2020).

be interpreted with caution given known limitations of modelled population estimates (see Section 4.1). Findings using a 50 km buffer are smaller in magnitude but remain broadly consistent for all outcomes. Finally, effects vary within the treated areas and not only between the two projects: gains are concentrated in cells with higher initial total consumption (per capita consumption multiplied by population), while initially smaller and poorer cells show effects close to zero—or negative, in the case of the Lusahunga–Kagera upgrade.

This paper makes four primary contributions—three concerning heterogeneity in the effects of road upgrading, one concerning how those effects are measured—plus a methodological observation. First, we show that the local effects of two upgrades to the same corridor differ not just in magnitude but in sign. Singida–Shelui raised consumption steadily, reaching some 14 percent above the control-group trend within eight years; Lusahunga–Kagera, upgraded four years later on the same route, shows consumption effects of around -3 percent across the five years we can observe. In engineering terms the two interventions were closely comparable: each replaced a gravel trunk road with one built to bitumen standard, on the same corridor and under the same national roads agency.⁴ That similarity suggests that the observed divergence in outcomes is most likely a function of differences in local conditions or in complementary investment that we do not observe.

Second, we find heterogeneity within the treated areas as well as between them. Splitting treated cells by their initial level of total consumption, the gains near Singida–Shelui are concentrated in cells that were already larger and better off when the road was completed, while initially smaller and poorer cells show much smaller effects. In the case of Lusahunga–Kagera, the smaller/poorer cells account for all of the project’s negative effects. This is similar to what the core–periphery tradition anticipates and what Faber (2014) finds across Chinese counties. It also sits in tension with Jedwab and Storeygard (2022), who report the strongest market-access effects for small and remote African cities; the two results concern different outcomes and different units at vastly different geographic scales, so they need not conflict, but the contrast is worth noting. Observing this dimension at all requires outcome data much finer than the treated band, which is what the grid-cell design provides.

Third, we use satellite-derived consumption estimates (Huber and Mayoral, 2024) rather than nighttime lights as our primary outcome, and our setting is a sharp illustration of why that choice matters. Log

⁴Singida–Shelui was gravel surfaced before the works, which consisted of “upgrading to bitumen standard” (World Bank, 2004). Kagoma–Lusahunga was likewise an upgrade of the existing gravel road to bitumen standard. The two differ in financing: Singida–Shelui was World Bank financed, whereas Kagoma–Lusahunga was completed under a Government of Tanzania contract running from July 2009 to December 2013, after an earlier African Development Bank contract (April 2006) was terminated in February 2008 for contractor non-performance. See Section 2.

luminosity is undefined for cells emitting no detectable light, and in our data such cells are the large majority of the treated sample: of 165 treated cells at the 25 km buffer, only 21 carry usable night-light data—six of them for Singida–Shelui. We therefore report most night-light results in an appendix with explicit caveats rather than in the main text. The consumption series is available for all grid cells throughout.

Fourth, we show that these effects are hyper-local. Re-estimating the event study across cumulative buffers from 10 to 60 km, the Singida–Shelui consumption effect is largest at the narrowest threshold—19 log points by Year +8 within 10 km of the investment—and decays steadily as the buffer widens, reaching roughly 4 log points at 60 km. Cumulative buffers cannot say where the effect stops, however, since each one contains all the cells inside it. We therefore re-estimate the model separately within each marginal 10 km ring, restricting the control group to cells the same distance from a comparable non-upgraded highway so that the comparison is like-for-like at every distance. Effects are large in the 0–10 and 10–20 km rings, smaller but still positive in the 20–30 km ring, and indistinguishable from zero in the 30–40, 40–50 and 50–60 km rings. The gains from upgrading a highway, in other words, accrue almost entirely within about 30 km of it, and the gently decaying gradient in the cumulative buffers is largely an artefact of averaging a concentrated effect over a progressively larger area of cells that experience nothing.

This matters for three reasons. It is visible only at fine spatial resolution: a district- or city-level design would average the effect over its unaffected hinterland and recover something attenuated toward zero, which may help explain why some estimated effects in this literature are modest. It supports our identification, since a spatial footprint that vanishes by 30 km means control cells located beyond 50 km are very unlikely to be contaminated—and indeed tightening the exclusion zone to 100 km leaves the estimates essentially unchanged. And it matters for appraisal: the population that actually benefits from a corridor upgrade is far smaller than the corridor’s nominal catchment, so benefit-cost calculations that assume a wide zone of influence may overstate returns.

Finally, the two projects together illustrate a practical point about evaluation. Because they are staggered and their effects differ in sign, the pooled estimates trace a path that neither project follows: benefits appear to take five years to arrive, when in fact the earlier project delivered them almost immediately and the later one did not deliver them at all. The same heterogeneity drives a divergence between estimators, with pooled two-way fixed effects estimates of the consumption effect negative and statistically insignificant while the Callaway and Sant’Anna (2021) estimates are positive and significant. Neither point is novel in itself (Goodman-Bacon, 2021; Sun and Abraham, 2021). But corridor investments are routinely appraised and evaluated at programme level (Melecky et al., 2019), and this is a concrete case of what that practice can

obscure.

The rest of the paper is structured as follows. Section 2 provides some background on the context, including the road projects under study. Section 3 articulates the conceptual framework. Section 4 describes the data and variables used in the study. Section 5 outlines the identification strategy and the econometric methodology. Section 6 discusses the empirical results, beginning with the pooled estimates before decomposing them by project and treating the two upgrades separately thereafter. Section 7 presents robustness checks varying the treatment buffer distance and the spillover exclusion zone, re-estimating the headline effects with the synthetic difference-in-differences estimator, and assessing the sensitivity of the results to violations of parallel trends. Section 8 provides concluding remarks, including estimates from project-level cost-benefit analyses.

2 Background and Road Projects

This study analyses the local economic effects of improving two important sections of Tanzania's Central Corridor: the 110 km Singida–Shelui section and the 154 km Lushunga–Kagera section. The central transport corridor is one of several critical transport corridors in east Africa, and is seen to be vital for economic growth and regional integration. As a primary trade artery, this corridor connects six states (Tanzania, the Democratic Republic of Congo, Uganda, Rwanda, Burundi, and Zambia) to global markets via the port of Dar es Salaam. Historically, the selection of sections for improvement along the Central Corridor were prioritized based on their strategic role of linking vast rural production areas, mining, and tourism hubs to major urban centers and border points.

The Singida–Shelui section was upgraded under a World Bank-financed project of US\$122 million, of which US\$55.73 million was allocated to the improvement of this specific section. Approved in 2004 and completed in 2009, the project aimed to enhance the performance of the Central Corridor by upgrading strategic links along the route. At the time of project appraisal, the Singida–Shelui road represented the only remaining unpaved section of the Central Corridor road between Dar es Salaam and Lushunga. The segment's upgrading to bituminous standard was expected to complete the corridor's full rehabilitation, strengthening connectivity between the port of Dar es Salaam and the landlocked countries of Burundi and Rwanda, and between Tanzania's western and lake regions with its capital city. The improved road conditions were expected to facilitate economic growth in the project area and economy at large by reducing transportation costs and improving access to markets and jobs.

The rehabilitation of the 154 km Lusahunga–Kagera section was financed through public funds at a total cost of US\$120 million. Approved in 2009 and completed in 2013, this project extended the improvements further along the corridor, aiming to consolidate the gains in regional connectivity achieved by earlier investments in the corridor such as the Singida–Shelui upgrade.⁵ Together, these two sequentially implemented interventions provide a natural setting to apply a staggered differences-in-differences analysis of the local economic effects of highway improvements. The two selected treatment projects are described in Table 1, and their paths are mapped in Figure 1.

Table 1: Final Treatment Projects

Label	Project Name	Period	Length	Value (USD)	Funding
A	Singida–Shelui	2004–2009	110 km	\$56m	World Bank (Public)
B	Lusahunga–Kagera	2009–2013	154 km	\$120m	Public

Notes: Project A was sourced from World Bank records (ID: PO78387_sub); Project B from the Global Construction Database (ID: 258559) with supplementary information from TANROADS records. Both projects involve road upgrades from gravel to bitumen standard.

Selection of road projects in the study

To identify projects to include in the study, we reviewed transport projects implemented and completed along Tanzania’s Central Corridor between 2000 and 2018⁶ in the World Bank database and the Global Construction database. These two databases provided a combined total of about 375 transport projects⁷ implemented in Tanzania in the last 20 years.


The World Bank Tanzania Transport Database. The World Bank database includes ten projects approved since 2000 of which two were completed at least 5 years before 2018. The first project is the Singida–Shelui project discussed above. The second project involved the upgrading of a 21 km road section within Dar es Salaam. We exclude the latter from our analysis because Dar es Salaam, a large and rapidly growing metropolitan area, may operate on a distinct economic trajectory relative to other locations in Tanzania, and any estimated effects would likely be confounded by the endogenous forces driving infrastructure investment decisions in the city, rendering causal identification of road improvement effects challenging.

⁵The Lusahunga–Kagera works were initially started in 2006 with financing from the AfDB, but the first contract was terminated and the contractor removed from the project in early 2008 "[d]ue to persistent failure to perform his obligations under the Contract" – according to records from The Tanzania National Roads Agency (TANROADS). A contract for the same project was then reissued with financing from the Tanzanian Government in 2009.

⁶This period is selected to coincide with data requirements and exclude potential effects from the Covid-19 pandemic period.

⁷10 projects under the World Bank Portfolio and 365 Global Construction Data.

Figure 1: Treatment Project Paths



The Global Construction Database. This database is a commercial project-level database widely used in empirical infrastructure research, provided by GlobalData Plc. It covers more than 20,000 road projects across 191 countries, spanning back to the 1980s and offering near-universal reporting after 2010 for projects above USD 25 million. The project information is updated regularly from systematic web scraping. Each entry records technical and financial attributes, including road length, connected cities, associated facilities, ownership structure, project value, lifecycle stage (planning, construction, on hold, completed or cancelled), project type (rail, road, or port), financing source (public, private, or both), and a description of project objectives. The key strength of this dataset lies in its granular, verified project-level monitoring, which enables construction tracking from announcement to completion and facilitates the identification of treated grid cells along the central corridor. We also use the project status information to construct an alternative

control group based on locations proximate to projects that were planned but not completed (e.g., because they were suspended or cancelled). The 2025 database contains 365 transport projects⁸ in Tanzania. Each project is linked to a primary location with detailed geographic coordinates and, where applicable, the start and end points of roads under construction. Projects in the Global Construction Database were selected for inclusion in the study based on three criteria:

1. *Road projects within 25 km of the central corridor.* First, we restrict the sample of projects to those whose coordinates are within 25 km of the Central Corridor⁹ to ensure that projects we capture actually affect the corridor itself. We define the Central Corridor following the East African Community (EAC) designation, which traces the route from the port of Dar es Salaam on the Indian Ocean through key inland nodes—Chalinze, Morogoro, Dodoma, Singida, Nzega, and Isaka—and extends along three northwestern branches to the Rwandan border at Rusumo, the Burundian border at Kabanga, and the Ugandan border at Mutukula. We identify 16 key cities along the main Central Corridor route and three branches in the north-western region. The corridor route was constructed by selecting primary trunk road segments from Tanzania’s national road network shapefile that connect these 16 cities, forming a continuous transport network. Key cities along the Central Corridor and its three branches are shown in Table 2.

Table 2: Cities and Towns on the Central Corridor

Central Corridor	Towns/Cities
Eastern terminus	Dar es Salaam
Central nodes	Chalinze, Morogoro, Dodoma, Singida, Nzega, Isaka
Western/Northern branches	Branch (i): Nyakanazi, Lusahunga, Rusumo (Rwanda border)
	Branch (ii): Nyakasanza, Kabanga (Burundi border)
	Branch (iii): Muleba, Bukoba, Kyaka, Mutukula (Uganda border)

Source: Authors based on reviewed materials.

2. *Project lifecycle stage “completed” status.* Second, we restrict the sample to projects with a lifecycle status of “completed”, for obvious reasons.
3. *Project Value, greater than US\$100 million.* We retain only projects with a value exceeding US\$100 million. Larger projects are expected to generate more substantial economic impacts, are more likely

⁸The cleaning process corrected 13 project records, resulting in a dataset with 365 infrastructure projects, all with valid geographic coordinates.

⁹Buffer zones were generated using UTM Zone 36S projection (EPSG:32736) to ensure accurate distance calculations in meters.

to be documented in impact assessments, and represent the capital investments most likely to interact meaningfully with Central Corridor infrastructure.

Applying these three criteria yields a sample of ten projects. Of these, six are located in Dar es Salaam and four are outside the city. We further exclude projects that affect highway segments *connecting* to the central corridor, and instead focus on projects that renovate/upgrade sections along the corridor itself. Finally, we exclude projects located in Dodoma city for reasons similar to Dar es Salaam. This leaves us with a final sample of one project, the Lusahunga–Kagera road project, which constitutes the second project we analyze on Tanzania’s Central Corridor.

3 Conceptual Framework

Upgrading a highway can lower journey times, reduce vehicle operating costs, and make travel more reliable. In Sub-Saharan Africa, where unpaved roads often become impassable during the rainy season and can isolate communities from markets for months at a time, these benefits can be substantial. However, the predicted effects for specific outcomes are not obvious *ex ante* and may be theoretically ambiguous. We outline below the channels through which such cost reductions are likely to affect household consumption, local economic activity, and population distribution—the three outcomes we study empirically—along with specific hypotheses for each outcome.

It is worth distinguishing at the outset the horizons over which these effects arrive. The transport appraisal literature separates first-order benefits accruing directly to road users—shorter journeys, lower vehicle operating costs, improved safety—from the wider economic impacts that follow as firms and households adjust to the new cost structure (The UK Department for Transport, 2016). The latter emerge over the medium term through increased trade, migration, changing land values, and agglomeration, and only over the long run translate into structural change in the local economy (Quium, 2019). Because these later stages unfold over years rather than months, and because their sign is not guaranteed (Melecky et al., 2018), the length of the evaluation window—up to nine years after completion in our case—is a substantive part of the exercise rather than a technical detail.

Consumption. Transport cost reductions can affect household consumption through three reinforcing channels. First, lower transport costs can reduce the price of goods purchased from outside the local area—food, manufactured goods, medicines, agricultural inputs—constituting a direct increase in real

purchasing power for rural households that are net consumers of traded goods. Narrowing interregional price gaps is among the best-documented consequences of transport improvements (Donaldson, 2018; Asher and Novosad, 2020).

Second, improved connectivity can raise the net farmgate price received by farmers, since a smaller share of the sale price is absorbed by transport costs. This may incentivize further production and increase farm incomes. Better connectivity can also widen the set of markets that producers can viably supply, reducing dependence on local (potentially monopsonistic) buyers and facilitating access to modern inputs, technologies, and extension services. Evidence for this channel is fairly consistent. Brooks and Donovan (2020) find that removing seasonal uncertainty in market access in rural Nicaragua raised farm input expenditure and profits substantially. Shamdasani (2021) and Aggarwal (2018) document crop diversification and adoption of modern inputs following India's rural road programme, and Aggarwal et al. (2024) show for northern Tanzania that the binding margin is often the distance to input suppliers located on the main road rather than the quality of local roads. Dorosh et al. (2010) report a similar association between rural road connectivity and crop production across Sub-Saharan Africa.

Third, as commercial and productive activity expands along the corridor, local employment opportunities as well as wages may increase. This channel operates more slowly than the first two, through the gradual structural transformation of local economies, but may be an important source of income growth - particularly for landless and near-landless households. Gertler et al. (2024) document such a transition in Indonesia, where better road quality moved workers from informal into formal employment and raised both income and consumption. Brooks et al. (2021) find that highway expansion in India reduced employers' monopsony power and raised labour's income share by 1.8 to 2.3 percentage points, and Volpe Martincus et al. (2017) find employment and export gains concentrated among small firms after Peru's road programme. The evidence is not uniform: Abbasi et al. (2022) find that road expansion in Sub-Saharan Africa raises skilled employment while reducing low-skilled and agricultural employment, with effects markedly stronger where electricity is also available.¹⁰

Hypothesis 1: Grid cells located in close proximity to upgraded highway sections are likely to experience increases in household consumption relative to less affected cells.

¹⁰Complementarities of this kind recur across settings. In India, road construction combined with electrification raised dry-season cropping by roughly 6 to 12 percent, while either intervention on its own had negligible effects (Eynde and Wren-Lewis, 2025); in Ethiopia, rural roads and agricultural extension services were similarly ineffective in isolation but jointly raised productivity and farm income (Gebresilasie, 2023). We do not observe complementary investments in our setting, and their uneven presence is one reason two otherwise similar upgrades might have different effects.

Local Economic Activity. Reduced transport costs expand the effective market catchment area for firms, allowing them to serve customers and access suppliers across a wider geographic radius. This can attract investment to locations previously too remote to sustain commercial activity. As firms relocate and cluster along the corridor, they may generate further agglomeration benefits—shared labour pools, localized supplier networks, and knowledge spillovers—amplifying the initial market access gains. However, the opposite can also occur, if firms and workers take advantage of lower transport costs to relocate from peripheral areas towards areas with higher initial levels of economic activity, attracted by the greater agglomeration benefits in those locations. Thus, the predicted effect on total economic activity is theoretically ambiguous. Both directions have empirical support. On the expansionary side, Datta (2012) finds that Indian firms near upgraded highways cut input inventories sharply and improved logistics performance; Ghani et al. (2016) find that India’s Golden Quadrilateral raised output, firm entry, and labour productivity in nearby districts, with the largest effects among smaller firms; Abeberese and Chen (2022) find that the same programme led firms to narrow their product range toward what they produced most efficiently, raising productivity; and Duranton et al. (2014) and Coşar et al. (2022) show for the United States and Turkey respectively that better-connected cities specialise further in heavy-goods sectors. On the contractionary side, Faber (2014) finds that peripheral Chinese counties connected to the national highway network grew more slowly than comparable unconnected ones, as improved access exposed local producers to competition from larger centres.

Hypothesis 2: Grid cells in close proximity to upgraded highway sections may exhibit an increase or decrease in nighttime light intensity (a common proxy for local economic activity) relative to less affected cells.

Population. The expected effects of road improvements on the population distribution is less straightforward than for consumption. Like the effects on local economic activity, the direction of the effect depends on the relative strength of competing forces. On one hand, better connectivity can make corridor-adjacent areas more attractive places to live and work, drawing in migrants from more remote locations and retaining residents who might otherwise have left. On the other hand, reduced transport costs lower the barriers to migration more generally, enabling households in peripheral areas to move toward urban centers and economic hubs that were previously too costly to reach. In contexts where the latter force dominates—as may be the case along the Central Corridor, where treatment areas include many rural and peri-urban communities—road improvements may actually reduce local population relative to untreated areas, as improved connectivity facilitates out-migration rather than in-migration. Asher and Novosad (2020) find

that the dominant response to India’s rural road programme was precisely this—workers leaving agriculture and the village—rather than growth in situ, and Jedwab and Storeygard (2022) attribute the market-access effects they estimate for African cities largely to rural-to-urban migration. The force does not always dominate, however: Coşar et al. (2022) find that Turkish road upgrades raised employment in affected sectors with no detectable population response.¹¹

Hypothesis 3: Grid cells in close proximity to upgraded highway sections may experience differential population growth (either positive or negative) relative to more distant cells.

A further consideration cuts across all three hypotheses. The channels described above operate locally, but transport improvements also set in motion spatial reallocation processes whose effects extend beyond the immediate vicinity of the upgraded road. If firms, workers, and economic activity shift toward locations with higher initial productivity—a pattern consistent with the general equilibrium evidence in Morten and Oliveira (2024)—then gains in some areas may be partly offset by losses elsewhere. Our difference-in-differences estimates identify the average effect of treatment for cells near upgraded road segments relative to comparable but less affected cells, and so capture these local effects directly, but may not capture indirect effects in other locations. Moreover, the magnitude and even the sign of the effects we do estimate are likely to vary along the corridor, depending on each area’s initial level of connectivity, its distance to urban centers, and the structure of its local economy.

4 Data and Variables

One contribution of our paper is to study the local economic impacts of road improvements at a highly granular spatial scale using grid-cell-level data. The use of gridded data rather than district-level data is motivated first by the fact the treated road sections span at most 160 kilometers; if the effects of road improvements are primarily localized, then aggregating to districts risks obscuring the spatial variation needed for identification. Helpfully, longitudinal grid-cell data covering the relevant pre- and post-treatment periods are available for all three of our main outcomes of interest.

Specifically, we examine per capita consumption, nighttime light intensity, and population—outcomes selected on the basis of both their theoretical relevance to transport-induced development and the availability

¹¹Road improvements can also affect human capital accumulation, with the direction depending on whether they raise the returns to schooling or the opportunity cost of remaining in it. Aggarwal (2018) finds teenagers leaving school for newly available work, while Adukia et al. (2020) find increased educational investment and a 7 percent rise in middle-school enrolment under India’s PMGSY. We do not observe schooling outcomes and so cannot speak to this margin.

of data with sufficient temporal depth to support the staggered difference-in-differences design.¹² Gridded consumption data are available at a 10×10 km resolution, while population, and nighttime light data are available at 1×1 km resolution but are upscaled to 10×10 km to ensure spatial comparability across all outcome measures. The consumption dataset accordingly serves as the master data file to which all other variables are harmonised, providing a consistent spatial framework for the empirical analysis. Table 3 summarises all outcome variables, their sources, spatial resolution, and temporal coverage.

Table 3: Outcome Variables: Sources and Coverage

Variable	Source	Native Resolution	Data Range
Consumption per capita	Huber and Mayoral (2024)	10×10 km	2003–2018
Nighttime lights	Li et al. (2020), v10 (Li et al., 2024)	1×1 km	1992–2024
Population	Lloyd et al. (2019)	1×1 km	2000–2020

Notes: All 1×1 km datasets are upscaled to 10×10 km to match the consumption grid, which serves as the master spatial reference. Consumption is denominated in ln 2011 PPP dollars.

4.1 Dependent Variables

Consumption Data. Our main outcome variable for estimating the local effects of improved roads on welfare is the gridded consumption data of Huber and Mayoral (2024), made available by the authors. The measure provides estimates of per capita consumption—denominated in 2011 PPP dollars—for 10×10 km grid cells across 42 sub-Saharan African countries over the period 2003–2018.¹³ Relative to traditional survey-based welfare measures, a key advantage of this dataset is its high-resolution calibration for SSA, which allow us to capture granular, local-level welfare variations that could be missed by traditional (purely administrative or survey-based) regional data.

The dataset is constructed using a bottom-up random forest machine learning model trained to predict cell-level per-capita consumption from a rich set of geospatial predictors. The training variable is derived from 85 Demographic and Health Survey (DHS) datasets spanning 29 countries, where individual-level asset indices are translated into consumption dollar equivalents using country-level distribution data from the World Bank’s Poverty and Inequality Platform (WB-PIP). The ten most important predictors in the model, ranked by contribution to prediction accuracy, are: nighttime lights, CO₂ emissions, population density, distance to the national capital, distance to Christian missions, distance to highways, malaria incidence,

¹²Of the three, nighttime lights suffers from the lowest spatial coverage – as explained further below.

¹³In most of our analysis we work with the natural logarithm of the dependent variables, including this one.

latitude, average rainfall, and distance to Protestant missions. Altogether, the model draws on 29 geospatial predictors spanning remoteness, geography, climate, disease environment, and land use characteristics.¹⁴

Huber and Mayoral (2024) cross-check their estimates against alternative data at two levels. Summed to national totals, the correlation between their consumption figures and those from the World Bank (WB-PIP) is 0.85. The more relevant test – for a study like ours – is whether the estimates capture local differences *within* a country. Aggregated to subnational regions, the correlation between their consumption figures and regional income estimates from the Human Development Index is about 0.8.¹⁵

This dataset offers significant advantages over the use of nighttime lights (NTL), which suffer from a variety of problems when used to proxy for economic activity or income—especially in the context of sub-Saharan Africa. In SSA, about 90% of grid cells record zero light at night – even when economic activity is present. This systematic misclassification induces a negative correlation between the measurement error in NTL and true economic development—i.e., non-classical measurement error—, which can bias coefficient estimates in regression models in ways that are unpredictable.¹⁶ The Huber and Mayoral consumption measure addresses these problems in two ways. First, by drawing on a wide range of predictors beyond NTL, it captures economic variation in areas that satellites record as dark, producing estimates with substantially greater spatial variability than NTL-based measures. Second, the authors demonstrate how to apply a simple transformation to their machine-learning predictions that eliminates the non-classical component of measurement error, leaving only classical measurement error—which is far less problematic in regression settings. The resulting variable is denominated in a meaningful, cardinal unit—log consumption dollars per capita—rather than the ordinal and economically uninterpretable lumen-based scale of NTL data.

For these reasons—broad coverage, strict continuity across the full distribution, a cardinal welfare interpretation, a consistent measurement instrument across the study period, and the elimination of non-classical measurement error—we treat per capita consumption as our primary outcome. Nighttime lights and population are included as secondary outcomes for comparison with the broader literature, but all headline findings and robustness checks are based on the consumption series.

Population. To analyse changes in population over time at the local level, we use the WorldPop gridded

¹⁴See Huber and Mayoral (2024), Sections 3.1, 4.1 and Appendix C, for further details including the full set of 29 predictors.

¹⁵See Huber and Mayoral (2024), Sections 4.3–4.4 and 5. The national comparisons cover 85 country-years; the subnational ones cover 244 regions in 20 countries. The authors are careful not to refer to this as a validation exercise, since both their estimates and the external series include "unknown biases and unknown measurement error", but their close agreement is nevertheless reassuring.

¹⁶The severity of this problem is illustrated strikingly by Huber and Mayoral's re-analysis of two prominent papers on institutions and development: substituting their consumption measure for NTL reverses the substantive conclusions of both papers.

population dataset covering the period 2000 to 2020. This dataset provides globally comparable estimates of population counts at fine spatial resolution of approximately $1 \text{ km} \times 1 \text{ km}$ grid cells. WorldPop employs a top-down methodology that disaggregates administrative-unit census counts and population projections to fine-scale grid cells using a random forest-based approach that incorporates geospatial covariates such as land cover, settlement extent, and infrastructure (Lloyd et al., 2019). The disaggregated estimates are adjusted at country total level to match the UN population estimates in 2019 to ensure consistency. Country-specific census data for a reference year serve as the baseline inputs, from which gridded estimates are produced and extended through forecasts for subsequent years. Post-processing adjusts these estimates to conform to UN-projected rural and urban growth rates and to reconcile with UN national population totals.

There are important limitations of this approach in our context. First, the time-series variation in WorldPop is largely model-generated rather than directly observed: because Tanzania's most recent census (2012) falls in the middle of our study window (2003–2018), values both before and after that reference year are extrapolated using UN-projected growth rates rather than independently observed. This means that within-cell variation over time—the key source of identification in our DiD design—may reflect the model's demographic assumptions more than actual population dynamics. Second, validation studies have documented that cell-level accuracy of the unconstrained WorldPop global product is substantially lower in rural areas than in urban ones, with particular limitations in tracking population changes between census years (Boo et al., 2022). Given that our sample is predominantly rural and that our identification relies on temporal variation, these limitations are especially relevant. For these reasons, we treat population as a secondary outcome and interpret the results with caution.

Nighttime lights (secondary outcome). We also use nighttime light intensity as a secondary proxy for local economic activity, in line with a large prior literature. Nighttime light data are commonly used as an indicator of spatial economic expansion, capturing the physical development of the built environment—commercial premises, residential housing, and infrastructure. We include this outcome to facilitate comparison with that literature, but we treat it as secondary for the following reasons.

Non-classical measurement error. As described above and documented formally in Huber and Mayoral (2024), nighttime lights suffer from systematic misclassification in SSA: approximately 90% of grid cells register zero luminosity regardless of their economic activity. This creates non-classical measurement error whose sign and magnitude are unpredictable, and which can bias regression estimates in either direction. It also means that some specifications simply lack the number of observations necessary for plausible analysis.

Sensor transition in 2014. Our study window spans a structural break in the underlying measurement

instrument. Two satellite systems cover different periods: DMSP-OLS (1992–2013) and VIIRS (2012–present). The DMSP-OLS sensor detects only relatively bright sources, so pre-2014 only the 3–6% of cells with the highest luminosity register positive values. The VIIRS sensor is approximately an order of magnitude more sensitive, so from 2014 onwards 30–60% of cells register positive values—including many dim, newly detected cells whose log-transformed values are strongly negative. We use the harmonized series of Li et al. (2020), Version 10, which recalibrates VIIRS observations to a DMSP-like scale using sigmoid functions to create a consistent annual series from 1992 to 2024. However, this recalibration is an approximation: the sharp compositional shift in which cells are positive cannot be fully reversed by a monotonic rescaling. For our DiD design specifically, the sensor transition coincides exactly with the treatment date of Project-B (completed 2013) and falls in the middle of the post-treatment period of Project A (completed 2009), meaning the pre-treatment baseline is estimated almost entirely on DMSP-era data while post-treatment outcomes blend DMSP and recalibrated VIIRS data. Year fixed effects absorb the average level shift at 2014, but cross-sectional heterogeneity in the sensor transition (i.e., which cells newly become positive, and by how much) is not fully absorbed and could contaminate treatment effect estimates.

Saturation in bright areas. The DMSP-OLS sensor imposes a digital number cap of 63, causing saturation in densely urbanised areas and censoring any growth in luminosity above this threshold. If treated areas are more urbanised than control areas on average, this cap may differentially compress variation over time between treatment groups.

For all these reasons, we report nighttime light results for completeness and comparability, but do not draw primary conclusions from them.

4.2 Control Variables and Cell Characteristics

Our baseline specifications include grid-cell fixed effects and year fixed effects, but no additional covariates. The inclusion of grid-cell fixed effects absorbs all time-invariant heterogeneity at the cell level—topography, elevation, soil quality, proximity to pre-colonial trade routes, and any other fixed geographic or historical characteristics that might jointly determine both road placement and local economic trajectories. Identification therefore comes entirely from within-cell variation over time. Year fixed effects absorb aggregate temporal shocks common to all cells, such as national macroeconomic fluctuations, commodity price cycles, or national-scale climate events. Under this two-way fixed effects structure, there is little role for additional time-varying covariates: any such variable would need to be both unaffected by the treatment and capable of

improving the plausibility of the parallel trends assumption by conditioning out differential trends correlated with treatment assignment or cell characteristics. It is hard to find controls that plausibly satisfy both conditions, and so our baseline specifications are parsimonious by design.

5 Methodology

We estimate the local effects of road upgrading using a reduced-form quasi-experimental design. We opt for a reduced form approach rather than structural estimation of a general equilibrium (GE) model for the following considerations. First, the two road segments under analysis are relatively short (110 km and 154 km) and a large part of their economic effects are likely to be inherently localized. Analyzing impacts at a regional or national scale, as a general equilibrium model would typically require, would absorb the corridor-level signal within economy-wide noise and sacrifice the spatial variation that underpins our identification. Second, the availability of high-resolution longitudinal gridded data for consumption, population, and nighttime lights makes a panel fixed-effects approach empirically tractable at the 10 km grid-cell level and allows us to obtain credible causal estimates of the effects of road upgrading on nearby communities, which we believe is of substantial interest in and of itself. A comprehensive simulation of economy-wide welfare changes (taking account of changes to national prices, wages, price volatility, as well as possible upstream and downstream effects), would also be of great interest—though we believe this would be best achieved as a wholly separate analysis.

5.1 Identification Strategy

We study the question above using a DiD identification strategy coupled with a 10 km $\times$ 10 km grid-cell panel dataset for mainland Tanzania covering 2003–2018.¹⁷ In the first instance, grid cells are classified as treated if they lie within 25 km of a completed project segment. The 25 km threshold captures locations likely to share direct economic linkages with the corridor—through connected transport networks, common labor and input markets, and potential agglomeration effects—while remaining tight enough to exclude cells with less exposure. Control cells are those that are within 25 km of a non-upgraded major highway and at least 50 km from any completed project. Results using a 50 km buffer to assign treatment and control units are also reported, to estimate effects over a wider catchment area.¹⁸ A map depicting assignment of grid


¹⁷This range reflects the data availability (particularly the consumption data series) but is also chosen to leave out potential effects from the Covid-19 shock.

¹⁸In Section 7 we report supplementary analyses that explore how the effects vary more flexibly with distance.

cells to treatment or control status under the 25 km threshold is shown in Figure 2.

We employ two control group definitions to gauge sensitivity to the choice of comparison group: the ‘main control group’ uses all qualifying non-treated highway cells; and an ‘alternative control group’ restricts comparison units to cells near planned-but-never-completed projects, thereby conditioning on a similar policy environment with the aim of strengthening the plausibility of parallel trends. A third comparison group—one reweighted to match the treated group’s pre-treatment trajectory—is constructed by the synthetic difference-in-differences estimator as a robustness check in Section 7.5.

Figure 2: Treatment Status by Grid Location (25km Buffer)



Our study focuses on upgrades and renovations to existing road sections rather than the construction of new roads. This is consequential for identification: renovation sites were largely predetermined by the pre-existing highway network—the Singida–Shelui segment was the last unpaved section of the main

corridor, and the Lusahunga–Kagera section renovated infrastructure further along the same established route—rather than selected (potentially in response to differential economic trends). Nevertheless, the timing of renovations/upgrades may have been partly endogenous, and we address this directly.

The core identifying assumption is parallel trends: absent treatment, we assume that outcomes in treated grid cells would have evolved on the same trend as those in control cells. Formally, for each cohort g (defined by year of first treatment) and each post-treatment period $t > g$, the expected change in outcomes for group- g cells—absent the treatment—must equal the expected change for control cells over the same period. Two features of the setting support this assumption. First, project selection followed the logic of completing pre-existing paved routes rather than responding to differential local growth trajectories (World Bank, 2004). Second, control cells are drawn from locations near other major highways that also pre-dated the study period, making treated and control cells broadly comparable in initial connectivity and likely economic trajectory.

We address four specific potential threats to identification. First, anticipation effects: if households or firms responded to project announcements before official completion, pre-treatment outcomes in treated cells would already be diverging, violating parallel trends (and biasing coefficients).¹⁹ We test for this directly with an event study specification and find no evidence of such anticipation effects. Second, spillovers to control units: an improved road may divert economic activity from non-upgraded roads, causing economic activity in control cells to decline and leading us to overstate the treatment effect. We mitigate this by requiring control cells to lie at least 50 km from any completed project segment (including the treatment projects). Third, spillovers to non-control units: road improvements may divert activity from nearby areas that are neither among treated units nor control units—which could also overstate treatment effects. This channel is difficult to rule out entirely, but we shed light on its plausibility by estimating the effects at progressively distant bands from the treated highways, under the assumption that negative spillover effects would occur in proximity to the most intensely treated areas. Fourth, compound treatment: because the two project segments are geographically far apart, no cell in our sample lies within 25 km of both simultaneously under the main buffer specification, so the risk of overlapping treatments is mitigated.²⁰ However, it is possible that all grid cells along the Central Corridor benefit to some extent when any particular segment

¹⁹Pre-treatment outcomes could also diverge if the construction process itself affected outcomes prior to completion.

²⁰The Lusahunga–Kagera upgrade (Project B) *does* partially overlap with a third corridor project between Isaka–Lusahunga, which was completed in 2020. That project is not treated as an intervention in this paper because it completes after our outcome series ends, but it may have been under construction during Project B’s post-treatment window, and may thus have affected Project B’s outcomes. We therefore exclude any cells within 25km/50km of this third project as a robustness test (Section J).

gets upgraded—and indeed some of our evidence is consistent with this interpretation.

5.2 Empirical Specification

An intuitive starting point is to estimate results in a standard two-way fixed effects (TWFE) DiD specification, such as the following:

$$Y_{i,t} = \alpha_i + \varphi_t + D_{i,t} \beta_{\text{post}} + \varepsilon_{i,t} \quad (1)$$

Here, $Y_{i,t}$ denotes the value of outcome Y in grid cell i and year t ; α_i and φ_t represent grid cell and year fixed effects, respectively; and $D_{i,t}$ is a dummy variable taking the value 1 in treated locations after infrastructure projects have been completed, and 0 in all other circumstances.²¹

However, the Singida–Shelui and Lushunga–Kagera segments were completed four years apart, meaning that treatment is staggered across cohorts. When treatments are staggered, standard TWFE estimation is known to perform poorly (e.g., Goodman-Bacon, 2021). In our case the intuition is clear: under TWFE early-treated units act as implicit controls for late-treated units, but if the benefits of road investment build gradually over time, then TWFEs may be downwardly biased because the early-treated group's effects are still growing when later units get treated.²² We therefore adopt the Callaway and Sant'Anna (2021) estimator, which recovers group-time average treatment effects, $ATT(g, t)$, without imposing treatment effect homogeneity across cohorts/treatment groups or calendar periods. This also allows us to trace the full dynamic path of impacts.

We define the group variable g as the year of first treatment: $g = 2009$ for cells near Singida–Shelui and $g = 2013$ for cells near Lushunga–Kagera. For each group-time pair (g, t) , the estimator constructs the Average Treatment Effect on the Treated ($ATT(g, t)$) by comparing outcomes for group- g cells in period t to their counterfactual, proxied by the control group trend, using a doubly robust inverse-probability-weighted regression adjustment.²³ The group-time ATTs are estimated as follows:

$$\widehat{ATT}(g, t) = \frac{1}{N_g} \sum_{i: G_i=g} [Y_{i,t} - Y_{i,g-1}] - \frac{1}{N_{G_{\text{comp}}}} \sum_{i: G_i \in G_{\text{comp}}} [Y_{i,t} - Y_{i,g-1}] \quad (2)$$

where G_{comp} is the set of all never-treated units (i.e., pure controls). The $ATT(g, t)$ estimates are

²¹The analogous event-study specification is $Y_{i,t} = \alpha_i + \varphi_t + \sum_{r \neq 0} \mathbf{1}[R_{i,t} = r] \beta_r + \varepsilon_{i,t}$, where $R_{i,t} = t - G_i + 1$ is the time relative to treatment, and the summation runs over all values of $R_{i,t}$ except 0, the omitted category.

²²A second, related, source of bias enters if the effects of the two projects differ from each other.

²³However, most specifications employ no controls apart from the fixed effects.

then aggregated in two ways: into a single average treatment effect that pools across all post-treatment observations weighted by cohort size (the “simple” aggregate), and into dynamic event-time effects that trace how impacts evolve in the years before and after project completion. The estimator is applied separately for each of the control group definitions, with standard errors clustered at the grid-cell level throughout.

For comparison, we also report standard TWFE estimates—corresponding to equation (1) above—in the Appendix. As discussed above, TWFE is expected to understate the true effect in this setting due to the negative-weighting problem that arises in staggered designs when treatment effects are heterogeneous over time, and this is what we find (Goodman-Bacon, 2021).

6 Results

6.1 Descriptive Statistics and Pre-treatment Balance

Table 4 reports pre-treatment (pre-2006) means for key outcomes, separately for control, treated, and ‘neither’ grid cells under both the 25 km and 50 km buffer thresholds. Treated and control units are broadly comparable in baseline consumption per capita in PPP dollars (543 vs. 504 under the 25 km buffer; 536 vs. 500 under the 50 km buffer), but treated units have substantially higher average values of nighttime light intensity and population density. Average distances from treated and control cells to their nearest road segment are broadly similar across buffers—as expected given the method by which these samples were constructed.

The table also reports the two upgrades separately, and it is worth noting at the outset how different their catchments are, since this bears on the divergent effects documented in Section 6.5. Baseline consumption per capita is similar across the two (553 vs. 536 under the 25 km buffer), but the Singida–Shelui catchment is considerably brighter at baseline (nighttime light intensity of 13.4 vs. 8.0) while the Lushunga–Kagera catchment is somewhat more populous. The two also sit in markedly different physical and economic geographies: the Singida–Shelui segment runs through flatter, drier terrain roughly 600 km from Dar es Salaam, whereas Lushunga–Kagera lies in hillier, greener country more than 1,200 km from the port.

Figure 3 maps the change in log consumption per capita from 2003 to 2018 across Tanzania. Consumption rose across most of the country, with the largest gains concentrated in Dar es Salaam and along major transport networks—including the Central Corridor route. This is to be expected given the hypothesis that corridor investments of the past decades have succeeded in generating localized welfare gains, but it also

Table 4: Balance Table: Treated Units (Overall and by Project) vs. Control Units, Pre-2006

Group	N (grids)	Cons. pc (PPP\$)	Nightlights	Population	Avg dist. treated (km)	Avg dist. control (km)
<i>Panel A: 25km Buffer</i>						
Control	1,399	503.94	8.94	3,527.80	—	13.03
Treated	165	542.80	13.22	8,149.24	13.46	—
Project A	66	553.35	13.40	7,702.39	13.99	—
Project B	99	535.77	8.00	8,447.14	13.11	—
Neither	7,936	509.84	12.24	4,166.09	—	—
<i>Panel B: 50km Buffer</i>						
Control	3,100	499.57	9.59	3,099.05	—	26.55
Treated	381	535.53	12.71	7,042.20	27.54	—
Project A	166	538.31	13.04	6,031.41	28.74	—
Project B	215	533.39	12.54	7,822.61	26.62	—
Neither	6,019	513.04	12.51	4,494.43	—	—

Notes: Pre-treatment averages computed over grid-years prior to 2006. Consumption per capita in 2011 PPP dollars. Nighttime lights are harmonised values (Li et al., 2020). Population figures are in persons per grid cell. “Distance” columns report average distance (km) to the nearest treated or control road segment.


highlights the need for selecting control units near other transport routes, since these seem to have been on different growth trajectories compared to the rest of the country.

6.2 Trend Analysis

Figure 4 displays unconditional (logged) outcome trends over time, separately by treatment project and for the control group, under the 25 km buffer.²⁴ Project A (Singida–Shelui, completed 2009) and Project B (Lusahunga–Kagera, completed 2013) are shown alongside the control group. Trends for treated and control cells seem to be on broadly similar trajectories prior to each project’s completion across all three outcomes, although the difference in average values pre-2009 is considerable for nightlights and population, and the behaviour across all outcomes differs markedly in the post-project completion period across treatment groups. Note that nighttime lights are missing for many observations and periods, which can be seen, for instance, in the short time series for grids near Project B, and in the erratic movement of values in certain periods. A striking feature of the nighttime lights trends is a sharp downward shift in average values around 2014, which is more pronounced for Project A cells than for the control group. As discussed in Section 4.1, this does not reflect a real decline in luminosity but is likely just a data artefact: the underlying source transitions

²⁴Analogous figures for the 50 km buffer can be found in Appendix C.

Figure 3: Change in Log Consumption Per Capita, 2003–2018




from DMSP-OLS to recalibrated VIIRS data around 2014, and because VIIRS is far more sensitive than DMSP, many previously dark cells register small positive values after the transition. When log-transformed, these very small values are strongly negative, pulling the conditional mean downward. The greater apparent drop for Project A relative to controls may reflect heterogeneity in the share of newly positive cells across locations, and is a further reason to treat the nighttime lights results with caution.

6.3 Event Study Plots

Figures 5 and 6 present event study plots from the Callaway and Sant’Anna (2021) estimator under the 25 km buffer. Figure 5 shows results for our primary outcome, log consumption; Figure 6 shows results for the secondary outcomes of log nighttime lights and log population respectively. Each figure stacks the two

Figure 4: Outcome Trends by Treatment Project (25km Buffer)




control group specifications (main control group above, alternative control group below). Coefficients in blue (pre-treatment periods) should trend neither up nor down if the treatment and control groups were on parallel trends prior to project completion; those in red trace the evolving treatment effect in the years after the projects were completed. These coefficients are a weighted average of the effects for each project (A and B) in each year subsequent to project completion (Callaway and Sant'Anna, 2021).

For log consumption (Figure 5), pre-treatment coefficients are noisy and do deviate from zero (especially for the main control group), but there is no clear differential trend before project completion. Post-treatment coefficients are zero or slightly negative in the 4 years after project completion. After about four years following project completion, they turn positive, rising to roughly 14–15 percent relative to the control-group outcomes by Year +8, the longest horizon we observe (all coefficients are reported in Table 20 of the Appendix). The pattern suggests a fairly lengthy period of disruption and costly adaptation following project completion, after which large benefits start to materialize, and may continue to grow even after 8 years. However, section 6.5 shows that this reading does not survive separating the two projects.

For log nighttime lights (Figure 6a), effects are more muted (even negative in some specifications) but a qualitatively similar pattern emerges: no major differential pre-trends followed by a possible initial dip after project completion, with rising coefficients after 5/6 years. However, the particular pattern observed and the long-run sign is sensitive to the control group specification. For log population (Figure 6b), pre-trends are clearly absent and post-treatment coefficients indicate a steadily deepening decline, reaching approximately 9 percent (main control) to 10 percent (alternative control) below the control-group trend by Year +8, consistent with net out-migration (relative to the control group) toward urban centres made more accessible

by the upgraded road. Results for all three outcomes (consumption, nightlights and population) are very similar—though largely muted—when using the 50 km buffer, as one would expect under the assumption that the effects of the road improvements are largely localized (Figures 25 and 26 in Appendix I).

Figure 5: Event Study: Log Consumption (25km Buffer)




6.4 Regression Tables

Table 5 presents the Callaway–Sant’Anna DiD estimates for all three outcomes under the 25 km buffer. For each outcome, the first column uses the full control group and the second the alternative control group. The pooled (aggregate) ATT weights each group-time estimate by cohort size and post-treatment exposure, and corresponds to the primary treatment effect estimate. The corresponding 50 km table, and the event-time disaggregation for both buffers, are in Appendix I.

For our primary outcome, log consumption, the pooled ATT is small—2.0 percent against the full control group and 3.3 percent against the alternative control group—because the estimate averages the initially negative short-run effects with the larger positive long-run effects; the event study in Figure 5 (and Table 20

Figure 6: Event Study: Secondary Outcomes (25km Buffer)



in Appendix I showing the event-study coefficients) better captures this dynamic pattern. For nighttime lights, the pooled ATTs are small, of inconsistent sign across the two control groups, and statistically indistinguishable from zero; they also rest on only 21 of the 165 treated grid cells, since log luminosity is undefined wherever a cell records no lights. As noted above, these estimates are more susceptible to measurement bias than the consumption results. For population, pooled ATTs are negative—around 4 to 5 percent—consistent with the persistent out-migration pattern in the event studies.

Table 5: CS-DiD Results: Pooled Estimates by Outcome (25km Buffer)

	Ln(Consumption)		Ln(Night Lights)		Ln(Population)	
	(1)	(2)	(3)	(4)	(5)	(6)
ATT	0.020*** (0.008)	0.033*** (0.008)	0.021 (0.058)	-0.010 (0.066)	-0.044*** (0.011)	-0.053*** (0.013)
Observations	25024	10464	535	221	29716	12426
Control Group	Full	Alt	Full	Alt	Full	Alt
Dep Variable Mean (Control Group pre 2009)	6.217	6.239	2.188	2.200	7.117	7.173
Treated locations	165	165	21	21	165	165
Estimator	CS-DID	CS-DID	CS-DID	CS-DID	CS-DID	CS-DID

Notes: Standard errors (in parentheses) are clustered at the grid level. Within each outcome, the first column uses the full control group and the second the alternative control group. "Treated locations" counts distinct grid cells with a non-missing outcome, not observations—note that log night lights is undefined for cells recording no luminosity, which is why that outcome rests on far fewer cells. The 50 km buffer is reported in Table 19, and the event-time disaggregation in Table 20. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.5 Decomposing the Pooled Effect: Singida–Shelui and Lusahunga–Kagera

The estimates reported so far pool the two upgrades, which the staggered-adoption estimator handles correctly but which implicitly assumes that a single average effect is a meaningful summary of both. That assumption turns out to be untenable, and the pooled dynamic path is actively misleading as a description of either project.

Table 6 makes the point in its simplest form. Estimating each project separately against the same never-treated control group, and restricting both to an identical five-year post-treatment window so that the comparison is like-for-like, consumption in the Singida–Shelui catchment rises by 2.4 percent while consumption in the Lusahunga–Kagera catchment falls by 2.8 percent. Both estimates are precisely determined, and they carry opposite signs. Column (3) shows that extending Singida–Shelui to its full nine-year window—which the data permit only for that project—raises its average effect to 6.0 percent. The difference between the two projects is therefore not an artefact of the shorter observation window available for the later upgrade; it is present even when the windows are made identical.

Once the projects are separated, the delayed payoff that appears in the pooled event study disappears as a feature of either one. Singida–Shelui is marginally negative in the year of completion and positive from the following year onward, rising steadily from roughly 3 percent to 14 percent by its eighth year. Lusahunga–Kagera is negative throughout the five years we observe it, at around 3 to 4 percent below the control group. The pooled path is a weighted average of these two series while both are in the sample, and—because the Lusahunga–Kagera cohort exits after its fifth year—becomes numerically identical to the Singida–Shelui path from the sixth year onward. What looks like a five-year gestation period is in fact the arithmetic of one project with positive effects being averaged against another with negative ones, followed by the negative project dropping out of the sample. For this reason, we treat the two upgrades separately for the remainder of the paper.

Figure 7 shows the two event studies side by side, each estimated against both the main and the alternative control group. Consumption in the Singida–Shelui catchment rises steadily away from zero from the first year after completion, while the Lusahunga–Kagera coefficients sit at or below zero throughout. Pre-treatment coefficients are close to zero in all four panels, and the divergence is unaffected by which control group is used—a reassurance that it reflects the projects rather than the comparison group.

Population behaves quite differently, in that the two projects agree. Figure 8 repeats the exercise for log population and finds the same steadily deepening decline in both catchments: pre-treatment coefficients are

Table 6: Like-for-Like Comparison of the Two Upgrades (25km Buffer)

	Ln(Consumption)			Ln(Population)		
	(1)	(2)	(3)	(4)	(5)	(6)
ATT	0.024*** (0.006)	-0.028*** (0.008)	0.060*** (0.010)	-0.027*** (0.008)	-0.039** (0.018)	-0.047*** (0.013)
Observations	17580	23968	23440	21975	28462	27835
Project	Project A	Project B	Project A	Project A	Project B	Project A
Post-treatment window	e = 0 to 4	e = 0 to 4	e = 0 to 8	e = 0 to 4	e = 0 to 4	e = 0 to 8
Last year in sample	2014	2018	2018	2014	2018	2018
Treated locations	66	99	66	66	99	66
Dep Variable Mean (Control, pre-completion)	6.217	6.242	6.217	7.117	7.166	7.117
Estimator	CS-DID	CS-DID	CS-DID	CS-DID	CS-DID	CS-DID

Notes: CS-DiD estimates against the main (never-treated) control group, aggregated with equal weight across post-treatment periods. Columns (1), (2), (4) and (5) restrict both projects to the first five post-treatment years ($e = 0$ to 4) by dropping observations after the fifth year, so that the two upgrades are observed over an identical window; the Lusahunga–Kagera sample is unaffected by this restriction because the outcome series ends in 2018 regardless. Columns (3) and (6) report Singida–Shelui over its full nine-year window. Standard errors (in parentheses) are clustered at the grid level. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

flat, and the post-treatment path falls steadily in both. Singida–Shelui reaches roughly 9 percent below the control-group trend by Year +8; Lusahunga–Kagera reaches 5 percent by Year +4, which is close to where Singida–Shelui stands at that same horizon. The alternative control group gives slightly larger declines in each case. Whatever distinguishes the two upgrades, then, it is not their effect on the number of people living nearby. Both are consistent with (relative) net out-migration toward the urban centres that the improved road makes more accessible; what differs is whether the residents who remain end up better or worse off.


Project-specific estimates for nighttime lights are reported in Appendix H, rather than here in the main text. We report them for completeness, but do not emphasize the results due to concerns regarding sample size.²⁵

²⁵Recall that, of the 165 treated grid cells at the 25 km buffer, only 21 ever register positive luminosity. Split by project, that is 6 cells for Singida–Shelui and 15 for Lusahunga–Kagera, which we believe is insufficient for analysis and inference.

Figure 7: Event Study by Project: Log Consumption (25km Buffer)

(a) Singida–Shelui (completed 2009)

(b) Lusahunga–Kagera (completed 2013)




Notes: CS-DiD (Callaway and Sant’Anna, 2021) event studies for log per capita consumption, 25 km buffer, estimated separately for each upgrade. Within each panel the upper plot uses the main (never-treated) control group and the lower plot the alternative control group of cells near planned-but-uncompleted projects. Blue estimates (left of zero) test for parallel pre-trends; red estimates (right of zero) trace the treatment effect. Treatment is dated from the first full post-completion year, so the horizontal axis runs to $e = +8$ for Singida–Shelui and $e = +4$ for Lusahunga–Kagera, the latter being limited by the end of the outcome series in 2018. Vertical axis units are log-points. Standard errors clustered at the grid level. Figure 8 reports the corresponding figures for log population; the 50 km versions are produced by the same routine and available on request.

Figure 8: Event Study by Project: Log Population (25km Buffer)

(a) Singida–Shelui (completed 2009)

(b) Lusahunga–Kagera (completed 2013)



Notes: As Figure 7, for log population. Both projects share a common vertical scale, so the two panels can be compared directly. The corresponding night-lights figures are reported in Appendix H, where the very small number of treated cells with observed luminosity is discussed.

6.6 Heterogeneous Effects by Initial Economic Activity

The conceptual framework in Section 3 gives reason to expect that average effects may conceal considerable heterogeneity: whether a transport improvement generates local gains or losses may depend on how much economic activity an area starts with. Places with deeper local markets and denser input–output linkages are better placed to exploit new connectivity, while thinner ones may find that better roads mainly lower the barrier to out-migration and expose local producers to competition from better-endowed neighbours.

We test this by splitting the treated samples on *initial total consumption*—per capita consumption multiplied by grid-cell population in 2003, the last year before Project A broke ground—which aims to proxy for the size of the local economy.²⁶ Cells are assigned to a ‘high’ or ‘low’ subgroup according to whether their 2003 total consumption lies above or below the median *within their own project*, so that each upgrade contributes roughly equal numbers to both subgroups. Control units enter both subgroup regressions in full.

Table 7 reports the results at the 25 km buffer, and Figure 9 shows the corresponding event studies. For Singida–Shelui both subgroups gain: consumption rises by 8.9 percent in initially larger and richer cells and by 3.1 percent in initially smaller and poorer ones, both precisely estimated. The gradient is steep—the high subgroup gains roughly three times as much—but no group is made worse off, and the pattern is one of unequal gains rather than of winners and losers. For Lusahunga–Kagera the pattern is different in kind. The initially larger and richer cells show effects close to zero (precisely estimated), while the initially smaller and poorer cells experience a consumption decline of 6.9 percent. That project’s negative average effect is therefore not spread evenly across its catchment; it is concentrated entirely among the areas that started with the least economic activity.

That gains accrue disproportionately to initially larger and richer areas is consistent with the logic of agglomeration economies: as transport costs fall, firms and households already operating in a thicker local economy can expand more easily, because the complementary suppliers, workers and customers are already in place. This is a central prediction of the new economic geography literature (Krugman, 1991; Redding and Venables, 2004). The Lusahunga–Kagera results push the point further. Read alongside Singida–Shelui, they suggest that the competitive-pressure and out-migration mechanisms need not merely offset the direct connectivity gains—where the local economy is sufficiently thin, they can dominate them, leaving initially weaker areas absolutely worse off rather than simply gaining less than their neighbours. Faber (2014) documents the same asymmetry for peripheral regions in China, while Morten and Oliveira (2024) suggest a

²⁶A similar analysis splitting by average per capita consumption instead of average total consumption yields nearly identical results (see 13 in Section E).

similar migration channel for Brazilian roads. The policy implication tempers the headline results: corridor investment can widen spatial inequality even while raising average welfare, and complementary investment in the absorptive capacity of smaller and more remote communities may be needed for the benefits of connectivity to be broadly shared.


Table 7: Heterogeneity by Initial Total Consumption, by Project (25km Buffer)

	Ln(Consumption)		Ln(Population)	
	High	Low	High	Low
<i>Panel A: Singida–Shelui (completed 2009)</i>				
ATT	0.089*** (0.015)	0.031*** (0.012)	-0.067*** (0.010)	-0.028 (0.022)
Treated cells	33	33	33	33
<i>Panel B: Lusahunga–Kagera (completed 2013)</i>				
ATT	0.013 (0.011)	-0.069*** (0.006)	-0.037*** (0.010)	-0.041 (0.033)
Treated cells	49	50	49	50


Notes: CS-DiD simple aggregate ATTs against the full never-treated control group, 25 km buffer. Treated cells are split at the median of 2003 total consumption (per capita consumption $\times$ population) computed within each project. “Treated cells” counts distinct grid cells, not observations. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 9: Event Study: Heterogeneity by Initial Total Consumption, by Project (25km Buffer)

(a) Singida–Shelui (completed 2009)



(b) Lusahunga–Kagera (completed 2013)



Notes: CS-DiD event studies for log per capita consumption, 25 km buffer, estimated separately for each upgrade and split by whether the grid cell's 2003 total consumption (per capita consumption $\times$ population) lies above ('High', left) or below ('Low', right) the median within that project. The control group is the full control group throughout. Blue estimates (left of zero) test for parallel pre-trends; red estimates (right of zero) trace the treatment effect. Vertical axis units are log-points, on a common scale within each panel. Standard errors clustered at the grid level.

7 Robustness Checks

7.1 Robustness to Distance Threshold


Figures 10 and 11 repeat the CS-DiD event study for log consumption using the main control group across six buffer thresholds ranging from 10 km to 60 km, separately for each project. For Singida–Shelui the results reveal a clear distance gradient: effects are strongest at the narrowest thresholds, reaching 19 log-points by Year +8 at 10 km, and dissipate progressively as the buffer expands, falling to around 4 log-points at 60 km. This pattern is consistent with localised economic spillovers that attenuate with distance from the upgraded highway. For Lusahunga–Kagera the estimates remain negative across all thresholds, with no comparable gradient. Pre-treatment coefficients oscillate around zero at every buffer for both projects. Finally, the small effects at 60 km suggest that a 100 km exclusion zone for the control group is a conservative bound on the range of treatment spillovers (we take advantage of this in Section 7.3 below). The pooled version of this exercise is reported in Appendix I.

7.2 Marginal Distance Bands

The analysis in Section 7.1 uses cumulative buffers (0–10 km, 0–20 km, . . . , 0–60 km), so each successive estimate includes all units from the previous one. As a complementary check, we estimate the CS-DiD separately for each *marginal* 10 km ring, isolating the treatment effect within a specific distance band rather than pooling across all distances up to some threshold. In each ring-specific estimation, both the treated group and the control group are restricted to units that fall within that same ring from their respective highway segments, ensuring a like-for-like comparison at each distance. The full set of event-study plots is shown in Figures 22 and 23 of Appendix G.


The results confirm and sharpen the distance-attenuation pattern documented above. For Singida–Shelui, effects are positive and economically meaningful in the 0–10 km and 10–20 km rings, with smaller but still positive estimates in the 20–30 km band (Figure 22). Beyond 30 km, estimated effects turn null and remain so for the 30–40, 40–50, and 50–60 km rings. This implies that positive welfare effects are concentrated within approximately 0–30 km of the upgraded highway, and that the apparent persistence of positive effects in wider cumulative buffers partly reflects dilution by an increasingly large share of units that experience little or no effect. A similar – if less pronounced – pattern is observed for Lusahunga–Kagera, where the negative coefficients become smaller and lose their significance as one moves farther from the

Figure 10: Robustness to Distance Threshold, Singida–Shelui: CS-DiD Event Study for Log Consumption



Notes: Each panel shows the CS-DiD (Callaway and Sant'Anna, 2021) event study for log consumption for the Singida–Shelui upgrade against the main control group. The buffer radius used to define treatment and control units is labelled in each panel. Blue estimates (left of zero) test for parallel pre-trends; red estimates (right of zero) trace the evolving treatment effect. Vertical axis units are log-points. Standard errors clustered at the grid level.

Figure 11: Robustness to Distance Threshold, Lusahunga–Kagera: CS-DiD Event Study for Log Consumption



Notes: As for Figure 10, for the Lusahunga–Kagera upgrade. The post-treatment window is shorter (five years) because this project completed in 2013 and the outcome series ends in 2018.

treated highway (Figure 23). Thus, highway upgrades appear more spatially localised in their impact than the headline buffer estimates suggest.

7.3 Robustness to Spillover Exclusion Zone

Our baseline specification excludes grid cells within 50 km of any treated road segment from the control group to guard against treatment spillovers. Figure 12 repeats the CS-DiD event study for log consumption using a stricter 100 km exclusion zone, separately by project, at the 25 km treatment buffer; Figure 29 in Appendix I does the same at 50 km. The results are materially unchanged relative to the baseline for both upgrades. This is unsurprising in light of the distance-threshold analysis above, which showed that detectable effects largely dissipate beyond 30 km: tightening the exclusion zone from 50 km to 100 km therefore removes very few units that could plausibly be contaminated by treatment spillovers, leaving the estimates essentially the same.


7.4 Overlap with the Subsequent Isaka–Lusahunga Works

Our treatment definition assumes that a treated cell's post-completion outcomes reflect its own upgrade and nothing else. For the Lusahunga–Kagera project, that assumption is strained by the fact that the project shares its Lusahunga endpoint with a third corridor upgrade, between Isaka and Lusahunga, that was completed in 2020. We do not study that project as an intervention in its own right because our outcome series ends before it was finished, but it was under construction for the length of Lusahunga–Kagera's post-treatment window. If its construction depressed local economic activity in the short run—that could explain or contribute to the negative effect we estimate. The exposure is specific to this one project: 19 of Lusahunga–Kagera's 99 treated grid cells lie within 25 km of the Isaka–Lusahunga segments and 34 within 50 km, whereas no Singida–Shelui cell lies within 127 km and no control cell within 50 km.²⁷

We therefore re-estimate Lusahunga–Kagera after dropping treated cells within 25 km, and then 50 km, of those segments. The negative consumption effect attenuates by roughly forty percent—from -2.8 to -1.7 percent—and weakens in significance somewhat. The negative population effect falls from -3.9 to -1.5 percent and is no longer significant at conventional levels. Part of the loss of precision is mechanical, since the exclusions remove up to a third of the treated cells. We read these results as indicating that exposure to the subsequent works may account for some—but not all—of the estimated consumption decline

²⁷By construction, no control cell is within 50 km of *any* completed transport infrastructure project, including this one.

Figure 12: Robustness to Spillover Exclusion Zone by Project: CS-DiD Event Study for Log Consumption (25km Buffer)



Notes: CS-DiD (Callaway and Sant'Anna, 2021) event studies for log consumption at the 25 km treatment buffer, estimated separately for each upgrade. Control units within 100 km of any treated road segment are excluded (compared to 50 km in the baseline). Blue dots (pre-treatment) test for parallel pre-trends; red dots (post-treatment) trace the evolving treatment effect. Both panels share a common vertical scale, in log-points. Standard errors clustered at the grid level. The 50 km buffer is shown in Figure 29.

at Lusahunga–Kagera, and perhaps all of the estimated population decline. Appendix J sets out the overlap in detail and reports the full estimates in Table 22.

7.5 Synthetic Difference-in-Differences

As a final robustness check, we re-estimate the treatment effect for each project using the synthetic difference-in-differences (SDID) estimator of Arkhangelsky et al. (2021), as implemented by Clarke et al. (2023). SDID combines the logic of difference-in-differences with that of the synthetic control method (Abadie et al., 2010), by constructing a separate synthetic comparison for each adoption cohort (2010 for Singida–Shelui and 2014 for Lusahunga–Kagera).²⁸ Rather than weighting all control units equally, the estimator chooses non-negative unit weights so that the *weighted* average of control cells runs parallel to the treated group’s pre-treatment trajectory, together with time weights that emphasise the pre-treatment years most comparable to the post-treatment period; the treatment effect is then obtained from difference-in-differences estimated on this reweighted panel. We report standard errors computed via the placebo procedure recommended by Arkhangelsky et al. (2021).

Unlike the classic synthetic control method, SDID tolerates a constant level gap between the treatment group and its synthetic control (the differencing step absorbs it). On the other hand, SDID does not tolerate missing outcome values, requiring a strongly balanced panel over the entire study period (2003–2018). For consumption and population this restriction is innocuous—all 165 treated cells in the 25 km buffer are retained—but nighttime lights must be omitted: log luminosity is undefined in any year with zero recorded lights, which eliminates all but a handful of cells and leaves too few treated units for meaningful estimation.

Table 8 reports the estimates for the 25 km buffer, estimated separately for each upgrade. The two projects’ effects differ under SDID in the same way they do under CS-DiD, but the contrast is sharper. For Singida–Shelui the consumption ATT is 0.040 (s.e. 0.007) against the full control group and 0.058 (s.e. 0.007) against the alternative—both highly significant, and both somewhat larger than the corresponding CS-DiD estimates. For Lusahunga–Kagera the consumption ATT is -0.010 (s.e. 0.009) and -0.003 (s.e. 0.009): negative in sign but small and statistically indistinguishable from zero. The negative population effects are of a similar size for both projects under reweighting, at -3 to -5 percent, consistent with the out-migration pattern documented in Section 6, though for Singida–Shelui the estimate is significant only against the alternative control group.

Figure 13 plots, for each adoption cohort, average log consumption for the treated cells against their

²⁸As elsewhere in the paper, we date treatment from the first full post-completion year.

synthetic control, with the estimated time weights shaded along the horizontal axis. The synthetic control tracks the treated group closely throughout the pre-treatment window in both panels. The 2010 cohort pulls away from its synthetic counterpart after adoption; the 2014 cohort exhibits no such divergence, with the usual caveat that its post-treatment window is only five years long.

Thus, reweighting the control group so that it tracks each treated group's pre-treatment trajectory leaves the Singida–Shelui consumption gains intact—indeed slightly larger—while the Lusahunga–Kagera consumption decline does not survive. One implication is that part of what the CS-DiD estimator attributes to the second upgrade is a pre-existing difference in trajectory between its catchment and the unweighted control group.

Table 8: Synthetic Difference-in-Differences by Project: Log Consumption and Log Population (25km Buffer)


	Singida–Shelui (2009)				Lusahunga–Kagera (2013)			
	(1) Ln(Cons.)	(2) Ln(Cons.)	(3) Ln(Pop.)	(4) Ln(Pop.)	(5) Ln(Cons.)	(6) Ln(Cons.)	(7) Ln(Pop.)	(8) Ln(Pop.)
treatment x post	0.040*** (0.007)	0.058*** (0.007)	-0.033 (0.022)	-0.049** (0.023)	-0.010 (0.009)	-0.003 (0.009)	-0.030** (0.015)	-0.043* (0.024)
Observations	23440	8880	23440	8880	23968	9408	23968	9408
Control Group	Full	Alt	Full	Alt	Full	Alt	Full	Alt
Dep Var Mean (Control, pre-compl.)	6.217	6.239	7.144	7.198	6.242	6.260	7.197	7.252
Treated locs (balanced)	66	66	66	66	99	99	99	99
Control locs (balanced)	1399	489	1399	489	1399	489	1399	489
Inference (vce)	placebo	placebo	placebo	placebo	placebo	placebo	placebo	placebo
Estimator	SDID	SDID	SDID	SDID	SDID	SDID	SDID	SDID

Notes: Synthetic difference-in-differences estimates (Arkhangelsky et al., 2021), implemented following Clarke et al. (2023), estimated separately for each upgrade. Columns (1)–(4) are Singida–Shelui (adoption 2010, the first full post-completion year); columns (5)–(8) are Lusahunga–Kagera (adoption 2014), whose post-treatment window is five years because the outcome series ends in 2018. Within each project the first two columns report log consumption and the last two log population, and within each outcome the first column uses the full control group and the second the alternative. The sample is restricted to grid cells with a complete series over 2003–2018; the balanced treated and control counts are reported in the table. Standard errors (in parentheses) are computed by the placebo method with 100 replications. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7.6 Robustness to Violations of Parallel Trends

Our identification strategy rests on the parallel-trends assumption. While this assumption is not testable, we can ask – following Rambachan and Roth (2023) – how large would a post-treatment violation of parallel trends have to be to overturn our results? Their method replaces the point-identified event-study coefficients with robust confidence sets, which stay valid when the unobserved post-treatment violation is allowed to

Figure 13: Synthetic Difference-in-Differences: Treated vs. Synthetic Control, Log Consumption
(a) 2010 adoption cohort (Singida–Shelui) (b) 2014 adoption cohort (Lusahunga–Kagera)



Notes: Each panel plots average log consumption for treated grid cells (solid red) against their synthetic control (dashed blue)—the unit-weighted average of control cells selected by the SDID estimator (full control group, 25 km buffer). The vertical dashed line marks the adoption year, defined as the first full post-completion year. The green shading along the horizontal axis shows the estimated time weights (right-hand axis), which determine the pre-treatment years that anchor the comparison. A constant level gap between the two series is permitted by the estimator; the treatment effect is identified from the change in the gap after adoption.


be anything inside a restriction class Δ whose size we increase gradually. We focus on the case of *relative magnitudes* (Δ^{RM}), under which the change in the differential trend from one post-treatment year to the next may be at most $\bar{M}$ times the largest such change observed between any two consecutive years in the pre-treatment period. The main outcome we report is the *breakdown value*: the smallest $\bar{M}$ at which the robust confidence set first contains zero. All confidence sets are 90 percent.

One feature of Δ^{RM} is particularly noteworthy because it determines how the breakdown values should be interpreted. The bound applies to the change in the differential trend from one year to the next — *and is allowed to go the same way every year*. The permitted bias therefore accumulates over time: at horizon e the differential trend may have drifted as far as $e + 1$ times the annual bound. At $e = +8$ (when some of our effects are largest) that is nine years of compounding. A given $\bar{M}$ consequently means something very different at a long horizon versus at a short one.

With that in mind, we apply this method to the CS-DiD event study for log consumption at the 25 km buffer, estimating each project separately against the never-treated control group and dating treatment from the first full post-completion year, as elsewhere in the paper. For each project we report its longest observable horizon: $e = +8$ for Singida–Shelui, and $e = +4$ for Lusahunga–Kagera, whose 2014 cohort runs out of outcome data after five years.

Panel (a) of Figure 14 shows results for the Singida–Shelui project. Under exact parallel trends the effect 8 years after project completion is +14.4 log points, with a 90 percent confidence interval of [11.2, 17.6]. It

Figure 14: Sensitivity to Parallel-Trends Violations: Log Consumption at Each Project's Longest Observed Horizon (Relative Magnitudes)



Notes: Each panel shows 90% robust confidence sets (Rambachan and Roth, 2023) for the CS-DiD log consumption effect at that project's longest observable horizon (25 km buffer, main control group). Each project is estimated separately against the never-treated control group. "Original" is the conventional confidence interval under exact parallel trends.

remains significantly positive at $\bar{M} = 0.25$ and loses significance at $\bar{M} = 0.50$. Read against the conventional benchmark of $\bar{M} = 1$, a breakdown value of 0.50 looks weak, so it is worth spelling out what it actually permits. Table 9 does this. The largest year-on-year change anywhere in the Singida–Shelui pre-treatment path is 3.84 log points. At $\bar{M} = 0.50$ the differential trend may therefore move 1.92 log points in a year—and, because the bound is on the annual change, may keep moving that way for all 9 post-treatment years, reaching 17.28 log points by $e = +8$.

However, the pre-treatment data look nothing like that path. In a typical pre-treatment year the differential trend moved 2.05 log points, but those movements largely cancelled out: the trend reversed direction 3 times in 5 years, never went the same way for more than 2 consecutive years, and finished only 4.42 log points from where it started—a net drift of 0.88 log points a year. The violation required to overturn the result is a drift of 1.92 log points a year that never once reverses – which we think is unlikely given the shape of the pre-period dynamics summarized in Panel A of Table 9.

Panel (b) shows results for the Lusahunga–Kagera project, and here the conclusion is different. The negative effect at $e = +4$ is -2.9 log points and breaks down at $\bar{M} = 0.25$ —half the Singida–Shelui value, and at half the horizon. As reported in Table 9, $\bar{M} = 0.25$ permits a drift of only 1.18 log points a year, or 5.90 log points over the five post-treatment years. That is a smaller violation, but is enough to erase the negative result.

Taken together with the other robustness exercises above (i.e., the negative Lusahunga–Kagera effect attenuates by roughly forty percent when cells near the Isaka–Lusahunga works are excluded, and loses

Table 9: What the Breakdown Values Permit, in Log Points

	Singida–Shelui ($e = +8$)	Lusahunga–Kagera ($e = +4$)
<i>A. Differential trend observed before treatment</i>		
Largest year-on-year change ($\bar{M}$ scaling factor)	3.84	4.72
Average year-on-year change	2.05	2.61
Longest run in a single direction (years)	2	3
Net drift per year	0.88	0.20
Total divergence over the 5 pre-treatment years	4.42	1.02
<i>B. Violation permitted at the breakdown value</i>		
Breakdown $\bar{M}$	0.50	0.25
Implied change per year, sustained in one direction	1.92	1.18
Implied total divergence by the final horizon	17.28	5.90

Notes: All entries except the breakdown $\bar{M}$ and the run length are in log points, that is $100\times$ the change in log consumption. Panel A describes the estimated pre-treatment differential trend between treated and control cells over the five years before treatment, taken from the CS-DiD event study for log consumption at the 25 km buffer. “Largest year-on-year change” is the maximum absolute change between consecutive pre-treatment event-study coefficients, with the omitted reference period entering as a zero; this is the quantity that $\bar{M}$ multiplies under Δ^{RM} . “Net drift per year” is the total divergence over the pre-treatment window divided by its length, and differs from the average year-on-year change because the year-on-year movements partly offset one another. Panel B converts the breakdown values in Figure 14 into the same units. Because Δ^{RM} bounds the annual change rather than the level, the implied total divergence is the annual figure multiplied by the number of post-treatment years: nine for Singida–Shelui and five for Lusahunga–Kagera.

significance entirely under synthetic difference-in-differences reweighting), we find that three independent checks point the same way. Our evidence for a positive effect at Singida–Shelui is considerably stronger and more robust than our evidence for a negative effect at Lusahunga–Kagera. It is safe to conclude that the two projects’ effects differed substantially, but it is less certain that the second one reduced welfare.

8 Conclusion

This paper studies the local economic effects of upgrading two segments of Tanzania’s Central Corridor—the Singida–Shelui section (completed 2009, US\$56 million) and the Lusahunga–Kagera section (completed 2013, US\$120 million)—using a panel of 10 km $\times$ 10 km grid cells from 2003 to 2018. Due to the staggered timing of treatments, we implement the Callaway and Sant’Anna (2021) staggered difference-in-differences estimator, comparing grid cells near these two segments with grid cells near other highway segments off the Central Corridor that did not experience road upgrading.

Our main results, using a 25 km treatment buffer, indicate that the two upgrades had markedly different

local effects, and that the pooled estimates are a poor description of either. Taken together, the projects appear to generate gains only after a substantial lag—benefits seem absent for roughly five years before rising to around 14 percent above the control-group trend by Year +8. Estimated separately, that lag turns out to be an artefact of aggregation rather than a property of either project. The Singida–Shelui upgrade generates positive effects from its second year onward, rising steadily to about 14 percent by Year +8, with no sustained initial dip. The Lusahunga–Kagera upgrade is negative throughout the five years for which we can observe it, at around 3 to 4 percent below its control group. Population density falls persistently near treated roads for both projects, consistent with net out-migration toward more accessible urban centres, though for Lusahunga–Kagera this decline does not survive the exclusion of cells near the subsequent Isaka–Lusahunga works. Night-light effects are smaller and less consistent across control groups.

A further finding concerns the distribution of benefits across space. Splitting the treated sample by initial total consumption—a proxy for local economic scale—reveals that areas initially larger and richer capture substantially greater welfare gains and avoid potential welfare losses (in the case of Lusahunga–Kagera). This pattern is consistent with agglomeration economies: thicker local markets are better positioned to exploit new connectivity, while smaller areas may struggle to contend with competitive pressure from better-endowed neighbors and out-migration (Krugman, 1991; Redding and Venables, 2004; Faber, 2014). The implication is that road investment along the Central Corridor may widen spatial inequality even as it raises average welfare in treated areas—underscoring the importance of complementary investments in human capital and local market infrastructure to ensure the benefits of connectivity reach smaller and more remote communities.

These patterns carry several important implications. The most immediate is methodological: evaluating a corridor as a programme rather than as a set of individual segments can produce a dynamic profile that describes none of its components. Here, pooling generated an apparent five-year gestation period that neither upgrade exhibits, and a positive average that conceals one project with sizeable gains and another with losses. Since corridor investments are routinely appraised and evaluated at programme level, this is a caution with some generality. It also matters for cost-benefit analysis: an appraisal anchored to the pooled path would understate the returns to Singida–Shelui while overstating those to Lusahunga–Kagera. At the same time, the persistent population decline near treated roads suggests that the primary channel is improved connectivity to urban markets and services rather than in-situ job creation; rural communities may benefit through better access rather than through attracting firms or workers. Finally, the results may suggest potential differences in total economic activity (proxied by nightlights) versus welfare/average consumption. Roads may have raised average consumption in local areas without substantially increasing

total economic activity.²⁹

Several caveats apply. First, our strategy is designed to identify local economic effects and will not pick up general equilibrium effects (e.g., changes to national prices, wages or price volatility) or direct effects on locations farther away on the central corridor (e.g., Dar es Salaam). Second, the satellite-derived consumption and night light measures are not necessarily perfect proxies of economic activity (especially night lights), so that some effects could be underestimated or overestimated. Third, the study period ends in 2018—which might be before the full dynamic effects have had time to stabilise. Fourth, with only two treatment projects, it is difficult to rule out the possibility that our results reflect idiosyncratic time-varying features of these corridor segments. Fifth, potential spillovers to control cells may exist—if improved connectivity propagates beyond the 25/50 km treatment zone. This would tend to attenuate our estimates, suggesting that the true effects might be even larger.

Sixth, we date treatment from completion and do not model the construction period. However, construction is itself a local economic event—workers are hired, spending may increase, and movement may be disrupted along the road being built. Both of our pre-treatment windows contain construction activity, and the Kagoma–Lusahunga works were unusually protracted.³⁰ If activity near the road was temporarily elevated while construction was underway and then fell back once it finished, our pre-treatment baseline would be too high and the estimated post-completion effect correspondingly too negative. That is one candidate explanation for the persistent decline we estimate for Lusahunga–Kagera, and it is consistent with that project’s estimates being the ones that fail to survive the synthetic difference-in-differences reweighting. We cannot separate a construction effect from a completion effect with annual grid-cell data, so we flag it here as a limitation rather than attempting to correct for it.

Finally, to gauge the scale of the welfare effects relative to project expenditures, Appendix F reports a back-of-the-envelope cost-benefit analysis of the two upgrades, using each project’s own event-study path to estimate benefits. Constructing a year-by-year stream of local welfare effects for populations within the 25 km buffer and comparing it against real capital and maintenance costs over a 25-year road lifetime, Singida–Shelui returns a benefit-cost ratio of approximately 5.9 at a standard 8% real discount rate, while Lusahunga–Kagera returns a negative ratio, since its estimated consumption effects are themselves negative. The contrast is stark, and it is worth repeating that this analysis inherits every caveat attaching to the underlying estimates—in particular that Lusahunga–Kagera’s negative effects are not very robust (they

²⁹Assuming that nightlights are a good enough proxy for total economic activity in this context.

³⁰An initial contract ran from April 2006 until its termination for contractor non-performance in February 2008, and the balance of the works was carried out between July 2009 and completion in December 2013.

weaken substantially once cells near the subsequent Isaka–Lusahunga works are excluded, and do not survive the synthetic difference-in-differences reweighting or Rambachan and Roth (2023) robustness tests at all). The Singida–Shelui figure should also be read with care: because its estimated effect is still rising when our data end, holding it flat thereafter extrapolates the largest coefficient in the series across a further seventeen years, which may overstate or understate the long-run effects (depending on whether effects continue on their trend or decrease over time). Sensitivity of the BCR to the discount rate and to sampling uncertainty in the treatment-effect estimates is examined in Table 17.

From a policy perspective, our findings – on balance – support the case for corridor infrastructure investment in low-income settings, but argue for evaluation frameworks that track outcomes over at least a decade. They also point to the value of complementary investments: the relative population decline near improved roads suggests that rural residents are responding primarily by migrating rather than diversifying into non-farm activity, which implies that road upgrades alone may be insufficient to foster broad-based rural economic development. Future work should examine whether the consumption gains are concentrated among specific groups (of households or businesses), whether they persist beyond the study window, and whether complementary investments in electricity, digital connectivity, or agricultural extension services amplify the returns to road upgrades.

References

- Abadie, Alberto, Alexis Diamond, and Jens Hainmueller**, “Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program,” *Journal of the American Statistical Association*, 2010, 105 (490), 493–505.
- Abbasi, Mansoureh, Mathilde Lebrand, Arcady Mongoue, Roland Pongou, and Fan Zhang**, “Roads, Electricity, and Jobs: Evidence of Infrastructure Complementarity in Sub-Saharan Africa,” Policy Research Working Paper 9976, World Bank 2022.
- Abeberese, Ama Baafra and Mary Chen**, “Intranational Trade Costs, Product Scope and Productivity: Evidence from India’s Golden Quadrilateral Project,” *Journal of Development Economics*, 2022, 156, 102791.

- Adukia, Anjali, Sam Asher, and Paul Novosad**, “Educational Investment Responses to Economic Opportunity: Evidence from Indian Road Construction,” *American Economic Journal: Applied Economics*, 2020, 12 (1), 348–376.
- Aggarwal, Shilpa**, “Do Rural Roads Create Pathways out of Poverty? Evidence from India,” *Journal of Development Economics*, 2018, 133, 375–395.
- , **Brian Giera, Dahyeon Jeong, Jonathan Robinson, and Alan Spearot**, “Market Access, Trade Costs, and Technology Adoption: Evidence from Northern Tanzania,” *Review of Economics and Statistics*, 2024, 106 (6), 1511–1528.
- Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager**, “Synthetic Difference-in-Differences,” *American Economic Review*, 2021, 111 (12), 4088–4118.
- Asher, Sam and Paul Novosad**, “Rural Roads and Local Economic Development,” *American Economic Review*, 2020, 110 (3), 797–823.
- Beuran, Monica, Marie Gachassin, and Gaël Raballand**, “Are There Myths on Road Impact and Transport in Sub-Saharan Africa?,” *Development Policy Review*, 2015, 33 (5), 673–700.
- Boo, Gianluca, Edith Darin, Douglas R. Leasure, Claire A. Dooley, Heather R. Chamberlain, Attila N. Lázár, Kevin Tschirhart, Cyrus Sinai, Nicole A. Hoff, Trevon Fuller, Kamy Musene, Arly Batumbo, Anne W. Rimoin, and Andrew J. Tatem**, “High-resolution population estimation using household survey data and building footprints,” *Nature Communications*, 2022, 13, 1330.
- Brooks, Wyatt and Kevin Donovan**, “Eliminating Uncertainty in Market Access: The Impact of New Bridges in Rural Nicaragua,” *Econometrica*, 2020, 88 (5), 1965–1997.
- , **Joseph P. Kaboski, Illenin O. Kondo, Yao Amber Li, and Wei Qian**, “Infrastructure Investment and Labor Monopsony Power,” *IMF Economic Review*, 2021, 69 (3), 470–504.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Clarke, Damian, Daniel Pailaňir, Susan Athey, and Guido Imbens**, “Synthetic Difference-in-Differences Estimation,” 2023. arXiv preprint arXiv:2301.11859.

- Coşar, A. Kerem, Banu Demir, Devaki Ghose, and Nathaniel Young**, "Road capacity, domestic trade and regional outcomes," *Journal of Economic Geography*, 2022, 22 (5), 901–929.
- Datta, Saugato**, "The Impact of Improved Highways on Indian Firms," *Journal of Development Economics*, 2012, 99 (1), 46–57.
- Donaldson, Dave**, "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, 2018, 108 (4–5), 899–934.
- Dorosh, Paul, Hyoung-Gun Wang, Liang You, and Emily Schmidt**, "Crop Production and Road Connectivity in Sub-Saharan Africa: A Spatial Analysis," Policy Research Working Paper 5385, World Bank 2010.
- Duranton, Gilles, Peter M. Morrow, and Matthew A. Turner**, "Roads and Trade: Evidence from the US," *Review of Economic Studies*, 2014, 81 (2), 681–724.
- Eynde, Oliver Vanden and Liam Wren-Lewis**, "Complementarities in Infrastructure: Evidence from Indian Agriculture," *World Bank Economic Review*, 2025. Forthcoming.
- Faber, Benjamin**, "Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System," *Review of Economic Studies*, 2014, 81 (3), 1046–1070.
- Foster, Vivien and Cecilia Briceño-Garmendia**, *Africa's Infrastructure: A Time for Transformation*, Washington, DC: Agence Française de Développement and World Bank, 2010.
- , **Nisan Gorgulu, Dhruv Jain, Stéphane Straub, and Maria Vagliasindi**, "The Impact of Infrastructure on Development Outcomes A Meta-Analysis," Technical Report 2023.
- Gebresilasse, Mesay**, "Rural roads, agricultural extension, and productivity," *Journal of Development Economics*, 2023, 162.
- Gertler, Paul J., Marco Gonzalez-Navarro, Tadeja Gračner, and Alexander D. Rothenberg**, "Road maintenance and local economic development: Evidence from Indonesia's highways," *Journal of Urban Economics*, 2024, 143.
- Ghani, Ejaz, Arti Grover Goswami, and William R. Kerr**, "Highway to Success: The Impact of the Golden Quadrilateral Project for the Location and Performance of Indian Manufacturing," *Economic Journal*, 2016, 126 (591), 317–357.

- Goodman-Bacon, Andrew**, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, 225 (2), 254–277.
- Huber, John D. and Laura Mayoral**, "Economic Development in Pixels: The Limitations of Nightlights and New Spatially Disaggregated Measures of Consumption and Poverty," Working Paper 1433, Barcelona School of Economics 2024.
- Jedwab, Rémi and Adam Storeygard**, "The Average and Heterogeneous Effects of Transportation Investments: Evidence from Sub-Saharan Africa 1960-2010," *Journal of the European Economic Association*, 2022, 20 (1), 1–38.
- Jedwab, Remi and Alexander Moradi**, "The permanent effects of transportation revolutions in poor countries: Evidence from Africa," *Review of Economics and Statistics*, 2016, 98 (2), 268–284.
- Kahn-Lang, Ariella and Kevin Lang**, "The Promise and Pitfalls of Differences-in-Differences: Reflections on 16 and Pregnant and Other Applications," *Journal of Business and Economic Statistics*, 2020, 38 (3), 613–620.
- Krugman, Paul**, "Increasing Returns and Economic Geography," *Journal of Political Economy*, 1991, 99 (3), 483–499.
- Li, Xuecao, Yuyu Zhou, Min Zhao, and Xia Zhao**, "A harmonized global nighttime light dataset 1992–2018," *Scientific Data*, 2020, 7 (1).
- , —, —, —, and —, "Harmonization of DMSP and VIIRS Nighttime Light Data from 1992–2024 at the Global Scale (Version 10)," Figshare 2024.
- Lloyd, Christopher T, Heather Chamberlain, David Kerr, Greg Yetman, Linda Pistolesi, Forrest R Stevens, Andrea E Gaughan, Jeremiah J Nieves, Graeme Hornby, Kytt Macmanus, Parmanand Sinha, Maksym Bondarenko, Alessandro Sorichetta, and Andrew J Tatem**, "Global spatio-temporally harmonised datasets for producing high-resolution gridded population distribution datasets," *Big Earth Data*, 2019, 3 (2), 108–139.
- Martincus, Christian Volpe, Jerónimo Carballo, and Ana Cusolito**, "Roads, Exports and Employment: Evidence from a Developing Country," *Journal of Development Economics*, 2017, 125, 21–39.

- Melecky, Martin, Mark Roberts, and Siddharth Sharma**, "The wider economic benefits of transport corridors: A policy framework and illustrative application to the China-Pakistan Economic Corridor," *Cambridge Journal of Regions, Economy and Society*, 2019, 12 (1), 17–44.
- , **Siddharth Sharma, and Hari Subhash**, "Wider Economic Benefits of Investments in Transport Corridors and the Role of Complementary Policies The WEB of Transport Corridors in South Asia," Policy Research Working Paper, World Bank 2018.
- Morten, Melanie and Jaqueline Oliveira**, "The Effects of Roads on Trade and Migration: Evidence from a Planned Capital City," *American Economic Journal: Applied Economics*, 2024, 16 (2), 389–421.
- Okonge, Flory Okandju and George Nyaronga**, "Long Term Impact of Information and Communication Technology on Corridors Performance," 2025, 8 (4), 2025–81.
- Quium, A. S. M.**, "Transport Corridors for Wider Socio-economic Development," *Sustainability (Switzerland)*, 2019, 11 (19).
- Rambachan, Ashesh and Jonathan Roth**, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, 90 (5), 2555–2591.
- Redding, Stephen J. and Anthony J. Venables**, "Economic Geography and International Inequality," *Journal of International Economics*, 2004, 62 (1), 53–82.
- **and Matthew A. Turner**, "Transportation Costs and the Spatial Organization of Economic Activity," in Gilles Duranton, J. Vernon Henderson, and William C. Strange, eds., *Handbook of Regional and Urban Economics*, Vol. 5, Elsevier, 2015, pp. 1339–1398.
- Shamdasani, Yogita**, "Rural road infrastructure & agricultural production: Evidence from India," *Journal of Development Economics*, 2021, 152.
- Storeygard, Adam**, "Farther on Down the Road: Transport Costs, Trade and Urban Growth in Sub-Saharan Africa," *Review of Economic Studies*, 2016, 83 (3), 1263–1295.
- Sun, Liyang and Sarah Abraham**, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, 225 (2), 175–199.
- The UK Department for Transport**, "Understanding and Valuing Impacts of Transport Investment: Updating Wider Economic Impacts Guidance," *Department for Transport, Great Minster House*, 2016.

World Bank, “Project Appraisal Document: Central Corridor Transport Project,” Technical Report, World Bank, Washington, DC 2004.

— , “Discounting Costs and Benefits in Economic Analysis of World Bank Projects,” Operations Policy and Sector Quality (OPSPQ) Note, World Bank Group, Washington, DC 2023.

Appendix

A Additional Maps

Figure 15: All Large Global Construction Projects in Tanzania




Figure 16: High-Value (>USD 100m) GC Projects on the CC Route

High Value Infrastructure Projects Within 25km of Central Corridor




Figure 17: (Log) Consumption Per Capita, 2003

Logged Consumption By Location (2003)





Figure 18: (Log) Consumption Per Capita, 2018


Logged Consumption By Location (2018)



B Trend Graphs with Alternative Control Group

The figures below replicate the trend graphs from Section 6 using the alternative control group definition (grid cells near planned-but-never-completed projects). These are provided as a robustness check on the parallel-trends patterns observed in the main analysis.

Figure 19: Outcome Trends by Treatment Project, Alternative Control Group (25km Buffer)



C 50km Buffer Results

Trend Graphs (50km Buffer)

Figure 20: Outcome Trends by Treatment Project (50km Buffer)



D Basic DiD (TWFE) Results

Table 10: Basic DiD Results (25km Buffer)

	(1) consump_rf	(2) consump_rf	(3) ln_gdp	(4) ln_gdp
treated	-0.008 (0.007)	0.007 (0.007)	-0.082*** (0.018)	-0.105*** (0.019)
Constant	6.296*** (0.000)	6.309*** (0.001)	9.553*** (0.001)	9.903*** (0.002)
Observations	27392	12384	35301	16128
Control Group	full	alt	full	alt
Dep Variable Mean (Control Group pre 2009)	6.217	6.239	9.064	9.106

Notes: Standard errors (in parentheses) are clustered at the grid level. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Basic DiD Results (50km Buffer)

	(1) consump_rf	(2) consump_rf	(3) ln_gdp	(4) ln_gdp
treated	-0.013*** (0.004)	0.001 (0.004)	-0.088*** (0.012)	-0.096*** (0.013)
Constant	6.279*** (0.000)	6.284*** (0.000)	9.335*** (0.000)	9.696*** (0.001)
Observations	60272	28560	77721	37191
Control Group	full	alt	full	alt
Dep Variable Mean (Control Group pre 2009)	6.206	6.215	8.834	8.936

Notes: Standard errors (in parentheses) are clustered at the grid level. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

E Heterogeneous Treatment Effects by Initial Economic Activity

The tables in this section of the Appendix supplement the results in Section 6.6. First, table 12 repeats the by-project split of Table 7 at the 50 km buffer. Then, table 13 reports results when splitting the treatment groups by average consumption per capita (i.e. above/below the median), rather than total consumption — to see whether effects are heterogeneous by total economic activity or average income.

Table 12: Heterogeneity by Initial Total Consumption, by Project (50km Buffer)

	Ln(Consumption)		Ln(Population)	
	High	Low	High	Low
<i>Panel A: Singida–Shelui (completed 2009)</i>				
ATT	0.053*** (0.009)	-0.007 (0.005)	-0.087*** (0.007)	-0.040*** (0.014)
Treated cells	83	83	83	83
<i>Panel B: Lushunga–Kagera (completed 2013)</i>				
ATT	0.009 (0.008)	-0.050*** (0.005)	-0.052*** (0.008)	0.011 (0.025)
Treated cells	107	108	107	108

Notes: As Table 7 but using the 50 km buffer for treatment assignment. Standard errors (in parentheses) are clustered at the grid level. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 13: Heterogeneity by Initial Average Consumption Per Capita, by Project (25km Buffer)

	Ln(Consumption)		Ln(Population)	
	High	Low	High	Low
<i>Panel A: Singida–Shelui (completed 2009)</i>				
ATT	0.076*** (0.015)	0.044*** (0.013)	-0.058*** (0.011)	-0.036 (0.022)
Treated cells	33	33	33	33
<i>Panel B: Lushunga–Kagera (completed 2013)</i>				
ATT	-0.006 (0.012)	-0.050*** (0.008)	-0.041** (0.021)	-0.038 (0.028)
Treated cells	49	50	49	50

Notes: As Table 7 but using average consumption per capita to split the treatment group. Standard errors (in parentheses) are clustered at the grid level. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Tables 14 and 15 report the *pooled* CS-DiD ATT estimates (i.e. both projects together) for log per capita consumption, separately for grid cells above and below the within-cohort median of initial total consumption (per capita consumption $\times$ population in 2003), for the 25 km and 50 km buffers respectively. See Section 6.6 for full discussion.

Table 14: CS-DiD Heterogeneity by Initial Total Consumption (25km Buffer)

	(1) High Initial total cons.	(2) Low Initial total cons.
ATT	0.055*** (0.011)	-0.015 (0.009)
Observations	23696	23712
Initial total cons. group	High	Low
Control Group	Full (unmatched)	Full (unmatched)
Estimator	CS-DID	CS-DID

Notes: Each column reports the CS-DiD pooled ATT for log per capita consumption estimated on the subsample of treated units whose 2003 total consumption (per capita consumption $\times$ population) lies above ('High') or below ('Low') the median within their treatment cohort (Project A or Project B). The control group is the full control group in both columns. Standard errors (in parentheses) are clustered at the grid level. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 15: CS-DiD Heterogeneity by Initial Total Consumption (50km Buffer)

	(1) High Initial total cons.	(2) Low Initial total cons.
ATT	0.034*** (0.007)	-0.025*** (0.004)
Observations	52640	52656
Initial total cons. group	High	Low
Control Group	Full (unmatched)	Full (unmatched)
Estimator	CS-DID	CS-DID

Notes: As Table 14 but using the 50 km buffer for treatment assignment. Standard errors (in parentheses) are clustered at the grid level. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F Cost-Benefit Analysis

This appendix translates the causal estimates from the main text into a back-of-the-envelope cost-benefit analysis (CBA) for the two road upgrades studied in this paper—the Singida–Shelui section (Project A, completed 2009) and the Lusahunga–Kagera section (Project B, completed 2013). The goal is to assess whether the local welfare gains identified by the CS-DiD estimator are large enough, relative to project costs, to clear a standard cost-benefit threshold.

Methodology

Costs. Capital costs are the road-specific expenditures from official project records. For Project A, the World Bank ICR for project P078387 reports US\$55.73 million allocated to the Singida–Shelui road component (the total WB project credit of US\$122 million also covered Zanzibar coastal roads, ferry rehabilitation, and railway support; see ICR_P078387.pdf). For Project B, the Global Construction database (ID 258559) reports a total cost of US\$120 million. Both figures are in nominal USD at the time of construction and are deflated to real 2011 USD using US CPI ratios from the U.S. Bureau of Labor Statistics (Project A: factor 1.049; Project B: factor 0.965) to match the price base of the Huber–Mayoral consumption estimates. Annual maintenance is assumed to be 1.5% of real capital cost per year, consistent with standard practice for sealed roads in Sub-Saharan Africa. All costs are discounted over a 25-year road lifetime from the project completion year.

Benefits. For each project, benefits are computed as the annual welfare gain accruing to the local population within the 25 km treatment buffer. Specifically, for each year t post-treatment:

$$B_t = N \times \bar{c}_0 \times (\exp(\widehat{ATT}_t) - 1), \quad (3)$$

where N is the total modelled population in treatment grid cells (fixed at the pre-treatment baseline), $\bar{c}_0$ is the population-weighted mean per capita real consumption in the pre-treatment period, and $\widehat{ATT}_t$ is the event-study coefficient for year t from Column (1) of Table 20 (full control group, 25 km buffer). For years beyond Year +9—the last year observed in our study window—the treatment effect is assumed to plateau at the Year +9 estimate, a conservative extrapolation given the upward trajectory visible in the event-study plots. The stream of annual benefits is discounted at a real rate of r over 25 years.

Net present value and BCR. The net present value (NPV) and benefit-cost ratio (BCR) are:

$$\text{NPV} = \sum_{t=1}^{25} \frac{B_t}{(1+r)^t} - C, \quad \text{BCR} = \frac{\sum_{t=1}^{25} \frac{B_t}{(1+r)^t}}{C}, \quad (4)$$

where C denotes the present value of total costs (capital plus discounted maintenance). The internal rate of return (IRR) is the value of r at which $\text{NPV} = 0$.

Results

Table 16 presents the main CBA results at a baseline discount rate of 8%. Following World Bank guidance, the appropriate social discount rate for a developing country is linked to its per capita growth prospects via the Ramsey formula, with rates of 6–10 percent recommended for countries with per capita growth in the range of 3–5 percent per year (World Bank, 2023). Tanzania sustained per capita GDP growth of approximately 3–5 percent annually over the study period, placing an 8 percent real discount rate comfortably within the recommended range. Sensitivity to the discount rate choice is examined in Table 17.

The two projects give sharply different answers. Singida–Shelui yields a benefit-cost ratio of about 5.9 and a net present value of roughly \$335 million at an 8% discount rate, with an internal rate of return above 30%. Lusahunga–Kagera yields a negative benefit-cost ratio: because its estimated consumption effects are negative throughout the observed window, the “benefit” stream is a stream of losses, and the internal rate of return is undefined. We would caution strongly against reading the second figure as a settled verdict on that investment. It rests on effects that attenuate by roughly forty percent once cells near the subsequent Isaka–Lusahunga works are excluded, and that lose statistical significance entirely under synthetic difference-in-differences; it also covers only five post-completion years, whereas Singida–Shelui’s gains did not reach their full size until considerably later. The honest reading is that Singida–Shelui clears a conventional cost-benefit threshold comfortably, while the evidence for Lusahunga–Kagera is inconclusive over the horizon we can observe.

Figure 21 plots the cumulative discounted benefits and costs over the road lifetime for each project. For Singida–Shelui the cumulative benefit line crosses cumulative costs within the first few years and diverges upward thereafter; for Lusahunga–Kagera it moves away from the cost line in the wrong direction. Note also that the Singida–Shelui calculation holds its effect flat beyond Year +8, which extrapolates the largest estimated coefficient across the remaining seventeen years of assumed road life—an assumption that flatters

the result rather than a conservative one.

Table 16: Cost-benefit analysis of Central Corridor road upgrades (Projects A and B)

Project	Completion year	Capital cost (real \$M)	Beneficiary pop. (M)	Consumption p.c. (\$)	PV benefits (\$M)	PV costs (\$M)	NPV (\$M)	BCR	IRR (%)
A	2009	58	0.57	620	235	68	167	3.5	18.1
B	2013	116	1.04	605	419	134	284	3.1	17.2

Notes: 25 km buffer panel. 8% real discount rate. 25-year road life. Annual maintenance = 1.5% of real capital cost. Benefits derived from CS-DiD event-study coefficients (Column 1, Table ??); plateau assumed at Year +9 for years outside the study window. Capital costs deflated from nominal USD to real 2011 USD. Project A cost is the road component of World Bank project P078387; Project B cost from Global Construction database (ID 258559).

Table 17 explores the sensitivity of the BCR to two key parameters: the discount rate (5%, 8%, 12%) and the size of the treatment effect (central estimate, lower 90% CI bound, upper 90% CI bound). Higher discount rates penalise the delayed benefit stream more heavily, so the BCR is considerably more sensitive to r than to sampling uncertainty in the ATT coefficients.

Table 17: Sensitivity of the benefit-cost ratio to discount rate and treatment-effect size


Project	Discount rate	BCR (lower)	BCR (central)	BCR (upper)
A	5%	3.10	5.10	7.20
A	8%	2.00	3.50	5.00
A	12%	1.10	2.10	3.10
B	5%	2.80	4.60	6.40
B	8%	1.80	3.10	4.50
B	12%	1.00	1.90	2.80

Notes: "Lower" / "upper" use the 5th and 95th percentile bounds of the CS-DiD event-study ATT coefficients ($\pm 1.645 \times \text{SE}$). All other parameters as in Table 16.

Interpretation and Caveats

Several caveats govern the interpretation of these figures. First, and most importantly, equation (3) captures only the *local* welfare gains to populations within the 25 km buffer—it deliberately excludes the general-equilibrium price and wage effects, trade-flow improvements, and wider network externalities that the corridor literature typically finds to be large. The BCRs in Table 16 should therefore be read as conservative lower bounds on the social return. Second, the plateau extrapolation beyond Year +9 may under- or over-state the long-run benefit depending on whether the consumption gains continue to compound or eventually diminish; this is an inherent limitation of any study whose panel does not extend to the end of the road's design life. Third, the consumption measure is derived from satellite-based estimates (Huber and Mayoral, 2024) that

Figure 21: Cumulative Discounted Benefits vs. Costs over the Road Lifetime




Note: 8% real discount rate. Benefits computed from the CS-DiD event-study ATT estimates (Column 1, Table 20), plateauing at the Year +9 estimate beyond the study window. Costs include real capital expenditure (nominal costs deflated to 2011 USD) plus discounted annual maintenance at 1.5% of capital per year. The crossing point of the two curves indicates the payback year.

may introduce measurement error, potentially biasing the benefit estimates in either direction. Fourth, the $\bar{c}_0$ baseline is computed from grid-cell averages and so captures both farm and non-farm households within the buffer; a disaggregated analysis might find heterogeneous returns across livelihood types. Taken together, these considerations suggest that the CBA figures are most informative as a rough order-of-magnitude check rather than a precise welfare accounting exercise.

G Marginal Distance Band Analysis


Figures 22 and 23 presents CS-DiD event studies for log per capita consumption estimated separately for each 10 km distance ring between 0 and 60 km from the highway for each project. Unlike the cumulative buffer robustness check in Section 7.1, each panel here isolates a single marginal band—0–10, 10–20, 20–30, 30–40, 40–50, and 50–60 km. In each estimation, the control group is also restricted to units within the same distance ring from their nearest non-upgraded highway segment, so that treated and control units are compared at comparable distances from a major road. See Section 7.2 for discussion.

Figure 22: Marginal Distance Bands, Singida-Shelui: CS-DiD Event Study for Log Consumption (Band-Matched Control)



Notes: Each panel shows the CS-DiD (Callaway and Sant'Anna, 2021) event study for log per capita consumption for the Singida-Shelui upgrade using the main control group, restricted to units within the labelled 10 km distance ring. The treated group contains grid cells whose distance to the nearest upgraded highway segment falls within the stated band; the control group is restricted to cells at the same distance from their nearest non-upgraded segment. Pre-treatment estimates (left of zero) test for parallel pre-trends; post-treatment estimates (right of zero) trace the evolving treatment effect. Vertical axis units are log-points. Standard errors clustered at the grid level.

Figure 23: Marginal Distance Bands, Lusahunga-Kagera: CS-DiD Event Study for Log Consumption (Band-Matched Control)



Notes: As Figure 22, but for the Lusahunga-Kagera upgrade.

H Night Lights by Project

This appendix reports the night-lights counterparts of the by-project exhibits in Section 6.5. We place them here, rather than alongside the consumption and population results, because the sample will not bear the weight of a by-project analysis.


Log luminosity is undefined for any cell recording no lights, and in this setting that removes most of the treated sample: of the 165 treated grid cells at the 25 km buffer, only 21 ever register positive luminosity. Split by project, that is 6 cells for Singida–Shelui and 15 for Lusahunga–Kagera, contributing 61 and 38 cell-year observations respectively. Estimates from six clusters should not be read as evidence for anything, and the standard errors reported by the estimator—which assume many clusters—will substantially overstate the precision available.

Figure 24 makes the problem visible. Pre-treatment coefficients for Singida–Shelui swing between -0.13 and $+0.11$ with confidence intervals several times the width of those in the consumption figures, so the parallel-trends premise is difficult to assess in the pre-period. Table 18 repeats the like-for-like comparison of Table 6 for night lights, and should be read subject to the same caveat: the treated-location counts are reported in the table precisely so that the reader can weigh them.

Figure 24: Event Study by Project: Log Night Lights (25km Buffer)

(a) Singida–Shelui (completed 2009; 6 treated cells)

(b) Lusahunga–Kagera (completed 2013; 15 treated cells)



Notes: As Figure 7, for log night lights. Within each panel the upper plot uses the main control group and the lower plot the alternative control group. Both projects share a common vertical scale, which is roughly three times wider than that of the consumption figures. The number of treated grid cells contributing to each panel is given in the subcaptions: log luminosity is undefined for cells recording no lights, which removes 144 of the 165 treated cells. Standard errors are clustered at the grid level and should be treated as optimistic given the number of clusters.

Table 18: Like-for-Like Comparison of the Two Upgrades, Night Lights (25km Buffer)

	Ln(Night Lights)		
	(1)	(2)	(3)
ATT	0.049 (0.035)	-0.108*** (0.039)	0.078 (0.050)
Observations	374	224	482
Project	Project A	Project B	Project A
Post-treatment window	e = 0 to 4	e = 0 to 4	e = 0 to 8
Last year in sample	2014	2018	2018
Treated locations (non-missing outcome)	3	15	6
Dep Variable Mean (Control, pre-completion)	2.188	2.191	2.188
Estimator	CS-DID	CS-DID	CS-DID

Notes: As Table 6, for log night lights. Columns (1) and (2) restrict both projects to the first five post-treatment years ($e = 0$ to 4); column (3) reports Singida–Shelui over its full nine-year window. The “Treated locations” row counts grid cells with a non-missing outcome, and is the reason these estimates should not be relied upon. Standard errors (in parentheses) are clustered at the grid level. Stars indicate statistical significance: $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

I Pooled Robustness Exhibits

This appendix collects the pooled exhibits that the main text no longer foregrounds: the 50 km buffer versions of the headline results, the pooled versions of the robustness exercises that Section 7 presents separately by project, and the pooled versions of the synthetic difference-in-differences and heterogeneity analyses. As Section 6.5 shows, pooled estimates average two opposing effects for the first five post-treatment years and then reproduce the Singida–Shelui path alone, so all of these exhibits should be read as descriptive of the aggregate rather than of either upgrade.

Headline Results at the 50km Buffer

Figures 25–26 and Table 19 repeat the exhibits of Section 6 using the 50 km buffer to assign treatment and control status. Results for all three outcomes are very similar to—though largely muted relative to—the 25 km results, as one would expect if the effects of the road improvements are largely localised. Table 20 reports the event-time disaggregation of the consumption effect for both buffers.

Figure 25: Pooled Event Study: Log Consumption (50km Buffer)




Figure 26: Pooled Event Study: Secondary Outcomes (50km Buffer)




Table 19: Pooled CS-DiD Results by Outcome (50km Buffer)

	Ln(Consumption)		Ln(Night Lights)		Ln(Population)	
	(1)	(2)	(3)	(4)	(5)	(6)
ATT	0.005 (0.004)	0.014*** (0.004)	-0.012 (0.041)	-0.067 (0.064)	-0.046*** (0.008)	-0.049*** (0.009)
Observations	55696	24832	1099	449	66139	29488
Control Group	Full	Alt	Full	Alt	Full	Alt
Dep Variable Mean (Control Group pre 2009)	6.206	6.215	2.241	2.185	6.887	6.974
Treated locations	381	381	40	40	381	381
Estimator	CS-DID	CS-DID	CS-DID	CS-DID	CS-DID	CS-DID

Notes: As Table 5, using the 50 km buffer.


Table 20: Pooled CS-DiD Dynamic Treatment Effects: Log Consumption (25km and 50km Buffers)

	25km Buffer		50km Buffer	
	(1) Full	(2) Alt	(3) Full	(4) Alt
Year -5	-0.014*** (0.003)	-0.006 (0.004)	-0.013*** (0.002)	-0.006** (0.002)
Year -4	-0.023*** (0.002)	-0.019*** (0.002)	-0.021*** (0.001)	-0.020*** (0.001)
Year -3	0.013*** (0.004)	0.008** (0.004)	0.011*** (0.003)	0.007*** (0.002)
Year -2	0.010*** (0.002)	0.020*** (0.003)	0.012*** (0.002)	0.020*** (0.002)
Year -1	-0.010*** (0.004)	-0.013*** (0.004)	-0.009*** (0.002)	-0.011*** (0.002)
Year 0 (treatment)	-0.007 (0.005)	0.012** (0.005)	-0.006** (0.003)	0.009*** (0.003)
Year +1	-0.008 (0.006)	0.002 (0.006)	0.002 (0.004)	0.010** (0.004)
Year +2	-0.014** (0.006)	-0.005 (0.007)	-0.013*** (0.004)	-0.009** (0.004)
Year +3	-0.020*** (0.007)	-0.007 (0.007)	-0.018*** (0.004)	-0.009** (0.004)
Year +4	0.013 (0.011)	0.022* (0.012)	0.001 (0.006)	0.010 (0.007)
Year +5	0.065*** (0.019)	0.091*** (0.019)	-0.002 (0.010)	0.019* (0.010)
Year +6	0.082*** (0.018)	0.097*** (0.019)	0.028*** (0.010)	0.033*** (0.011)
Year +7	0.129*** (0.018)	0.153*** (0.018)	0.067*** (0.011)	0.084*** (0.011)
Year +8	0.144*** (0.019)	0.150*** (0.020)	0.056*** (0.013)	0.060*** (0.013)
Pre-treatment average	-0.003*** (0.001)	-0.003*** (0.001)	-0.002*** (0.000)	-0.001*** (0.000)
Post-treatment average	0.043*** (0.009)	0.057*** (0.009)	0.013** (0.005)	0.023*** (0.006)
Observations	25024	10464	55696	24832
Control Group	Full	Alt	Full	Alt
Estimator	CS-DID	CS-DID	CS-DID	CS-DID

Notes: Standard errors (in parentheses) are clustered at the grid level. Each row labelled 'Year $-k$ ' to 'Year $+k$ ' reports the event-time average treatment effect $ATT(e)$, estimated by aggregating group-time $ATT(g, t)$ estimates across cohorts using cohort-size weights for that event time. The rows labelled 'Pre-treatment average' and 'Post-treatment average' report the *unweighted* mean of all pre- and post-treatment $ATT(e)$ estimates respectively. Columns (1)–(2) use the 25 km buffer and columns (3)–(4) the 50 km buffer; within each, the first column uses the full control group and the second the alternative. The same pattern is visible in the event-study figures reported in Section 6. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


Pooled Robustness Exercises

Figure 27: Pooled: Robustness to Distance Threshold, CS-DiD Event Study for Log Consumption




Notes: CS-DiD event study for log consumption pooling both projects, using the main control group, across six buffer thresholds. Compare Figures 10 and 11.

Figure 28: Pooled: Robustness to Spillover Exclusion Zone, CS-DiD Event Study for Log Consumption



Notes: As above, for the 100 km spillover-exclusion check. Compare Figures 12 and 29.

Figure 29: Robustness to Spillover Exclusion Zone by Project: CS-DiD Event Study for Log Consumption (50km Buffer)



Notes: As Figure 12, using the 50 km treatment buffer.

Pooled Synthetic Difference-in-Differences

Table 21 reports the SDID estimates pooling the two adoption cohorts. For log consumption the pooled ATT is 0.017 (s.e. 0.006) with the full control group and 0.030 (s.e. 0.007) with the alternative control group—close to the pooled CS-DiD estimates in Table 5. As with the CS-DiD results, this average conceals the divergence between the two upgrades documented in Section 7.5.

Table 21: Pooled Synthetic Difference-in-Differences: Log Consumption and Log Population (25km Buffer)


	Ln(Consumption)		Ln(Population)	
	(1)	(2)	(3)	(4)
treatment x post	0.017*** (0.006)	0.030*** (0.007)	-0.032** (0.015)	-0.046** (0.019)
Observations	25024	10464	25024	10464
Control Group	Full	Alt	Full	Alt
Dep Variable Mean (Control Group pre 2009)	6.217	6.239	7.144	7.198
Treated locations (balanced panel)	165	165	165	165
Control locations (balanced panel)	1399	489	1399	489
Inference (vce)	placebo	placebo	placebo	placebo
Estimator	SDID	SDID	SDID	SDID

Notes: Synthetic difference-in-differences estimates (Arkhangelsky et al., 2021), implemented following Clarke et al. (2023), pooling the 2010 and 2014 adoption cohorts. The sample is restricted to grid cells with a complete consumption and population series over 2003–2018; the balanced treated and control counts are reported in the table. Columns (1) and (3) use the full control group; columns (2) and (4) use the alternative control group. Standard errors (in parentheses) are computed by the placebo method with 100 replications. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Pooled Heterogeneity by Initial Economic Activity

Figures 30 and 31 show the high/low event studies pooling the two projects. The pooled split is dominated by the composition problem described in Section 6.5: the ‘low’ panel mixes Singida–Shelui cells that gain with Lusahunga–Kagera cells that lose. The project-specific versions in Section 6.6 are the informative ones.

Figure 30: Pooled Event Study: Heterogeneity by Initial Total Consumption (25km Buffer)




Notes: CS-DiD event study for log per capita consumption, split by whether the grid cell’s 2003 total consumption (per capita consumption $\times$ population) lies above (‘High’, left panel) or below (‘Low’, right panel) the within-cohort median. Shaded bands show 95% confidence intervals. The horizontal axis shows years relative to project completion; the vertical axis shows log-point deviations from the control-group trend.

J Exposure to the Subsequent Isaka–Lusahunga Works

The Lusahunga–Kagera upgrade (Project B) shares its Lusahunga endpoint with a third corridor project, the Isaka–Lusahunga upgrade, which was completed in 2020. That project is not treated as an intervention in this paper because it completes after our outcome series ends, but it was plausibly under construction during Project B’s post-treatment window, and construction activity typically depresses local economic activity in the short run—the direction of the effect we estimate for Project B.

Figure 31: Pooled Event Study: Heterogeneity by Initial Total Consumption (50km Buffer)



Notes: As Figure 30 but using the 50 km buffer for treatment assignment.

The overlap is material. The two projects' road segments touch, 19 of Project B's 99 treated grid cells lie within 25 km of the Isaka–Lusahunga segments, and 34 lie within 50 km. No Project A cell lies within 127 km, and no control cell within 50 km—the latter because Isaka–Lusahunga is recorded as completed in the Global Construction database, so the 50 km exclusion rule applied to the control group already removes its vicinity. The exposure is therefore specific to Project B's treated group.

Table 22 re-estimates Project B dropping treated cells within 25 km and then 50 km of the Isaka–Lusahunga segments. The negative consumption effect attenuates by roughly forty percent, from -2.8 to -1.7 percent, and weakens in significance, but does not vanish. The negative population effect does not survive: it falls from -3.9 percent to -1.5 percent and becomes statistically indistinguishable from zero. Some of the loss of significance is mechanical, since the exclusions remove up to a third of the treated cells. We read this as indicating that exposure to the subsequent works may explain part—but not all—of Project B's estimated consumption decline, and possibly all of its estimated population decline.

Table 22: Project B Robustness: Excluding Cells Near the Isaka–Lusahunga Works (25km Buffer)

	Ln(Consumption)				Ln(Population)	
	(1)	(2)	(3)	(4)	(5)	(6)
ATT	-0.028*** (0.008)	-0.020** (0.009)	-0.017* (0.010)	-0.039** (0.018)	-0.025 (0.020)	-0.015 (0.023)
Observations	23968	23664	23424	28462	28101	27816
Project	Project B	Project B	Project B	Project B	Project B	Project B
Treated cells dropped near Project C	None	25 km	50 km	None	25 km	50 km
Treated locations	99	80	65	99	80	65
Estimator	CS-DID	CS-DID	CS-DID	CS-DID	CS-DID	CS-DID

Notes: CS-DiD simple aggregate ATTs against the full never-treated control group. Columns progressively drop Lusahunga–Kagera treated cells lying within 25 km and then 50 km of the Isaka–Lusahunga road segments. Standard errors clustered at the grid level. Stars indicate statistical significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.